

GaussDB

Developer Guide(Centralized_V2.0-2.x)

Issue	01
Date	2025-08-20



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

- 1 Database System Overview..... 1**
 - 1.1 Database Logical Architecture..... 1
 - 1.2 Query Request Handling Process..... 3
 - 1.3 Managing Transactions..... 3
 - 1.4 Concepts..... 5
- 2 Database Security Management.....6**
 - 2.1 Users and Permissions..... 6
 - 2.1.1 Default Permission Mechanism..... 6
 - 2.1.2 Administrator..... 7
 - 2.1.3 Separation of Duties..... 9
 - 2.1.4 Users..... 11
 - 2.1.5 Roles..... 12
 - 2.1.6 Schemas..... 14
 - 2.1.7 Setting User Permissions..... 15
 - 2.1.8 Row-Level Security Policy..... 16
 - 2.2 Database Audit..... 18
- 3 Database Quick Start..... 20**
 - 3.1 Connecting to a Database..... 20
 - 3.1.1 Using gsql to Connect to a Database..... 20
 - 3.1.2 APIs..... 22
 - 3.2 Operating a Database..... 22
 - 3.2.1 Creating a Database Account..... 22
 - 3.2.2 Creating and Managing Databases..... 23
 - 3.2.3 Creating and Managing Tablespaces..... 25
 - 3.2.4 Creating and Managing Tables..... 28
 - 3.2.4.1 Creating Tables..... 28
 - 3.2.4.2 Inserting Data to Tables..... 28
 - 3.2.4.3 Updating Data in a Table..... 30
 - 3.2.4.4 Viewing Data..... 31
 - 3.2.4.5 Deleting Data from a Table..... 32
 - 3.2.5 Querying a System Catalog..... 32
 - 3.2.6 Other Operations..... 34

3.2.6.1 Creating and Managing Schemas.....	34
3.2.6.2 Creating and Managing Partitioned Tables.....	37
3.2.6.3 Creating and Managing Indexes.....	41
3.2.6.4 Creating and Managing Views.....	45
3.2.6.5 Creating and Managing Sequences.....	46
3.2.6.6 Creating and Managing Scheduled Jobs.....	47
4 Development and Design Proposal.....	51
4.1 Overview.....	51
4.2 Database Design Specifications.....	52
4.2.1 General Specifications.....	52
4.2.2 Deployment Specifications.....	54
4.2.3 Database Object Naming Specifications.....	54
4.2.4 Database and Schema Design Specifications.....	56
4.2.5 Permission Design Specifications.....	58
4.2.6 Character Set Design Specifications.....	59
4.2.7 Table Design Specifications.....	59
4.2.8 Column Design Specifications.....	61
4.2.9 Index Design Specifications.....	65
4.2.10 Function/Stored Procedure Design Specifications.....	66
4.2.11 Constraint Design.....	66
4.2.12 View and Joined Table Design.....	67
4.3 Database Design Specifications.....	67
4.3.1 GUC Parameter Programming Specifications.....	67
4.3.2 Object Access Programming Specifications.....	68
4.3.3 WHERE.....	68
4.3.4 SELECT.....	70
4.3.5 INSERT.....	70
4.3.6 UPDATE.....	71
4.3.7 DELETE.....	71
4.3.8 Join Query.....	72
4.3.9 Subquery.....	72
4.3.10 Transaction.....	73
4.3.11 SQL Compilation.....	74
4.4 Client Programming Specifications.....	78
4.4.1 JDBC.....	78
5 Application Development Guide.....	81
5.1 GaussDB Application Development Guide.....	81
5.2 Development Specifications.....	84
5.3 Development Based on JDBC.....	85
5.3.1 Development Process.....	85
5.3.2 Development Process.....	85
5.3.2.1 Obtaining the Driver JAR Package and Configuring JDK 1.8.....	86

5.3.2.2 Connecting to a Database.....	87
5.3.2.2.1 Connection Methods.....	87
5.3.2.2.2 Connection Parameter Reference.....	89
5.3.2.2.3 Connecting to a Database in Non-Encrypted Mode.....	110
5.3.2.2.4 Connecting to a Database in SSL Mode.....	113
5.3.2.2.5 Connecting to a Database in UDS Mode.....	115
5.3.2.3 Running SQL Statements.....	116
5.3.2.4 Processing Data in a Result Set.....	121
5.3.2.5 Closing a Database Connection.....	124
5.3.3 Typical Application Development Examples.....	124
5.3.3.1 Configuring Connection Parameters in Different Scenarios.....	124
5.3.3.2 Creating and Calling a Stored Procedure.....	126
5.3.3.3 Obtaining the Return Value of a Function.....	128
5.3.3.4 Batch Query.....	130
5.3.3.5 SQL Retry at an Application Layer.....	132
5.3.3.6 Importing and Exporting Data Through Local Files.....	135
5.3.3.7 Migrating Data from a MySQL Database.....	137
5.3.3.8 Logical Replication.....	139
5.3.4 JDBC API Reference.....	145
5.3.4.1 java.sql.Connection.....	145
5.3.4.2 java.sql.CallableStatement.....	147
5.3.4.3 java.sql.DatabaseMetaData.....	149
5.3.4.4 java.sql.Driver.....	158
5.3.4.5 java.sql.PreparedStatement.....	159
5.3.4.6 java.sql.ResultSet.....	163
5.3.4.7 java.sql.ResultSetMetaData.....	170
5.3.4.8 java.sql.Statement.....	171
5.3.4.9 javax.sql.ConnectionPoolDataSource.....	173
5.3.4.10 javax.sql.DataSource.....	174
5.3.4.11 javax.sql.PooledConnection.....	174
5.3.4.12 javax.naming.Context.....	175
5.3.4.13 javax.naming.spi.InitialContextFactory.....	175
5.3.4.14 CopyManager.....	175
5.3.4.15 PGReplicationConnection.....	177
5.3.4.16 PGReplicationStream.....	177
5.3.4.17 ChainedStreamBuilder.....	179
5.3.4.18 ChainedCommonStreamBuilder.....	179
5.4 Development Based on ODBC.....	180
5.4.1 Development Process.....	180
5.4.2 Development Procedure.....	181
5.4.2.1 Obtaining a Source Code Package, ODBC Packages, and Dependent Libraries.....	181
5.4.2.2 Connecting to a Database.....	182

5.4.2.2.1 Configuring a Data Source in the Linux OS.....	182
5.4.2.2.2 Configuring a Data Source in the Windows OS.....	191
5.4.2.2.3 APIs for Connecting to a Database.....	195
5.4.2.3 Executing SQL Statements.....	196
5.4.2.4 Processing Data in a Result Set.....	197
5.4.2.5 Closing a Connection.....	198
5.4.3 Typical Application Development Examples.....	198
5.4.3.1 Typical Application Scenarios and Configurations.....	198
5.4.3.2 Obtaining and Processing Data in a Database.....	199
5.4.3.3 Batch Binding.....	201
5.4.3.4 High-Performance Binding.....	207
5.4.4 ODBC Interface Reference.....	214
5.4.4.1 SQLAllocEnv.....	214
5.4.4.2 SQLAllocConnect.....	214
5.4.4.3 SQLAllocHandle.....	214
5.4.4.4 SQLAllocStmt.....	216
5.4.4.5 SQLBindCol.....	216
5.4.4.6 SQLBindParameter.....	217
5.4.4.7 SQLColAttribute.....	218
5.4.4.8 SQLConnect.....	220
5.4.4.9 SQLDisconnect.....	221
5.4.4.10 SQLExecDirect.....	222
5.4.4.11 SQLExecute.....	223
5.4.4.12 SQLFetch.....	224
5.4.4.13 SQLFreeStmt.....	225
5.4.4.14 SQLFreeConnect.....	225
5.4.4.15 SQLFreeHandle.....	225
5.4.4.16 SQLFreeEnv.....	226
5.4.4.17 SQLPrepare.....	226
5.4.4.18 SQLGetData.....	227
5.4.4.19 SQLGetDiagRec.....	228
5.4.4.20 SQLSetConnectAttr.....	231
5.4.4.21 SQLSetEnvAttr.....	232
5.4.4.22 SQLSetStmtAttr.....	233
5.5 Development Based on libpq.....	234
5.5.1 Development Process.....	234
5.5.2 Development Procedure.....	234
5.5.2.1 Obtaining a Driver Package, a Dependency Library, and Header Files.....	234
5.5.2.2 Connecting to a Database.....	235
5.5.2.3 Executing SQL Statements.....	236
5.5.2.4 Processing Data in a Result Set.....	236
5.5.2.5 Closing a Connection.....	237

5.5.3 Typical Application Development Examples.....	237
5.5.3.1 Establishing a Database Connection, Executing SQL Statements, and Returning Results.....	237
5.5.3.2 Executing Prepared Statements.....	240
5.5.3.3 Binding Parameters and Returning a Binary Result.....	241
5.5.4 libpq API Reference.....	244
5.5.4.1 Database Connection Control Functions.....	244
5.5.4.1.1 PQconnectdbParams.....	245
5.5.4.1.2 PQconnectdb.....	245
5.5.4.1.3 PQconninfoParse.....	246
5.5.4.1.4 PQconnectStart.....	247
5.5.4.1.5 PQerrorMessage.....	247
5.5.4.1.6 PQsetdbLogin.....	248
5.5.4.1.7 PQfinish.....	249
5.5.4.1.8 PQreset.....	249
5.5.4.1.9 PQstatus.....	250
5.5.4.2 Database Statement Execution Functions.....	251
5.5.4.2.1 PQclear.....	251
5.5.4.2.2 PQexec.....	252
5.5.4.2.3 PQexecParams.....	253
5.5.4.2.4 PQexecParamsBatch.....	253
5.5.4.2.5 PQexecPrepared.....	254
5.5.4.2.6 PQexecPreparedBatch.....	255
5.5.4.2.7 PQprepare.....	256
5.5.4.3 Asynchronous Command Processing Functions.....	257
5.5.4.3.1 PQsendQuery.....	258
5.5.4.3.2 PQsendQueryParams.....	258
5.5.4.3.3 PQsendPrepare.....	259
5.5.4.3.4 PQsendQueryPrepared.....	260
5.5.4.3.5 PQflush.....	261
5.5.4.4 Functions for Canceling Query Processing.....	262
5.5.4.4.1 PQgetCancel.....	262
5.5.4.4.2 PQfreeCancel.....	263
5.5.4.4.3 PQcancel.....	263
5.5.4.5 Functions for Processing Database Result.....	264
5.5.4.5.1 PQgetvalue.....	264
5.5.4.5.2 PQfname.....	265
5.5.4.5.3 PQnfields.....	265
5.5.4.5.4 PQntuples.....	266
5.5.4.5.5 PQresultStatus.....	266
5.6 Psycopg-based Development.....	267
5.6.1 Development Process.....	269
5.6.2 Psycopg Package.....	269

5.6.3 Example: Common Operations.....	272
5.6.4 Psycpg API Reference.....	274
5.6.4.1 psycpg2.connect().....	274
5.6.4.2 connection.cursor().....	276
5.6.4.3 cursor.execute(query,vars_list).....	276
5.6.4.4 cursor.executemany(query,vars_list).....	277
5.6.4.5 connection.commit().....	278
5.6.4.6 connection.rollback().....	278
5.6.4.7 cursor.fetchone().....	279
5.6.4.8 cursor.fetchall().....	279
5.6.4.9 cursor.close().....	280
5.6.4.10 connection.close().....	280
5.7 Development Based on the Go Driver.....	281
5.7.1 Go Driver Package, Environment Class, and Driver Class.....	281
5.7.2 Development Process.....	282
5.7.3 Connecting to the Database.....	284
5.7.4 Connecting to the Database (Using SSL).....	291
5.7.5 Go API Reference.....	292
5.7.5.1 sql.Open.....	292
5.7.5.2 type DB.....	293
5.7.5.3 type Stmt.....	295
5.7.5.4 type Tx.....	296
5.7.5.5 type Rows.....	298
5.7.5.6 type Row.....	299
5.7.5.7 type ColumnType.....	299
5.7.5.8 type Result.....	300
5.8 Appendix.....	300
5.8.1 JDBC.....	300
5.8.1.1 Data Type Mapping.....	300
5.8.1.2 Log Management.....	301
5.8.1.3 Troubleshooting.....	306
5.8.1.3.1 Incorrect batchSize Settings.....	306
5.8.1.3.2 Error Is Reported or Connection Is Blocked in SSL Mode.....	306
5.8.2 libpq.....	307
5.8.2.1 Connection Parameters.....	308
5.8.3 Parameters Related to Log Output.....	313
6 SQL Optimization.....	318
6.1 Query Execution Process.....	318
6.2 Introduction to the SQL Execution Plan.....	321
6.2.1 Overview.....	321
6.2.2 Description.....	322
6.3 Tuning Process.....	329

6.4 Updating Statistics.....	330
6.5 Reviewing and Modifying a Table Definition.....	330
6.5.1 Overview.....	330
6.5.2 Selecting a Storage Model.....	331
6.5.3 Using Partitioned Tables.....	331
6.5.4 Selecting a Data Type.....	332
6.6 Typical SQL Optimization Methods.....	332
6.6.1 Optimizing SQL Self-Diagnosis.....	332
6.6.2 Optimizing Subqueries.....	333
6.6.3 Optimizing Statistics.....	341
6.6.4 Optimizing Operators.....	343
6.7 Experience in Rewriting SQL Statements.....	344
6.8 Configuring Key Parameters for SQL Optimization.....	346
6.9 Hint-based Tuning.....	348
6.9.1 Plan Hint Optimization.....	348
6.9.2 Join Order Hints.....	353
6.9.3 Join Operation Hints.....	354
6.9.4 Rows Hints.....	355
6.9.5 Scan Operation Hints.....	356
6.9.6 Sublink Name Hints.....	357
6.9.7 Hint Errors, Conflicts, and Other Warnings.....	358
6.9.8 Optimizer GUC Parameter Hints.....	359
6.9.9 Hint for Selecting the Custom Plan or Generic Plan.....	360
6.9.10 Hint Specifying Not to Expand Subqueries.....	361
6.9.11 Hint Specifying Not to Use Global Plan Cache.....	362
6.9.12 Hint of Parameterized Paths at the Same Level.....	362
6.9.13 Hint for Materializing a Sub-plan Result.....	363
6.10 Tuning with SQL PATCH.....	364
6.11 Optimization Cases.....	367
6.11.1 Case: Modifying the GUC Parameter rewrite_rule.....	367
6.11.2 Case: Creating an Appropriate Index.....	370
6.11.3 Case: Adding NOT NULL for the JOIN Columns.....	371
6.11.4 Case: Modifying a Partitioned Table.....	372
6.11.5 Case: Rewriting SQL Statements to Eliminate Subqueries.....	372
6.11.6 Case: Rewriting SQL Statements and Deleting in-clause.....	373
7 SQL Reference.....	376
7.1 GaussDB SQL.....	376
7.2 Keywords.....	377
7.3 Data Type.....	412
7.3.1 Numeric Types.....	413
7.3.2 Monetary Types.....	419
7.3.3 Boolean Types.....	420

7.3.4 Character Types.....	421
7.3.5 Binary Types.....	423
7.3.6 Date/Time Types.....	425
7.3.7 Geometric.....	434
7.3.8 Network Address Types.....	437
7.3.9 Bit String Types.....	440
7.3.10 UUID.....	441
7.3.11 JSON/JSONB Types.....	442
7.3.12 HLL.....	446
7.3.13 Range Types.....	450
7.3.14 OID Types.....	455
7.3.15 Pseudo-Types.....	457
7.3.16 ACLItem.....	458
7.4 Constant and Macro.....	459
7.5 Functions and Operators.....	460
7.5.1 Logical Operators.....	460
7.5.2 Comparison Operators.....	461
7.5.3 Character Processing Functions and Operators.....	462
7.5.4 Binary String Functions and Operators.....	480
7.5.5 Bit String Functions and Operators.....	483
7.5.6 Pattern Matching Operators.....	486
7.5.7 Mathematical Functions and Operators.....	490
7.5.8 Date and Time Processing Functions and Operators.....	505
7.5.9 Type Conversion Functions.....	528
7.5.10 Geometric Functions and Operators.....	545
7.5.11 Network Address Functions and Operators.....	556
7.5.12 JSON/JSONB Functions and Operators.....	561
7.5.13 HLL Functions and Operators.....	573
7.5.14 SEQUENCE Functions.....	586
7.5.15 Array Functions and Operators.....	588
7.5.16 Range Functions and Operators.....	596
7.5.17 Aggregate Functions.....	601
7.5.18 Window Functions.....	614
7.5.19 Security Functions.....	619
7.5.20 Set Returning Functions.....	624
7.5.21 Conditional Expression Functions.....	626
7.5.22 System Information Functions.....	628
7.5.23 System Administration Functions.....	657
7.5.23.1 Configuration Settings Functions.....	657
7.5.23.2 Universal File Access Functions.....	657
7.5.23.3 Server Signal Functions.....	659
7.5.23.4 Backup and Restoration Control Functions.....	661

7.5.23.5 DR Control Functions for Dual Database Instances.....	667
7.5.23.6 DR Query Functions for Dual-Database Instances.....	669
7.5.23.7 Snapshot Synchronization Functions.....	673
7.5.23.8 Database Object Functions.....	673
7.5.23.9 Advisory Lock Functions.....	679
7.5.23.10 Logical Replication Functions.....	682
7.5.23.11 Segment-Page Storage Functions.....	692
7.5.23.12 Other Functions.....	695
7.5.23.13 Undo System Functions.....	716
7.5.24 Statistics Information Functions.....	730
7.5.25 Trigger Functions.....	774
7.5.26 Hash Function.....	775
7.5.27 Prompt Message Function.....	778
7.5.28 Global Temporary Table Functions.....	778
7.5.29 Fault Injection System Function.....	781
7.5.30 AI Feature Functions.....	781
7.5.31 Dynamic Data Masking Functions.....	783
7.5.32 Hierarchical Recursion Query Functions.....	784
7.5.33 Internal Functions.....	785
7.5.34 Global SysCache Feature Functions.....	791
7.5.35 Data Damage Detection and Repair Functions.....	793
7.5.36 Other System Functions.....	799
7.5.37 Obsolete Functions.....	821
7.6 Expressions.....	822
7.6.1 Simple Expressions.....	822
7.6.2 Condition Expressions.....	824
7.6.3 Subquery Expressions.....	828
7.6.4 Array Expressions.....	831
7.6.5 Row Expressions.....	832
7.7 Type Conversion.....	834
7.7.1 Overview.....	834
7.7.2 Operators.....	836
7.7.3 Functions.....	838
7.7.4 Value Storage.....	840
7.7.5 UNION, CASE, and Related Constructs.....	841
7.8 System Operation.....	845
7.9 Controlling Transactions.....	846
7.10 DDL Syntax Overview.....	847
7.11 DML Syntax Overview.....	852
7.12 DCL Syntax Overview.....	853
7.13 SQL Syntax.....	854
7.13.1 ABORT.....	854

7.13.2 ALTER AGGREGATE.....	855
7.13.3 ALTER AUDIT POLICY.....	856
7.13.4 ALTER DATABASE.....	858
7.13.5 ALTER DEFAULT PRIVILEGES.....	860
7.13.6 ALTER DIRECTORY.....	862
7.13.7 ALTER FUNCTION.....	863
7.13.8 ALTER GLOBAL CONFIGURATION.....	866
7.13.9 ALTER GROUP.....	867
7.13.10 ALTER INDEX.....	868
7.13.11 ALTER LANGUAGE.....	870
7.13.12 ALTER MASKING POLICY.....	870
7.13.13 ALTER MATERIALIZED VIEW.....	872
7.13.14 ALTER OPERATOR.....	873
7.13.15 ALTER PACKAGE.....	874
7.13.16 ALTER PROCEDURE.....	875
7.13.17 ALTER RESOURCE LABEL.....	878
7.13.18 ALTER ROLE.....	879
7.13.19 ALTER ROW LEVEL SECURITY POLICY.....	881
7.13.20 ALTER SCHEMA.....	883
7.13.21 ALTER SEQUENCE.....	884
7.13.22 ALTER SERVER.....	885
7.13.23 ALTER SESSION.....	887
7.13.24 ALTER SYNONYM.....	889
7.13.25 ALTER SYSTEM KILL SESSION.....	890
7.13.26 ALTER TABLE.....	890
7.13.27 ALTER TABLE PARTITION.....	903
7.13.28 ALTER TABLE SUBPARTITION.....	909
7.13.29 ALTER TABLESPACE.....	912
7.13.30 ALTER TRIGGER.....	913
7.13.31 ALTER TYPE.....	914
7.13.32 ALTER USER.....	916
7.13.33 ALTER USER MAPPING.....	918
7.13.34 ALTER VIEW.....	919
7.13.35 ANALYZE ANALYSE.....	921
7.13.36 BEGIN.....	925
7.13.37 CALL.....	926
7.13.38 CHECKPOINT.....	927
7.13.39 CLEAN CONNECTION.....	928
7.13.40 CLOSE.....	929
7.13.41 CLUSTER.....	930
7.13.42 COMMENT.....	932
7.13.43 COMMIT END.....	935

7.13.44 COMMIT PREPARED.....	936
7.13.45 COPY.....	937
7.13.46 CREATE AGGREGATE.....	951
7.13.47 CREATE AUDIT POLICY.....	953
7.13.48 CREATE CAST.....	955
7.13.49 CREATE CONVERSION.....	957
7.13.50 CREATE DATABASE.....	958
7.13.51 CREATE DIRECTORY.....	966
7.13.52 CREATE FUNCTION.....	967
7.13.53 CREATE GROUP.....	975
7.13.54 CREATE INCREMENTAL MATERIALIZED VIEW.....	976
7.13.55 CREATE INDEX.....	978
7.13.56 CREATE LANGUAGE.....	986
7.13.57 CREATE MASKING POLICY.....	986
7.13.58 CREATE MATERIALIZED VIEW.....	988
7.13.59 CREATE OPERATOR.....	989
7.13.60 CREATE OPERATOR CLASS.....	991
7.13.61 CREATE PACKAGE.....	993
7.13.62 CREATE PROCEDURE.....	995
7.13.63 CREATE RESOURCE LABEL.....	998
7.13.64 CREATE ROLE.....	999
7.13.65 CREATE ROW LEVEL SECURITY POLICY.....	1004
7.13.66 CREATE RULE.....	1008
7.13.67 CREATE SCHEMA.....	1010
7.13.68 CREATE SEQUENCE.....	1011
7.13.69 CREATE SERVER.....	1014
7.13.70 CREATE SYNONYM.....	1015
7.13.71 CREATE TABLE.....	1017
7.13.72 CREATE TABLE AS.....	1036
7.13.73 CREATE TABLE PARTITION.....	1039
7.13.74 CREATE TABLESPACE.....	1058
7.13.75 CREATE TABLE SUBPARTITION.....	1060
7.13.76 CREATE TRIGGER.....	1079
7.13.77 CREATE TYPE.....	1084
7.13.78 CREATE USER.....	1091
7.13.79 CREATE USER MAPPING.....	1093
7.13.80 CREATE VIEW.....	1094
7.13.81 CREATE WEAK PASSWORD DICTIONARY.....	1096
7.13.82 CURSOR.....	1097
7.13.83 DEALLOCATE.....	1098
7.13.84 DECLARE.....	1099
7.13.85 DELETE.....	1100

7.13.86 DO.....	1103
7.13.87 DROP AGGREGATE.....	1104
7.13.88 DROP AUDIT POLICY.....	1105
7.13.89 DROP CAST.....	1105
7.13.90 DROP DATABASE.....	1106
7.13.91 DROP DIRECTORY.....	1107
7.13.92 DROP FUNCTION.....	1108
7.13.93 DROP GLOBAL CONFIGURATION.....	1108
7.13.94 DROP GROUP.....	1109
7.13.95 DROP INDEX.....	1110
7.13.96 DROP LANGUAGE.....	1111
7.13.97 DROP MASKING POLICY.....	1111
7.13.98 DROP MATERIALIZED VIEW.....	1111
7.13.99 DROP OPERATOR.....	1112
7.13.100 DROP OWNED.....	1113
7.13.101 DROP PACKAGE.....	1114
7.13.102 DROP PROCEDURE.....	1114
7.13.103 DROP RESOURCE LABEL.....	1115
7.13.104 DROP ROLE.....	1115
7.13.105 DROP ROW LEVEL SECURITY POLICY.....	1116
7.13.106 DROP RULE.....	1117
7.13.107 DROP SCHEMA.....	1118
7.13.108 DROP SEQUENCE.....	1118
7.13.109 DROP SERVER.....	1119
7.13.110 DROP SYNONYM.....	1120
7.13.111 DROP TABLE.....	1121
7.13.112 DROP TABLESPACE.....	1122
7.13.113 DROP TRIGGER.....	1123
7.13.114 DROP TYPE.....	1123
7.13.115 DROP USER.....	1124
7.13.116 DROP USER MAPPING.....	1125
7.13.117 DROP VIEW.....	1126
7.13.118 DROP WEAK PASSWORD DICTIONARY.....	1127
7.13.119 EXECUTE.....	1127
7.13.120 EXPLAIN.....	1128
7.13.121 EXPLAIN PLAN.....	1133
7.13.122 FETCH.....	1134
7.13.123 GRANT.....	1138
7.13.124 INSERT.....	1147
7.13.125 LOCK.....	1152
7.13.126 MERGE INTO.....	1155
7.13.127 MOVE.....	1158

7.13.128 PREPARE.....	1159
7.13.129 PREPARE TRANSACTION.....	1160
7.13.130 PURGE.....	1161
7.13.131 REASSIGN OWNED.....	1163
7.13.132 REFRESH INCREMENTAL MATERIALIZED VIEW.....	1164
7.13.133 REFRESH MATERIALIZED VIEW.....	1164
7.13.134 REINDEX.....	1165
7.13.135 RELEASE SAVEPOINT.....	1168
7.13.136 RESET.....	1169
7.13.137 REVOKE.....	1170
7.13.138 ROLLBACK.....	1173
7.13.139 ROLLBACK PREPARED.....	1174
7.13.140 ROLLBACK TO SAVEPOINT.....	1175
7.13.141 SAVEPOINT.....	1176
7.13.142 SELECT.....	1178
7.13.143 SELECT INTO.....	1201
7.13.144 SET.....	1203
7.13.145 SET CONSTRAINTS.....	1205
7.13.146 SET ROLE.....	1206
7.13.147 SET SESSION AUTHORIZATION.....	1207
7.13.148 SET TRANSACTION.....	1208
7.13.149 SHOW.....	1209
7.13.150 SHUTDOWN.....	1210
7.13.151 START TRANSACTION.....	1211
7.13.152 TIMECAPSULE TABLE.....	1212
7.13.153 TRUNCATE.....	1215
7.13.154 UPDATE.....	1217
7.13.155 VACUUM.....	1220
7.13.156 VALUES.....	1224
7.14 Appendix.....	1225
7.14.1 Extended Functions.....	1225
7.14.2 Extended Syntax.....	1225
8 Best Practices.....	1227
8.1 Best Practices for Table Design.....	1227
8.1.1 Selecting a Storage Model.....	1227
8.1.2 Using Partitioned Tables.....	1227
8.1.3 Selecting a Data type.....	1228
8.2 Best Practices for SQL Queries.....	1228
9 User-defined Functions.....	1231
9.1 PL/SQL Functions.....	1231
10 Stored Procedure.....	1232

10.1 Stored Procedure.....	1232
10.2 Data Types.....	1232
10.3 Data Type Conversion.....	1232
10.4 Arrays, Sets, and Records.....	1234
10.4.1 Arrays.....	1234
10.4.2 Collections.....	1235
10.4.3 Records.....	1237
10.5 DECLARE Syntax.....	1239
10.5.1 Basic Structure.....	1239
10.5.2 Anonymous Blocks.....	1240
10.5.3 Subprogram.....	1241
10.6 Basic Statements.....	1241
10.6.1 Variable Definition Statements.....	1241
10.6.2 Assignment Statements.....	1243
10.6.3 Call Statements.....	1245
10.7 Dynamic Statements.....	1246
10.7.1 Executing Dynamic Query Statements.....	1246
10.7.2 Executing Dynamic Non-Query Statements.....	1248
10.7.3 Dynamically Calling Stored Procedures.....	1249
10.7.4 Dynamically Calling Anonymous Blocks.....	1250
10.8 Control Statements.....	1252
10.8.1 RETURN Statements.....	1252
10.8.1.1 RETURN.....	1252
10.8.1.2 RETURN NEXT and RETURN QUERY.....	1253
10.8.2 Conditional Statements.....	1254
10.8.3 Loop Statements.....	1256
10.8.4 Branch Statements.....	1259
10.8.5 NULL Statements.....	1260
10.8.6 Error Trapping Statements.....	1260
10.8.7 GOTO Statements.....	1262
10.9 Transaction Management.....	1264
10.10 Other Statements.....	1272
10.10.1 Lock Operations.....	1272
10.10.2 Cursor Operations.....	1272
10.11 Cursors.....	1272
10.11.1 Overview.....	1272
10.11.2 Explicit Cursor.....	1273
10.11.3 Implicit Cursor.....	1278
10.11.4 Cursor Loop.....	1279
10.12 Advanced Packages.....	1280
10.12.1 Basic APIs.....	1280
10.12.1.1 PKG_SERVICE.....	1280

10.12.1.2 PKG_UTIL.....	1291
10.12.2 Secondary Encapsulation APIs (Recommended).....	1313
10.12.2.1 DBE_LOB.....	1314
10.12.2.2 DBE_RANDOM.....	1327
10.12.2.3 DBE_OUTPUT.....	1328
10.12.2.4 DBE_RAW.....	1330
10.12.2.5 DBE_TASK.....	1333
10.12.2.6 DBE_SCHEDULER.....	1345
10.12.2.7 DBE_SQL.....	1384
10.12.2.8 DBE_FILE.....	1413
10.12.2.9 DBE_UTILITY.....	1434
10.12.2.10 DBE_SESSION.....	1435
10.12.2.11 DBE_MATCH.....	1437
10.12.2.12 DBE_APPLICATION_INFO.....	1437
10.13 Retry Management.....	1438
10.14 Debugging.....	1439
10.15 Package.....	1442
11 Autonomous Transaction.....	1444
11.1 Stored Procedure Supporting Autonomous Transaction.....	1444
11.2 Anonymous Block Supporting Autonomous Transaction.....	1445
11.3 Function Supporting Autonomous Transaction.....	1445
11.4 Package Supporting Autonomous Transaction.....	1446
11.5 Restrictions.....	1447
12 System Catalogs and System Views.....	1451
12.1 Overview of System Catalogs and System Views.....	1451
12.2 System Catalogs.....	1451
12.2.1 GS_ASP.....	1452
12.2.2 GS_AUDITING_POLICY.....	1453
12.2.3 GS_AUDITING_POLICY_ACCESS.....	1454
12.2.4 GS_AUDITING_POLICY_FILTERS.....	1455
12.2.5 GS_AUDITING_POLICY_PRIVILEGES.....	1455
12.2.6 GS_DB_PRIVILEGE.....	1456
12.2.7 GS_GLOBAL_CONFIG.....	1456
12.2.8 GS_JOB_ARGUMENT.....	1457
12.2.9 GS_JOB_ATTRIBUTE.....	1457
12.2.10 GS_MASKING_POLICY.....	1458
12.2.11 GS_MASKING_POLICY_ACTIONS.....	1458
12.2.12 GS_MASKING_POLICY_FILTERS.....	1459
12.2.13 GS_MATVIEW.....	1460
12.2.14 GS_MATVIEW_DEPENDENCY.....	1460
12.2.15 GS_OPT_MODEL.....	1461
12.2.16 GS_PACKAGE.....	1463

12.2.17 GS_POLICY_LABEL.....	1464
12.2.18 GS_RECYCLEBIN.....	1464
12.2.19 GS_SQL_PATCH.....	1466
12.2.20 GS_TXN_SNAPSHOT.....	1467
12.2.21 GS_UID.....	1467
12.2.22 PG_AGGREGATE.....	1468
12.2.23 PG_AM.....	1469
12.2.24 PG_AMOP.....	1472
12.2.25 PG_AMPROC.....	1473
12.2.26 PG_APP_WORKLOADGROUP_MAPPING.....	1474
12.2.27 PG_ATTRDEF.....	1474
12.2.28 PG_ATTRIBUTE.....	1475
12.2.29 PG_AUTHID.....	1477
12.2.30 PG_AUTH_HISTORY.....	1479
12.2.31 PG_AUTH_MEMBERS.....	1480
12.2.32 PG_CAST.....	1480
12.2.33 PG_CLASS.....	1481
12.2.34 PG_COLLATION.....	1486
12.2.35 PG_CONSTRAINT.....	1486
12.2.36 PG_CONVERSION.....	1489
12.2.37 PG_DATABASE.....	1490
12.2.38 PG_DB_ROLE_SETTING.....	1491
12.2.39 PG_DEFAULT_ACL.....	1492
12.2.40 PG_DEPEND.....	1492
12.2.41 PG_DESCRIPTION.....	1493
12.2.42 PG_DIRECTORY.....	1494
12.2.43 PG_ENUM.....	1494
12.2.44 PG_FOREIGN_SERVER.....	1495
12.2.45 PG_HASHBUCKET.....	1496
12.2.46 PG_INDEX.....	1496
12.2.47 PG_INHERITS.....	1499
12.2.48 PG_JOB.....	1499
12.2.49 PG_JOB_PROC.....	1501
12.2.50 PG_LANGUAGE.....	1501
12.2.51 PG_LARGEOBJECT.....	1502
12.2.52 PG_LARGEOBJECT_METADATA.....	1503
12.2.53 PG_NAMESPACE.....	1503
12.2.54 PG_OBJECT.....	1504
12.2.55 PG_OPCLASS.....	1505
12.2.56 PG_OPERATOR.....	1506
12.2.57 PG_OPFAMILY.....	1507
12.2.58 PG_PARTITION.....	1508

12.2.59 PG_PLTEMPLATE.....	1510
12.2.60 PG_PROC.....	1511
12.2.61 PG_RANGE.....	1515
12.2.62 PG_REPLICATION_ORIGIN.....	1516
12.2.63 PG_RESOURCE_POOL.....	1516
12.2.64 PG_REWRITE.....	1517
12.2.65 PG_RLSPOLICY.....	1518
12.2.66 PG_SECLABEL.....	1519
12.2.67 PG_SHDEPEND.....	1519
12.2.68 PG_SHDESCRIPTION.....	1520
12.2.69 PG_SHSECLABEL.....	1521
12.2.70 PG_STATISTIC.....	1521
12.2.71 PG_STATISTIC_EXT.....	1523
12.2.72 PG_SYNONYM.....	1525
12.2.73 PG_TABLESPACE.....	1525
12.2.74 PG_TRIGGER.....	1526
12.2.75 PG_TS_CONFIG.....	1527
12.2.76 PG_TS_CONFIG_MAP.....	1527
12.2.77 PG_TS_DICT.....	1528
12.2.78 PG_TS_PARSER.....	1529
12.2.79 PG_TS_TEMPLATE.....	1529
12.2.80 PG_TYPE.....	1530
12.2.81 PG_USER_MAPPING.....	1533
12.2.82 PG_USER_STATUS.....	1534
12.2.83 PG_WORKLOAD_GROUP.....	1534
12.2.84 PGXC_CLASS.....	1535
12.2.85 PGXC_GROUP.....	1536
12.2.86 PGXC_NODE.....	1536
12.2.87 PGXC_SLICE.....	1538
12.2.88 PLAN_TABLE_DATA.....	1539
12.2.89 STATEMENT_HISTORY.....	1540
12.2.90 STREAMING_STREAM.....	1545
12.2.91 STREAMING_CONT_QUERY.....	1545
12.2.92 STREAMING_REAPER_STATUS.....	1546
12.3 System Views.....	1547
12.3.1 ADM_COL_COMMENTS.....	1547
12.3.2 ADM_CONS_COLUMNS.....	1547
12.3.3 ADM_CONSTRAINTS.....	1548
12.3.4 ADM_DATA_FILES.....	1549
12.3.5 ADM_HIST_SNAPSHOT.....	1549
12.3.6 ADM_HIST_SQL_PLAN.....	1549
12.3.7 ADM_HIST_SQLSTAT.....	1550

12.3.8 ADM_IND_COLUMNS.....	1551
12.3.9 ADM_IND_EXPRESSIONS.....	1551
12.3.10 ADM_IND_PARTITIONS.....	1551
12.3.11 ADM_IND_SUBPARTITIONS.....	1553
12.3.12 ADM_INDEXES.....	1554
12.3.13 ADM_OBJECTS.....	1554
12.3.14 ADM_PART_INDEXES.....	1555
12.3.15 ADM_PART_TABLES.....	1556
12.3.16 ADM_PROCEDURES.....	1557
12.3.17 ADM_SCHEDULER_JOBS.....	1558
12.3.18 ADM_SEQUENCES.....	1559
12.3.19 ADM_SOURCE.....	1560
12.3.20 ADM_SYNONYMS.....	1560
12.3.21 ADM_TAB_COLUMNS.....	1561
12.3.22 ADM_TAB_COMMENTS.....	1562
12.3.23 ADM_TAB_PARTITIONS.....	1562
12.3.24 ADM_TAB_SUBPARTITIONS.....	1563
12.3.25 ADM_TABLES.....	1564
12.3.26 ADM_TABLESPACES.....	1565
12.3.27 ADM_TRIGGERS.....	1565
12.3.28 ADM_TYPE_ATTRS.....	1566
12.3.29 ADM_USERS.....	1567
12.3.30 ADM_VIEWS.....	1567
12.3.31 DB_ALL_TABLES.....	1567
12.3.32 DB_COL_COMMENTS.....	1568
12.3.33 DB_CONS_COLUMNS.....	1568
12.3.34 DB_CONSTRAINTS.....	1569
12.3.35 DB_DEPENDENCIES.....	1569
12.3.36 DB_IND_COLUMNS.....	1570
12.3.37 DB_IND_EXPRESSIONS.....	1570
12.3.38 DB_IND_PARTITIONS.....	1571
12.3.39 DB_IND_SUBPARTITIONS.....	1572
12.3.40 DB_INDEXES.....	1573
12.3.41 DB_OBJECTS.....	1574
12.3.42 DB_PART_INDEXES.....	1574
12.3.43 DB_PART_TABLES.....	1576
12.3.44 DB_PROCEDURES.....	1577
12.3.45 DB_SEQUENCES.....	1577
12.3.46 DB_SOURCE.....	1577
12.3.47 DB_SYNONYMS.....	1578
12.3.48 DB_TAB_COLUMNS.....	1579
12.3.49 DB_TAB_COMMENTS.....	1579

12.3.50 DB_TAB_PARTITIONS.....	1580
12.3.51 DB_TAB_SUBPARTITIONS.....	1580
12.3.52 DB_TABLES.....	1581
12.3.53 DB_TRIGGERS.....	1582
12.3.54 DB_USERS.....	1582
12.3.55 DB_VIEWS.....	1583
12.3.56 DV_SESSIONS.....	1583
12.3.57 DV_SESSION_LONGOPS.....	1584
12.3.58 GET_GLOBAL_PREPARED_XACTS (Deprecated).....	1584
12.3.59 GS_ALL_CONTROL_GROUP_INFO.....	1584
12.3.60 GS_AUDITING.....	1584
12.3.61 GS_AUDITING_ACCESS.....	1585
12.3.62 GS_AUDITING_PRIVILEGE.....	1586
12.3.63 GS_CLUSTER_RESOURCE_INFO.....	1586
12.3.64 GS_COMM_PROXY_THREAD_STATUS.....	1586
12.3.65 GS_DB_PRIVILEGES.....	1587
12.3.66 GS_FILE_STAT.....	1587
12.3.67 GS_GET_CONTROL_GROUP_INFO.....	1588
12.3.68 GS_GSC_MEMORY_DETAIL.....	1588
12.3.69 GS_INSTANCE_TIME.....	1589
12.3.70 GS_LABELS.....	1589
12.3.71 GS_LSC_MEMORY_DETAIL.....	1590
12.3.72 GS_MASKING.....	1590
12.3.73 GS_MATVIEWS.....	1591
12.3.74 GS_OS_RUN_INFO.....	1591
12.3.75 GS_REDO_STAT.....	1592
12.3.76 GS_SESSION_MEMORY.....	1592
12.3.77 GS_SESSION_MEMORY_CONTEXT.....	1593
12.3.78 GS_SESSION_MEMORY_DETAIL.....	1594
12.3.79 GS_SESSION_STAT.....	1595
12.3.80 GS_SESSION_TIME.....	1595
12.3.81 GS_SQL_COUNT.....	1596
12.3.82 GS_STAT_SESSION_CU.....	1597
12.3.83 GS_THREAD_MEMORY_CONTEXT.....	1598
12.3.84 GS_TOTAL_MEMORY_DETAIL.....	1598
12.3.85 GS_TOTAL_NODEGROUP_MEMORY_DETAIL.....	1599
12.3.86 GV_SESSION.....	1600
12.3.87 MPP_TABLES.....	1601
12.3.88 MY_COL_COMMENTS.....	1602
12.3.89 MY_CONS_COLUMNS.....	1602
12.3.90 MY_CONSTRAINTS.....	1603
12.3.91 MY_IND_COLUMNS.....	1603

12.3.92 MY_IND_EXPRESSIONS.....	1604
12.3.93 MY_IND_PARTITIONS.....	1604
12.3.94 MY_IND_SUBPARTITIONS.....	1605
12.3.95 MY_INDEXES.....	1606
12.3.96 MY_JOBS.....	1607
12.3.97 MY_OBJECTS.....	1608
12.3.98 MY_PART_INDEXES.....	1609
12.3.99 MY_PART_TABLES.....	1610
12.3.100 MY_PROCEDURES.....	1611
12.3.101 MY_SEQUENCES.....	1611
12.3.102 MY_SOURCE.....	1612
12.3.103 MY_SYNONYMS.....	1612
12.3.104 MY_TAB_COLUMNS.....	1613
12.3.105 MY_TAB_COMMENTS.....	1614
12.3.106 MY_TAB_PARTITIONS.....	1614
12.3.107 MY_TAB_SUBPARTITIONS.....	1615
12.3.108 MY_TABLES.....	1616
12.3.109 MY_TRIGGERS.....	1617
12.3.110 MY_VIEWS.....	1617
12.3.111 PG_CURSORS.....	1618
12.3.112 PG_COMM_DELAY.....	1618
12.3.113 PG_COMM_RECV_STREAM.....	1619
12.3.114 PG_COMM_SEND_STREAM.....	1620
12.3.115 PG_COMM_STATUS.....	1621
12.3.116 PG_CONTROL_GROUP_CONFIG.....	1622
12.3.117 PG_EXT_STATS.....	1622
12.3.118 PG_GET_INVALID_BACKENDS.....	1625
12.3.119 PG_GET_SENDERS_CATCHUP_TIME.....	1625
12.3.120 PG_GROUP.....	1626
12.3.121 PG_GTT_ATTACHED_PIDS.....	1626
12.3.122 PG_GTT_RELSTATS.....	1627
12.3.123 PG_GTT_STATS.....	1627
12.3.124 PG_INDEXES.....	1628
12.3.125 PG_LOCKS.....	1629
12.3.126 PG_NODE_ENV.....	1630
12.3.127 PG_OS_THREADS.....	1631
12.3.128 PG_PREPARED_STATEMENTS.....	1631
12.3.129 PG_PREPARED_XACTS.....	1632
12.3.130 PG_REPLICATION_ORIGIN_STATUS.....	1633
12.3.131 PG_REPLICATION_SLOTS.....	1633
12.3.132 PG_RLSPOLICIES.....	1634
12.3.133 PG_ROLES.....	1635

12.3.134 PG_RULES.....	1637
12.3.135 PG_RUNNING_XACTS.....	1638
12.3.136 PG_SECLABELS.....	1638
12.3.137 PG_SETTINGS.....	1639
12.3.138 PG_SHADOW.....	1640
12.3.139 PG_STATS.....	1642
12.3.140 PG_STAT_ACTIVITY.....	1645
12.3.141 PG_STAT_ACTIVITY_NG.....	1648
12.3.142 PG_STAT_ALL_INDEXES.....	1651
12.3.143 PG_STAT_ALL_TABLES.....	1652
12.3.144 PG_STAT_BAD_BLOCK.....	1653
12.3.145 PG_STAT_BGWRITER.....	1654
12.3.146 PG_STAT_DATABASE.....	1655
12.3.147 PG_STAT_DATABASE_CONFLICTS.....	1656
12.3.148 PG_STAT_USER_FUNCTIONS.....	1657
12.3.149 PG_STAT_USER_INDEXES.....	1657
12.3.150 PG_STAT_USER_TABLES.....	1658
12.3.151 PG_STAT_REPLICATION.....	1659
12.3.152 PG_STAT_SYS_INDEXES.....	1661
12.3.153 PG_STAT_SYS_TABLES.....	1662
12.3.154 PG_STAT_XACT_ALL_TABLES.....	1663
12.3.155 PG_STAT_XACT_SYS_TABLES.....	1664
12.3.156 PG_STAT_XACT_USER_FUNCTIONS.....	1665
12.3.157 PG_STAT_XACT_USER_TABLES.....	1665
12.3.158 PG_STATIO_ALL_INDEXES.....	1666
12.3.159 PG_STATIO_ALL_SEQUENCES.....	1666
12.3.160 PG_STATIO_ALL_TABLES.....	1667
12.3.161 PG_STATIO_SYS_INDEXES.....	1667
12.3.162 PG_STATIO_SYS_SEQUENCES.....	1668
12.3.163 PG_STATIO_SYS_TABLES.....	1668
12.3.164 PG_STATIO_USER_INDEXES.....	1669
12.3.165 PG_STATIO_USER_SEQUENCES.....	1670
12.3.166 PG_STATIO_USER_TABLES.....	1670
12.3.167 PG_TABLES.....	1671
12.3.168 PG_TDE_INFO.....	1672
12.3.169 PG_TIMEZONE_ABBREVS.....	1672
12.3.170 PG_TIMEZONE_NAMES.....	1673
12.3.171 PG_TOTAL_MEMORY_DETAIL.....	1673
12.3.172 PG_TOTAL_USER_RESOURCE_INFO.....	1673
12.3.173 PG_TOTAL_USER_RESOURCE_INFO_OID.....	1675
12.3.174 PG_USER.....	1676
12.3.175 PG_USER_MAPPINGS.....	1678

12.3.176 PG_VARIABLE_INFO.....	1678
12.3.177 PG_VIEWS.....	1679
12.3.178 PGXC_PREPARED_XACTS.....	1679
12.3.179 PGXC_THREAD_WAIT_STATUS.....	1680
12.3.180 PLAN_TABLE.....	1680
12.3.181 SYS_DUMMY.....	1681
13 Schemas.....	1682
13.1 Information Schema.....	1684
13.1.1 _PG_FOREIGN_DATA_WRAPPERS.....	1684
13.1.2 _PG_FOREIGN_SERVERS.....	1685
13.1.3 _PG_FOREIGN_TABLE_COLUMNS.....	1686
13.1.4 _PG_FOREIGN_TABLES.....	1686
13.1.5 _PG_USER_MAPPINGS.....	1687
13.1.6 INFORMATION_SCHEMA_CATALOG_NAME.....	1687
13.2 DBE_PERF Schema.....	1687
13.2.1 OS.....	1688
13.2.1.1 OS_RUNTIME.....	1688
13.2.1.2 GLOBAL_OS_RUNTIME.....	1688
13.2.1.3 OS_THREADS.....	1689
13.2.1.4 GLOBAL_OS_THREADS.....	1689
13.2.1.5 NODE_NAME.....	1690
13.2.2 Instance.....	1690
13.2.2.1 INSTANCE_TIME.....	1690
13.2.2.2 GLOBAL_INSTANCE_TIME.....	1690
13.2.3 Memory.....	1691
13.2.3.1 GS_SHARED_MEMORY_DETAIL.....	1691
13.2.3.2 GLOBAL_MEMORY_NODE_DETAIL.....	1691
13.2.3.3 GLOBAL_SHARED_MEMORY_DETAIL.....	1693
13.2.3.4 MEMORY_NODE_DETAIL.....	1693
13.2.3.5 SHARED_MEMORY_DETAIL.....	1695
13.2.4 File.....	1695
13.2.4.1 FILE_IOSTAT.....	1695
13.2.4.2 SUMMARY_FILE_IOSTAT.....	1696
13.2.4.3 GLOBAL_FILE_IOSTAT.....	1697
13.2.4.4 FILE_REDO_IOSTAT.....	1698
13.2.4.5 SUMMARY_FILE_REDO_IOSTAT.....	1698
13.2.4.6 GLOBAL_FILE_REDO_IOSTAT.....	1699
13.2.4.7 LOCAL_REL_IOSTAT.....	1699
13.2.4.8 GLOBAL_REL_IOSTAT.....	1700
13.2.4.9 SUMMARY_REL_IOSTAT.....	1700
13.2.5 Object.....	1701
13.2.5.1 STAT_USER_TABLES.....	1701

13.2.5.2 SUMMARY_STAT_USER_TABLES.....	1702
13.2.5.3 GLOBAL_STAT_USER_TABLES.....	1703
13.2.5.4 STAT_USER_INDEXES.....	1705
13.2.5.5 SUMMARY_STAT_USER_INDEXES.....	1705
13.2.5.6 GLOBAL_STAT_USER_INDEXES.....	1706
13.2.5.7 STAT_SYS_TABLES.....	1706
13.2.5.8 SUMMARY_STAT_SYS_TABLES.....	1707
13.2.5.9 GLOBAL_STAT_SYS_TABLES.....	1709
13.2.5.10 STAT_SYS_INDEXES.....	1710
13.2.5.11 SUMMARY_STAT_SYS_INDEXES.....	1711
13.2.5.12 GLOBAL_STAT_SYS_INDEXES.....	1711
13.2.5.13 STAT_ALL_TABLES.....	1712
13.2.5.14 SUMMARY_STAT_ALL_TABLES.....	1713
13.2.5.15 GLOBAL_STAT_ALL_TABLES.....	1714
13.2.5.16 STAT_ALL_INDEXES.....	1715
13.2.5.17 SUMMARY_STAT_ALL_INDEXES.....	1716
13.2.5.18 GLOBAL_STAT_ALL_INDEXES.....	1716
13.2.5.19 STAT_DATABASE.....	1717
13.2.5.20 SUMMARY_STAT_DATABASE.....	1718
13.2.5.21 GLOBAL_STAT_DATABASE.....	1720
13.2.5.22 STAT_DATABASE_CONFLICTS.....	1722
13.2.5.23 SUMMARY_STAT_DATABASE_CONFLICTS.....	1722
13.2.5.24 GLOBAL_STAT_DATABASE_CONFLICTS.....	1722
13.2.5.25 STAT_XACT_ALL_TABLES.....	1723
13.2.5.26 SUMMARY_STAT_XACT_ALL_TABLES.....	1724
13.2.5.27 GLOBAL_STAT_XACT_ALL_TABLES.....	1724
13.2.5.28 STAT_XACT_SYS_TABLES.....	1725
13.2.5.29 SUMMARY_STAT_XACT_SYS_TABLES.....	1726
13.2.5.30 GLOBAL_STAT_XACT_SYS_TABLES.....	1726
13.2.5.31 STAT_XACT_USER_TABLES.....	1727
13.2.5.32 SUMMARY_STAT_XACT_USER_TABLES.....	1727
13.2.5.33 GLOBAL_STAT_XACT_USER_TABLES.....	1728
13.2.5.34 STAT_XACT_USER_FUNCTIONS.....	1729
13.2.5.35 SUMMARY_STAT_XACT_USER_FUNCTIONS.....	1729
13.2.5.36 GLOBAL_STAT_XACT_USER_FUNCTIONS.....	1730
13.2.5.37 STAT_BAD_BLOCK.....	1730
13.2.5.38 SUMMARY_STAT_BAD_BLOCK.....	1731
13.2.5.39 GLOBAL_STAT_BAD_BLOCK.....	1731
13.2.5.40 STAT_USER_FUNCTIONS.....	1732
13.2.5.41 SUMMARY_STAT_USER_FUNCTIONS.....	1732
13.2.5.42 GLOBAL_STAT_USER_FUNCTIONS.....	1733
13.2.6 Workload.....	1733

13.2.6.1 WORKLOAD_SQL_COUNT.....	1733
13.2.6.2 SUMMARY_WORKLOAD_SQL_COUNT.....	1734
13.2.6.3 WORKLOAD_TRANSACTION.....	1735
13.2.6.4 SUMMARY_WORKLOAD_TRANSACTION.....	1735
13.2.6.5 GLOBAL_WORKLOAD_TRANSACTION.....	1736
13.2.6.6 WORKLOAD_SQL_ELAPSE_TIME.....	1737
13.2.6.7 SUMMARY_WORKLOAD_SQL_ELAPSE_TIME.....	1738
13.2.6.8 USER_TRANSACTION.....	1739
13.2.6.9 GLOBAL_USER_TRANSACTION.....	1740
13.2.7 Session/Thread.....	1741
13.2.7.1 SESSION_STAT.....	1741
13.2.7.2 GLOBAL_SESSION_STAT.....	1742
13.2.7.3 SESSION_TIME.....	1742
13.2.7.4 GLOBAL_SESSION_TIME.....	1742
13.2.7.5 SESSION_MEMORY.....	1743
13.2.7.6 GLOBAL_SESSION_MEMORY.....	1743
13.2.7.7 SESSION_MEMORY_DETAIL.....	1744
13.2.7.8 GLOBAL_SESSION_MEMORY_DETAIL.....	1744
13.2.7.9 SESSION_STAT_ACTIVITY.....	1745
13.2.7.10 GLOBAL_SESSION_STAT_ACTIVITY.....	1748
13.2.7.11 THREAD_WAIT_STATUS.....	1751
13.2.7.12 GLOBAL_THREAD_WAIT_STATUS.....	1752
13.2.7.13 LOCAL_THREADPOOL_STATUS.....	1753
13.2.7.14 GLOBAL_THREADPOOL_STATUS.....	1754
13.2.7.15 LOCAL_ACTIVE_SESSION.....	1754
13.2.8 Transaction.....	1756
13.2.8.1 TRANSACTIONS_PREPARED_XACTS.....	1756
13.2.8.2 SUMMARY_TRANSACTIONS_PREPARED_XACTS.....	1757
13.2.8.3 GLOBAL_TRANSACTIONS_PREPARED_XACTS.....	1757
13.2.8.4 TRANSACTIONS_RUNNING_XACTS.....	1757
13.2.8.5 SUMMARY_TRANSACTIONS_RUNNING_XACTS.....	1758
13.2.8.6 GLOBAL_TRANSACTIONS_RUNNING_XACTS.....	1759
13.2.9 Query.....	1760
13.2.9.1 STATEMENT.....	1760
13.2.9.2 SUMMARY_STATEMENT.....	1762
13.2.9.3 STATEMENT_COUNT.....	1765
13.2.9.4 GLOBAL_STATEMENT_COUNT.....	1766
13.2.9.5 SUMMARY_STATEMENT_COUNT.....	1768
13.2.9.6 STATEMENT_RESPONSETIME_PERCENTILE.....	1769
13.2.9.7 STATEMENT_HISTORY.....	1770
13.2.9.8 GS_SLOW_QUERY_INFO (Deprecated).....	1773
13.2.9.9 GS_SLOW_QUERY_HISTORY (Deprecated).....	1775

13.2.9.10 GLOBAL_SLOW_QUERY_HISTORY (Deprecated).....	1775
13.2.9.11 GLOBAL_SLOW_QUERY_INFO (Deprecated).....	1775
13.2.10 Cache/IO.....	1775
13.2.10.1 STATIO_USER_TABLES.....	1775
13.2.10.2 SUMMARY_STATIO_USER_TABLES.....	1776
13.2.10.3 GLOBAL_STATIO_USER_TABLES.....	1777
13.2.10.4 STATIO_USER_INDEXES.....	1778
13.2.10.5 SUMMARY_STATIO_USER_INDEXES.....	1778
13.2.10.6 GLOBAL_STATIO_USER_INDEXES.....	1778
13.2.10.7 STATIO_USER_SEQUENCES.....	1779
13.2.10.8 SUMMARY_STATIO_USER_SEQUENCES.....	1779
13.2.10.9 GLOBAL_STATIO_USER_SEQUENCES.....	1780
13.2.10.10 STATIO_SYS_TABLES.....	1780
13.2.10.11 SUMMARY_STATIO_SYS_TABLES.....	1781
13.2.10.12 GLOBAL_STATIO_SYS_TABLES.....	1782
13.2.10.13 STATIO_SYS_INDEXES.....	1782
13.2.10.14 SUMMARY_STATIO_SYS_INDEXES.....	1783
13.2.10.15 GLOBAL_STATIO_SYS_INDEXES.....	1783
13.2.10.16 STATIO_SYS_SEQUENCES.....	1784
13.2.10.17 SUMMARY_STATIO_SYS_SEQUENCES.....	1784
13.2.10.18 GLOBAL_STATIO_SYS_SEQUENCES.....	1785
13.2.10.19 STATIO_ALL_TABLES.....	1785
13.2.10.20 SUMMARY_STATIO_ALL_TABLES.....	1786
13.2.10.21 GLOBAL_STATIO_ALL_TABLES.....	1786
13.2.10.22 STATIO_ALL_INDEXES.....	1787
13.2.10.23 SUMMARY_STATIO_ALL_INDEXES.....	1788
13.2.10.24 GLOBAL_STATIO_ALL_INDEXES.....	1788
13.2.10.25 STATIO_ALL_SEQUENCES.....	1789
13.2.10.26 SUMMARY_STATIO_ALL_SEQUENCES.....	1789
13.2.10.27 GLOBAL_STATIO_ALL_SEQUENCES.....	1789
13.2.10.28 GLOBAL_STAT_DB_CU.....	1790
13.2.10.29 GLOBAL_STAT_SESSION_CU.....	1790
13.2.11 Utility.....	1791
13.2.11.1 REPLICATION_STAT.....	1791
13.2.11.2 GLOBAL_REPLICATION_STAT.....	1792
13.2.11.3 REPLICATION_SLOTS.....	1793
13.2.11.4 GLOBAL_REPLICATION_SLOTS.....	1794
13.2.11.5 BGWRITER_STAT.....	1795
13.2.11.6 GLOBAL_BGWRITER_STAT.....	1795
13.2.11.7 GLOBAL_CKPT_STATUS.....	1796
13.2.11.8 GLOBAL_DOUBLE_WRITE_STATUS.....	1797
13.2.11.9 GLOBAL_PAGEWRITER_STATUS.....	1798

13.2.11.10 GLOBAL_RECORD_RESET_TIME.....	1798
13.2.11.11 GLOBAL_REDO_STATUS.....	1799
13.2.11.12 GLOBAL_RECOVERY_STATUS.....	1800
13.2.11.13 CLASS_VITAL_INFO.....	1801
13.2.11.14 USER_LOGIN.....	1801
13.2.11.15 SUMMARY_USER_LOGIN.....	1802
13.2.11.16 GLOBAL_SINGLE_FLUSH_DW_STATUS.....	1802
13.2.11.17 GLOBAL_CANDIDATE_STATUS.....	1803
13.2.12 Lock.....	1803
13.2.12.1 LOCKS.....	1803
13.2.12.2 GLOBAL_LOCKS.....	1805
13.2.13 Wait Events.....	1806
13.2.13.1 WAIT_EVENTS.....	1806
13.2.13.2 GLOBAL_WAIT_EVENTS.....	1807
13.2.14 Configuration.....	1808
13.2.14.1 CONFIG_SETTINGS.....	1808
13.2.14.2 GLOBAL_CONFIG_SETTINGS.....	1809
13.2.15 Operator.....	1810
13.2.15.1 OPERATOR_HISTORY_TABLE.....	1810
13.2.15.2 OPERATOR_HISTORY.....	1811
13.2.15.3 OPERATOR_RUNTIME.....	1812
13.2.15.4 GLOBAL_OPERATOR_HISTORY.....	1813
13.2.15.5 GLOBAL_OPERATOR_HISTORY_TABLE.....	1815
13.2.15.6 GLOBAL_OPERATOR_RUNTIME.....	1815
13.2.16 Workload Manager.....	1817
13.2.16.1 WLM_USER_RESOURCE_CONFIG.....	1817
13.2.16.2 WLM_USER_RESOURCE_RUNTIME.....	1817
13.2.17 Global Plancache.....	1818
13.2.17.1 GLOBAL_PLANCACHE_STATUS.....	1818
13.2.17.2 GLOBAL_PLANCACHE_CLEAN.....	1819
13.2.18 RTO & RPO.....	1819
13.2.18.1 global_rto_status.....	1819
13.2.18.2 global_streaming_hadr_rto_and_rpo_stat.....	1819
13.3 DBE_PLDEBUGGER Schema.....	1820
13.3.1 DBE_PLDEBUGGER.turn_on.....	1824
13.3.2 DBE_PLDEBUGGER.turn_off.....	1824
13.3.3 DBE_PLDEBUGGER.local_debug_server_info.....	1825
13.3.4 DBE_PLDEBUGGER.attach.....	1825
13.3.5 DBE_PLDEBUGGER.info_locals.....	1826
13.3.6 DBE_PLDEBUGGER.next.....	1826
13.3.7 DBE_PLDEBUGGER.continue.....	1827
13.3.8 DBE_PLDEBUGGER.abort.....	1827

13.3.9 DBE_PLDEBUGGER.print_var.....	1828
13.3.10 DBE_PLDEBUGGER.info_code.....	1828
13.3.11 DBE_PLDEBUGGER.step.....	1829
13.3.12 DBE_PLDEBUGGER.add_breakpoint.....	1829
13.3.13 DBE_PLDEBUGGER.delete_breakpoint.....	1830
13.3.14 DBE_PLDEBUGGER.info_breakpoints.....	1830
13.3.15 DBE_PLDEBUGGER.backtrace.....	1830
13.3.16 DBE_PLDEBUGGER.enable_breakpoint.....	1831
13.3.17 DBE_PLDEBUGGER.disable_breakpoint.....	1831
13.3.18 DBE_PLDEBUGGER.finish.....	1831
13.3.19 DBE_PLDEBUGGER.set_var.....	1832
13.4 DBE_PLDEVELOPER.....	1832
13.4.1 DBE_PLDEVELOPER.gs_source.....	1832
13.4.2 DBE_PLDEVELOPER.gs_errors.....	1833
13.5 DBE_SQL_UTIL Schema.....	1834
13.5.1 DBE_SQL_UTIL.create_hint_sql_patch.....	1834
13.5.2 DBE_SQL_UTIL.create_abort_sql_patch.....	1834
13.5.3 DBE_SQL_UTIL.drop_sql_patch.....	1835
13.5.4 DBE_SQL_UTIL.enable_sql_patch.....	1835
13.5.5 DBE_SQL_UTIL.disable_sql_patch.....	1836
13.5.6 DBE_SQL_UTIL.show_sql_patch.....	1836
14 Logical Replication.....	1837
14.1 Logical Decoding.....	1837
14.1.1 Overview.....	1838
14.1.2 Logical Decoding Options.....	1841
14.1.3 Logical Decoding by SQL Functions.....	1846
14.1.4 Replicating Data Using the Logical Replication Tool.....	1847
15 Materialized View.....	1848
15.1 Complete-refresh Materialized View.....	1848
15.1.1 Overview.....	1848
15.1.2 Usage.....	1848
15.1.3 Support and Constraints.....	1849
15.2 Fast-refresh Materialized View.....	1849
15.2.1 Overview.....	1849
15.2.2 Usage.....	1850
15.2.3 Support and Constraints.....	1851
16 Error Log Reference.....	1852
16.1 Kernel Error Information.....	1852
16.2 CM Error Information.....	1875
17 Configuring GUC Parameters.....	1885
17.1 Viewing Parameter Values.....	1885

17.2 Resetting Parameters.....	1886
17.3 GUC Parameters.....	1894
17.3.1 GUC Parameter Usage.....	1894
17.3.2 File Location.....	1894
17.3.3 Connection and Authentication.....	1896
17.3.3.1 Connection Settings.....	1896
17.3.3.2 Security and Authentication (postgresql.conf).....	1902
17.3.3.3 Communication Library Parameters.....	1912
17.3.4 Resource Consumption.....	1915
17.3.4.1 Memory.....	1915
17.3.4.2 Disk Space.....	1925
17.3.4.3 Kernel Resource Usage.....	1926
17.3.4.4 Cost-based Vacuum Delay.....	1927
17.3.4.5 Background Writer.....	1929
17.3.4.6 Asynchronous I/O.....	1932
17.3.5 Data Import and Export.....	1933
17.3.6 Write Ahead Log.....	1934
17.3.6.1 Settings.....	1934
17.3.6.2 Checkpoints.....	1943
17.3.6.3 Log Replay.....	1945
17.3.6.4 Archiving.....	1948
17.3.7 HA Replication.....	1949
17.3.7.1 Sending Server.....	1949
17.3.7.2 Primary Server.....	1959
17.3.7.3 Standby Server.....	1966
17.3.8 Query Planning.....	1970
17.3.8.1 Optimizer Method Configuration.....	1970
17.3.8.2 Optimizer Cost Constants.....	1977
17.3.8.3 Genetic Query Optimizer.....	1979
17.3.8.4 Other Optimizer Options.....	1981
17.3.9 Error Reporting and Logging.....	1996
17.3.9.1 Logging Destination.....	1997
17.3.9.2 Logging Time.....	2001
17.3.9.3 Logging Content.....	2004
17.3.9.4 Using CSV Log Output.....	2013
17.3.10 Alarm Detection.....	2015
17.3.11 Statistics During the Database Running.....	2016
17.3.11.1 Query and Index Statistics Collector.....	2016
17.3.11.2 Performance Statistics.....	2020
17.3.12 Automatic Vacuuming.....	2021
17.3.13 Default Settings of Client Connection.....	2025
17.3.13.1 Statement Behavior.....	2025


17.3.13.2 Locale and Formatting.....	2031
17.3.13.3 Other Default Parameters.....	2035
17.3.14 Lock Management.....	2036
17.3.15 Version and Platform Compatibility.....	2040
17.3.15.1 Compatibility with Earlier Versions.....	2040
17.3.15.2 Platform and Client Compatibility.....	2044
17.3.16 Fault Tolerance.....	2060
17.3.17 Connection Pool Parameters.....	2062
17.3.18 Transaction.....	2062
17.3.19 Replication Parameters of Two Database Instances.....	2066
17.3.20 Developer Options.....	2067
17.3.21 Auditing.....	2077
17.3.21.1 Audit Switch.....	2077
17.3.21.2 User and Permission Audit.....	2080
17.3.21.3 Operation Audit.....	2082
17.3.22 CM Parameters.....	2089
17.3.22.1 CM Agent Parameters.....	2089
17.3.22.2 Parameters Related to CM Server.....	2095
17.3.23 Upgrade Parameters.....	2104
17.3.24 Miscellaneous Parameters.....	2105
17.3.25 Wait Events.....	2110
17.3.26 Query.....	2111
17.3.27 System Performance Snapshot.....	2116
17.3.28 Security Configuration.....	2118
17.3.29 Global Temporary Table.....	2119
17.3.30 HyperLogLog.....	2120
17.3.31 User-defined Functions.....	2121
17.3.32 Scheduled Task.....	2122
17.3.33 Thread Pool.....	2123
17.3.34 Backup and Restoration.....	2126
17.3.35 Undo.....	2127
17.3.36 DCF Parameters Settings.....	2127
17.3.37 Flashback.....	2135
17.3.38 Rollback Parameters.....	2136
17.3.39 AI Features.....	2137
17.3.40 Global SysCache Parameters.....	2137
17.3.41 Reserved Parameters.....	2138

1 Database System Overview

1.1 Database Logical Architecture

Data nodes (DNs) in GaussDB store data on disks. This section describes the objects on each DN from the logical view and the relationship between these objects. [Figure 1-1](#) shows the database logical structure.

Figure 1-1 Database logical architecture




NOTE

- A tablespace is a directory storing physical files of one or more databases. An instance may contain multiple tablespaces.
- Database: A database manages various data objects and is isolated from each other. Objects managed by a database can be distributed to multiple tablespaces.
- Datafile Segment: Data file, each of which stores data of only one table. A table containing more than 1 GB of data is stored in multiple datafile segments.
- One table belongs to only one database and one tablespace. The datafile segments storing the data of the same table must be in the same tablespace.
- Block: Basic unit of database management. Its default size is 8 KB.

1.2 Query Request Handling Process

Figure 1-2 GaussDB service response process



1.3 Managing Transactions

A transaction is a customized sequence of database operations, which form an integral unit of work. In GaussDB, you can start, set, commit, and roll back transactions. A GaussDB database supports the following transaction isolation levels: READ COMMITTED, READ UNCOMMITTED (not recommended), REPEATABLE READ, and SERIALIZABLE. SERIALIZABLE is equivalent to REPEATABLE READ.

Controlling Transactions

The following describes transaction operations supported by the database:

- Starting transactions

You can use the **START TRANSACTION** or **BEGIN** syntax to start a transaction. For details, see [START TRANSACTION](#) and [BEGIN](#).

- Setting transactions

You can use the **SET TRANSACTION** or **SET LOCAL TRANSACTION** syntax to set transactions. For details, see [SET TRANSACTION](#).

- **Committing transactions**
You can commit all operations of a transaction using **COMMIT** or **END**. For details, see [COMMIT | END](#).
- **Rolling back transactions**
Rollback indicates that the system cancels all changes that a transaction has made to a database if the transaction fails to be executed due to a fault. For details, see [ROLLBACK](#).

Transaction Isolation Levels

A transaction isolation level specifies how concurrent transactions process the same object.

NOTE

The isolation level cannot be changed after data is modified using **SELECT**, **INSERT**, **DELETE**, **UPDATE**, **FETCH**, or **COPY** in the transaction.

- **READ COMMITTED:** At this level, a transaction can access only committed data. This is the default level.

The **SELECT** statement accesses the snapshot of the database taken when the query begins. The **SELECT** statement can also access the data modifications in its transaction, regardless of whether they have been committed. Note that different database snapshots may be available to two consecutive **SELECT** statements for the same transaction, because data may be committed for other transactions while the first **SELECT** statement is executed.

At the **READ COMMITTED** level, the execution of each statement begins with a new snapshot, which contains all the transactions that have been committed by the execution time. Therefore, during a transaction, a statement can access the result of other committed transactions. Note whether a single statement always accesses absolutely consistent data in a database.

Transaction isolation at this level meets the requirements of many applications, and is fast and easy to use. However, applications performing complicated queries and updates may require data that is more consistent than this level can provide.
- **READ UNCOMMITTED:** This level is not recommended, because it may result in data inconsistency. You are advised not to use it to write transactions, as this may lead to data inconsistency. However, you can use it to read transactions in emergency cases.
- **REPEATABLE READ:** At this level, a transaction can only read data committed before it starts. Uncommitted data or data committed in other concurrent transactions cannot be read. However, a query can read earlier data modifications in its transaction, regardless of whether they have been committed. **READ COMMITTED** differs from this level in that a transaction reads the snapshot taken at the start of the transaction, not at the beginning of the current query within the transaction. Therefore, the **SELECT** statement within a transaction always reads the same data, and cannot read data committed by other concurrent transactions after the transaction starts. Applications at this level must be able to retry transactions, because serialization failures may occur.
- **SERIALIZABLE:** Currently, GaussDB does not support this isolation level. Setting this isolation level is equivalent to **REPEATABLE READ**.

1.4 Concepts

Database

Databases manage various data objects and are isolated from each other. While creating a database, you can specify a tablespace. If you do not specify it, the object will be saved to the **PG_DEFAULT** tablespace by default. Objects managed by a database can be distributed to multiple tablespaces.

Tablespace

In GaussDB, a tablespace is a directory storing physical files of databases. An instance may contain multiple tablespaces. Files are physically isolated using tablespaces and managed by a file system.

Schema

GaussDB schemas logically separate databases. All database objects are created under certain schemas. In GaussDB, schemas and users are loosely bound. When you create a user, a schema with the same name as the user will be created automatically. You can also create a schema or specify another schema.

User and Role

GaussDB uses users and roles to control the access to databases. A role can be a database user or a group of database users, depending on role settings. In GaussDB, the difference between roles and users is that a role does not have the **LOGIN** permission by default. In GaussDB, one user can have only one role, but you can put a user's role under a parent role to grant multiple permissions to the user.

Transaction

In GaussDB, transactions are managed by multi-version concurrency control (MVCC) and two-phase locking (2PL). It enables smooth data reads and writes. In GaussDB, Astore stores various versions of historical data together with the current tuple version. In GaussDB, Astore uses a VACUUM thread instead of rollback segments to periodically delete historical data. Generally, you do not need to pay special attention to the VACUUM thread unless you need to optimize the performance. In GaussDB, Ustore stores various versions of historical data to the undo rollback segments. The thread for purging undo deletes various versions of historical data in a unified manner. In addition, GaussDB automatically commits transactions for single-statement queries (without using statements such as BEGIN to explicitly start a transaction block).

2 Database Security Management

2.1 Users and Permissions

2.1.1 Default Permission Mechanism

A user who creates an object is the owner of this object. By default, [separation of duties](#) is disabled after database installation. A database system administrator has the same permissions as object owners. After an object is created, only the object owner or system administrator can query, modify, and delete the object, and grant permissions for the object to other users through [GRANT](#) by default.

To enable another user to use the object, grant required permissions to the user or the role that contains the user.

GaussDB supports the following permissions: SELECT, INSERT, UPDATE, DELETE, TRUNCATE, REFERENCES, CREATE, CONNECT, EXECUTE, USAGE, ALTER, DROP, COMMENT, INDEX, and VACUUM. Permission types are associated with object types. For permission details, see [GRANT](#).

To revoke permissions, run [REVOKE](#). Object owners have implicit permissions (such as ALTER, DROP, COMMENT, INDEX, VACUUM, GRANT, and REVOKE) on objects. That is, once becoming the owner of an object, the owner is immediately granted the implicit permissions on the object. Object owners can remove their own common permissions, for example, making tables read-only to themselves or others, except the system administrator.

System catalogs and views are visible to either system administrators or all users. System catalogs and views that require system administrator permissions can be queried only by system administrators. For details, see [System Catalogs and System Views](#).

The database provides the object isolation feature. If this feature is enabled, users can view only the objects (tables, views, columns, and functions) that they have the permission to access. System administrators are not affected by this feature. For details, see [ALTER DATABASE](#).

You are advised not to modify the permissions on system catalogs or system views.

2.1.2 Administrator

Initial User

The account automatically generated during database installation is called an initial user. The initial user has the highest-level permissions in the system, including SYSADMIN, CREATEROLE, AUDITADMIN, MONADMIN, OPRADMIN, and POLADMIN, and can perform all operations. If the initial username is not specified during installation, the username is the same as the name of the OS user who installs the database. If the password of the initial user is not specified during the installation, the password is empty after the installation. In this case, you need to change the password of the initial user on the GSQL client before performing other operations. If the initial user password is empty, you cannot perform any SQL operations, such as upgrade, scaling, and node replacement, except changing the password.

The initial user bypasses all permission checks. You are advised to use the initial user as a database administrator only for database management other than service running.

System Administrator

A system administrator is an account with the SYSADMIN attribute. By default, a system administrator has the same permissions as the object owner but does not have the object permissions in the dbf_perf schema or the permission to use Roach to perform backup and restoration.

To create a database administrator, connect to the database as an administrator and run the **CREATE USER** or **ALTER USER** statement with **SYSADMIN** specified.

```
openGauss=# CREATE USER sysadmin WITH SYSADMIN password '*****';
```

or

```
openGauss=# ALTER USER joe SYSADMIN;
```

To run the **ALTER USER** statement, the user must exist.

Security Administrator

A security administrator is an account with the CREATEROLE attribute. It has the permission to create, modify, and delete users or roles, and grant or revoke the permission of any non-system administrator, built-in role, permanent user, or O&M administrator.

If you want to create a security administrator when separation of duties is disabled, connect to the database as a system administrator or security administrator. When separation of duties is enabled, connect to the database as a security administrator and use the **CREATE USER** or **ALTER USER** statement with the **CREATEROLE** option.

```
openGauss=# CREATE USER createrole WITH CREATEROLE password '*****';
```

or

```
openGauss=# ALTER USER joe CREATEROLE;
```

To run the **ALTER USER** statement, the user must exist.

Audit Administrator

An audit administrator is an account with the AUDITADMIN attribute, which has the permission to view and delete audit logs.

If you want to create an audit administrator when separation of duties is disabled, connect to the database as a system administrator or security administrator. When separation of duties is enabled, connect to the database only as the initial user and use the **CREATE USER** or **ALTER USER** statement with the **AUDITADMIN** option.

```
openGauss=# CREATE USER auditadmin WITH AUDITADMIN password '*****';
```

or

```
openGauss=# ALTER USER joe AUDITADMIN;
```

To run the **ALTER USER** statement, the user must exist.

Monitor Administrator

A monitor administrator is an account with the MONADMIN attribute and has the permission to view views and functions in the db_perf schema. The monitor administrator can also grant or revoke object permissions in the db_perf schema.

To create a monitor administrator, connect to the database as a system administrator and run the **CREATE USER** or **ALTER USER** statement with **MONADMIN** specified.

```
openGauss=# CREATE USER monadmin WITH MONADMIN password '*****';
```

Alternatively,

```
openGauss=# ALTER USER joe MONADMIN;
```

To run the **ALTER USER** statement, the user must exist.

O&M Administrator

An O&M administrator is an account with the OPRADMIN attribute and has the permission to use Roach to perform backup and restoration.

To create an O&M administrator, connect to the database as an initial user and run the **CREATE USER** or **ALTER USER** statement with **OPRADMIN** specified.

```
openGauss=# CREATE USER opradmin WITH OPRADMIN password '*****';
```

Alternatively,

```
openGauss=# ALTER USER joe OPRADMIN;
```

To run the **ALTER USER** statement, the user must exist.

Security Policy Administrator

A security policy administrator is an account with the POLADMIN attribute and has the permission to create resource tags, masking policies, and unified audit policies.

To create a security policy administrator, connect to the database as an administrator and run the **CREATE USER** or **ALTER USER** statement with **POLADMIN** specified.

```
openGauss=# CREATE USER poladmin WITH POLADMIN password '*****';
```

Alternatively,

```
openGauss=# ALTER USER joe POLADMIN;
```

To run the **ALTER USER** statement, the user must exist.

2.1.3 Separation of Duties

Descriptions in **Default Permission Mechanism** and **Administrator** are about the initial situation after the database system is created. By default, a system administrator with the SYSADMIN attribute has the highest-level permissions.

To avoid risks caused by centralized permissions, you can enable separation of duties to assign the system administrator's user creation permission to security administrators and audit management permission to audit administrators.

After separation of duties, the system administrator does not have the CREATEROLE attribute (security administrator) or the AUDITADMIN attribute (audit administrator). That is, the system administrator can neither create roles or users, nor view or maintain database audit logs. For details about the CREATEROLE and AUDITADMIN attributes, see **CREATE ROLE**.

Separation of duties does not take effect for an initial user. Therefore, you are advised to use an initial user as a database administrator only for database management other than service running.

To enable separation of duties, set GUC parameter **enableSeparationOfDuty** to **on**.



If you need to use the separation of duties model, specify it during database initialization. You are advised not to switch the permission management model back and forth. In particular, if you want to switch from a non-separation-of-duties permission management model to the separation-of-duties permission management model, you need to review the permission set of existing users. If a user has the system administrator permission and audit administrator permission, the permissions need to be tailored.

After separation of duties, the system administrator does not have permissions for non-system schemas of other users. Therefore, the system administrator cannot access the objects in other users' schemas before being granted the permissions. For details about permission changes before and after enabling separation of duties, see **Table 2-1** and **Table 2-2**.

Table 2-1 Default user permissions

Object Name	Initial User (ID: 10)	System Administrator	Security Administrator	Audit Administrator	Common User
Tablespaces	Has all permissions.	Can create, modify, delete, access, or grant permissions for tablespaces.	Cannot create, modify, delete, or grant permissions for tablespaces and can access tablespaces if the access permission is granted.		
Schemas		Has all permissions for all schemas except dbperf .	Has all permissions for their own schemas, but does not have permissions for non-system schemas of other users.		
User-defined functions		Has all permissions for all user-defined functions.	Has all permissions for their own functions, and has only the call permission for other users' functions.		
User-defined tables or views		Has all permissions for all user-defined tables or views.	Has all permissions for their own tables or views, but does not have permissions for other users' tables or views.		

Table 2-2 Changes in permissions after separation of duties

Object Name	Initial User (ID: 10)	System Administrator	Security Administrator	Audit Administrator	Common User
Tablespaces	N/A Has all permissions.	N/A	N/A		
Schemas		Permissions reduced Has all permissions for their own schemas, but does not have permissions for non-system schemas of other users.	N/A		
User-defined functions		Cannot access functions in non-system schemas of other users before being granted the permissions.	N/A		

Object Name	Initial User (ID: 10)	System Administrator	Security Administrator	Audit Administrator	Common User
User-defined tables or views		Cannot access tables or views in non-system schemas of other users before being granted the permissions.	N/A		

NOTICE

PG_STATISTIC and **PG_STATISTIC_EXT** store sensitive information about statistical objects, such as high-frequency MCVs. After separation of duties is enabled, a system administrator can still access the two system catalogs to obtain sensitive information in statistics.

2.1.4 Users

You can use **CREATE USER** and **ALTER USER** to create and manage database users, respectively. A database system contains one or more databases. Users and roles are shared within the entire database system, but their data is not shared. That is, a user can connect to any database, but after the connection is successful, any user can access only the database declared in the connection request.

In non-**Separation of Duties** scenarios, GaussDB user accounts can be created and deleted only by a system administrator or a security administrator with the **CREATEROLE** attribute. In separation-of-duties scenarios, a user account can be created only by an initial user or a security administrator.

When a user logs in, GaussDB authenticates the user. A user can own databases and database objects (such as tables), and grant permissions on these objects to other users and roles. In addition to system administrators, users with the **CREATEDB** attribute can create databases and grant permissions on these databases.

Adding, Modifying, and Deleting Users

- To create a user, use the SQL statement **CREATE USER**.
For example, create a user **joe** and set the **CREATEDB** attribute for the user.

```
openGauss=# CREATE USER joe WITH CREATEDB PASSWORD '*****';  
CREATE ROLE
```
- To create a system administrator, use the **CREATE USER** statement with the **SYSADMIN** parameter.
- To delete an existing user, use **DROP USER**.
- To change a user account (for example, rename the user or change the password), use **ALTER USER**.

- To view the user list, query the view **PG_USER**:
`openGauss=# SELECT * FROM pg_user;`
- To view user attributes, query the system catalog **PG_AUTHID**.
`openGauss=# SELECT * FROM pg_authid;`

Permanent User

GaussDB provides the permanent user solution. That is, create a permanent user with the **PERSISTENCE** attribute.

```
openGauss=# CREATE USER user_persistence WITH PERSISTENCE IDENTIFIED BY "*****";
```

Only the initial user is allowed to create, modify, and delete permanent users with the **PERSISTENCE** attribute.

2.1.5 Roles

After a role is granted to a user through **GRANT**, the user will have all the permissions of the role. It is recommended that roles be used to efficiently grant permissions. For example, you can create different roles of design, development, and maintenance personnel, grant the roles to users, and then grant specific data permissions required by different users. When permissions are granted or revoked at the role level, these changes take effect on all members of the role.

GaussDB provides an implicitly defined group **PUBLIC** that contains all roles. By default, all new users and roles have the permissions of **PUBLIC**. For details about the default permissions of **PUBLIC**, see **GRANT**. To revoke permissions of **PUBLIC** from a user or role, or re-grant these permissions to them, add the **PUBLIC** keyword in the **REVOKE** or **GRANT** statement.

To view all roles, query the system catalog **PG_ROLES**.

```
SELECT * FROM PG_ROLES;
```

Adding, Modifying, and Deleting Roles

In scenarios other than separation of duties, a role can be created, modified, and deleted only by a system administrator or a user with the **CREATEROLE** attribute. In **separation-of-duties** scenarios, a role can be created, modified, and deleted only by the initial user or a user with the **CREATEROLE** attribute.

- To create a role, use **CREATE ROLE**.
- To add or delete users in an existing role, use **ALTER ROLE**.
- To delete a role, use **DROP ROLE**. **DROP ROLE** deletes only a role, rather than role members.

Built-in Roles

GaussDB provides a group of default roles whose names start with "gs_role_". These roles are provided to access to specific, typically high-privileged operations. You can grant these roles to other users or roles within the database so that they can use specific functions. These roles should be given with great care to ensure that they are used where they are needed. **Table 2-3** describes the permissions of built-in roles.

Table 2-3 Permission description of built-in roles

Roles	Permission
gs_role_copy_files	Permission to run the copy... to/from filename command. However, the GUC parameter enable_copy_server_files must be set first to enable the function of copying server files.
gs_role_signal_backend	Permission to call the pg_cancel_backend , pg_terminate_backend , and pg_terminate_session functions to cancel or terminate other sessions. However, this role cannot perform operations on sessions of the initial user or PERSISTENCE user.
gs_role_tablespace	Permission to create a tablespace.
gs_role_replication	Permission to call logical replication functions, such as kill_snapshot , pg_create_logical_replication_slot , pg_create_physical_replication_slot , pg_drop_replication_slot , pg_replication_slot_advance , pg_create_physical_replication_slot_extern , pg_logical_slot_get_changes , pg_logical_slot_peek_changes , pg_logical_slot_get_binary_changes , and pg_logical_slot_peek_binary_changes .
gs_role_account_lock	Permission to lock and unlock users. However, this role cannot lock or unlock the initial user or users with the PERSISTENCE attribute.
gs_role_pldebugger	Permission to debug functions in db_pldebugger .
gs_role_directory_create	Permission to create directory objects. However, this role needs to enable the GUC parameter enable_access_server_directory first.
gs_role_directory_drop	Permission to delete directory objects. However, this role needs to enable the GUC parameter enable_access_server_directory first.

The restrictions on built-in roles are as follows:

- The role names starting with "gs_role_" are reserved for built-in roles in the database. Do not create users or roles starting with "gs_role_" or rename existing users or roles to names starting with "gs_role_".
- Do not perform the ALTER or DROP operation on built-in roles.
- By default, built-in roles do not have the LOGIN permission and do not have preset passwords.
- The gsql meta-commands **\du** and **\dg** do not display information about built-in roles. However, if *pattern* is specified as a specific built-in role, the information is displayed.
- When separation-of-duties is disabled, the initial user, users with the SYSADMIN permission, and users with the ADMIN OPTION built-in role

permission have the permission to perform GRANT and REVOKE operations on built-in roles. When separation of duties is enabled, the initial user and users with the ADMIN OPTION built-in role permission have the permission to perform GRANT and REVOKE operations on built-in roles. Example:

```
GRANT gs_role_signal_backend TO user1;  
REVOKE gs_role_signal_backend FROM user1;
```

2.1.6 Schemas

Schemas function as models. Schema management allows multiple users to use the same database without mutual impacts, to organize database objects as manageable logical groups, and to add third-party applications to the same schema without causing conflicts.

Each database has one or more schemas. Each schema contains tables and other types of objects. When a database is created, a public schema named **public** is created by default, and all users have the **USAGE** permission on this schema. In addition, each database has a **pg_catalog** schema, which contains system catalogs and all built-in data types, functions, and operators. Only system administrators and the initial user can create common functions, aggregate functions, stored procedures, and synonym objects in public and pg_catalog schemas. Only the initial user can create operators in public and pg_catalog schemas. Other users cannot create the preceding five types of objects even if they are granted the CREATE permission on the public and pg_catalog schemas. You can group database objects by schema. A schema is similar to an OS directory but cannot be nested. By default, only the initial user can create objects in **pg_catalog**.

The same database object name can be used in different schemas of the same database without causing conflicts. For example, both **a_schema** and **b_schema** can contain a table named **mytable**. Users with required permissions can access objects across multiple schemas of the same database.

When you run the **CREATE USER** command to create a user, the system creates a schema with the same name as the user in the database where the command is executed.

Database objects are generally created in the first schema in a database search path. For details about the first schema and how to change the schema order, see [Search Path](#).

Creating, Modifying, and Deleting Schemas

- To create a schema, use **CREATE SCHEMA**. By default, the initial user and system administrators can create schemas. Other users can create schemas in the database only when they have the CREATE permission on the database. For details about how to grant the permission, see the syntax in [GRANT](#).
- To change the name or owner of a schema, use **ALTER SCHEMA**. The schema owner can change the schema.
- To delete a schema and its objects, use **DROP SCHEMA**. Schema owners can delete schemas.
- To create a table in a schema, use the *schema_name.table_name* format to specify the table. If *schema_name* is not specified, the table will be created in the first schema in [search path](#).

- To view the owner of a schema, perform the following join query on the system catalogs **PG_NAMESPACE** and **PG_USER**. Replace *schema_name* in the statement with the name of the schema to be queried.

```
openGauss=# SELECT s.nspname,u.username AS nspowner FROM pg_namespace s, pg_user u WHERE  
nspname='schema_name' AND s.nspowner = u.usesysid;
```
- To view a list of all schemas, query the system catalog **PG_NAMESPACE**.

```
openGauss=# SELECT * FROM pg_namespace;
```
- To view a list of tables in a schema, query the system catalog **PG_TABLES**. For example, the following query will return a table list from **PG_CATALOG** in the schema.

```
openGauss=# SELECT distinct(tablename),schemaname from pg_tables where schemaname =  
'pg_catalog';
```

Search Path

A search path is defined in the **search_path** parameter. The parameter value is a list of schema names separated by commas (,). If no target schema is specified during object creation, the object will be added to the first schema listed in the search path. If there are objects with the same name across different schemas and no schema is specified for an object query, the object will be returned from the first schema containing the object in the search path.

- To view the current search path, use **SHOW**.

```
openGauss=# SHOW SEARCH_PATH;  
search_path  
-----  
"$user",public  
(1 row)
```

The default value of **search_path** is **"\$user",public**. *\$user* indicates the name of the schema with the same name as the current session user. If the schema does not exist, *\$user* will be ignored. By default, after a user connects to a database that has schemas with the same name, objects will be added to all the schemas. If there are no such schemas, objects will be added only to the **public** schema.

- To change the default schema of the current session, run the **SET** statement.

Run the following command to set the search path to **myschema,public** (**myschema** will be searched before **public**):

```
openGauss=# SET SEARCH_PATH TO myschema, public;  
SET
```

2.1.7 Setting User Permissions

- To grant permissions for an object to a user, use **GRANT**.

When permissions for a table or view in a schema are granted to a user or role, the **USAGE** permission of the schema must be granted together. Otherwise, the user or role can only see these objects but cannot access them.

In the following example, permissions for the schema **tpcds** are first granted to user **joe**, and then the **SELECT** permission for the **tpcds.web_returns** table is also granted.

```
openGauss=# GRANT USAGE ON SCHEMA tpcds TO joe;  
openGauss=# GRANT SELECT ON TABLE tpcds.web_returns to joe;
```

- Grant a role to a user to allow the user to inherit the object permissions of the role.

- a. Create a role.

Create a role **lily** and grant the system permission **CREATEDB** to the role.

```
openGauss=# CREATE ROLE lily WITH CREATEDB PASSWORD '*****!';
```

- b. Grant object permissions to the role by using **GRANT**.

For example, first grant permissions for the schema **tpcds** to the role **lily**, and then grant the **SELECT** permission of the **tpcds.web_returns** table to **lily**.

```
openGauss=# GRANT USAGE ON SCHEMA tpcds TO lily;
openGauss=# GRANT SELECT ON TABLE tpcds.web_returns TO lily;
```

- c. Grant the role permissions to a user.

```
openGauss=# GRANT lily TO joe;
```

NOTE

When the permissions of a role are granted to a user, the attributes of the role are not transferred together.

- To revoke user permissions, use **REVOKE**.

2.1.8 Row-Level Security Policy

The row-level security feature enables database access control to be accurate to each row of data tables. In this way, the same SQL query may return different results for different users.

You can create a row-level security policy for a data table. The policy defines an expression that takes effect only for specific database users and SQL operations. When a database user accesses the data table, if a SQL statement meets the specified row-level security policies of the data table, the expressions that meet the specified condition will be combined by using **AND** or **OR** based on the attribute type (**PERMISSIVE** | **RESTRICTIVE**) and applied to the execution plan in the query optimization phase.

Row-level security policy is used to control the visibility of row-level data in tables. By predefining filters for data tables, the expressions that meet the specified condition can be applied to execution plans in the query optimization phase, which will affect the final execution result. Currently, the SQL statements that can be affected include **SELECT**, **UPDATE**, and **DELETE**.

Scenario 1: A table summarizes the data of different users. Users can view only their own data.

```
-- Create users alice, bob, and peter.
openGauss=# CREATE USER alice PASSWORD '*****!';
openGauss=# CREATE USER bob PASSWORD '*****!';
openGauss=# CREATE USER peter PASSWORD '*****!';

-- Create the all_data table that contains user information.
openGauss=# CREATE TABLE all_data(id int, role varchar(100), data varchar(100));

-- Insert data into the data table.
openGauss=# INSERT INTO all_data VALUES(1, 'alice', 'alice data');
openGauss=# INSERT INTO all_data VALUES(2, 'bob', 'bob data');
openGauss=# INSERT INTO all_data VALUES(3, 'peter', 'peter data');

-- Grant the read permission for the all_data table to users alice, bob, and peter.
openGauss=# GRANT SELECT ON all_data TO alice, bob, peter;

-- Enable the row-level security policy.
openGauss=# ALTER TABLE all_data ENABLE ROW LEVEL SECURITY;
```

```
-- Create a row-level security policy to specify that the current user can view only their own data.
openGauss=# CREATE ROW LEVEL SECURITY POLICY all_data_rls ON all_data USING(role =
CURRENT_USER);

-- View table details.
openGauss=# \d+ all_data
Table "public.all_data"
Column |      Type      | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
id      | integer         |           |         |              |
role    | character varying(100) |         | extended |              |
data    | character varying(100) |         | extended |              |
Row Level Security Policies:
  POLICY "all_data_rls"
    USING (((role)::name = "current_user"()))
Has OIDs: no
Location Nodes: ALL DATANODES
Options: orientation=row, compression=no, enable_rowsecurity=true

-- Switch to user alice and run SELECT * FROM public.all_data.
openGauss=# SELECT * FROM public.all_data;
id | role | data
---+---+---
1  | alice | alice data
(1 row)

openGauss=# EXPLAIN(COSTS OFF) SELECT * FROM public.all_data;
QUERY PLAN
-----
Streaming (type: GATHER)
Node/s: All datanodes
-> Seq Scan on all_data
   Filter: ((role)::name = 'alice'::name)
Notice: This query is influenced by row level security feature
(5 rows)

-- Switch to user peter and run SELECT * FROM public.all_data.
openGauss=# SELECT * FROM public.all_data;
id | role | data
---+---+---
3  | peter | peter data
(1 row)

openGauss=# EXPLAIN(COSTS OFF) SELECT * FROM public.all_data;
QUERY PLAN
-----
Streaming (type: GATHER)
Node/s: All datanodes
-> Seq Scan on all_data
   Filter: ((role)::name = 'peter'::name)
Notice: This query is influenced by row level security feature
(5 rows)
```

NOTICE

PG_STATISTIC and **PG_STATISTIC_EXT** store sensitive information about statistical objects, such as high-frequency MCVs. If the permission to query the two system catalogs is granted to a common user after the row-level security policy is created, the common user can still access the two system catalogs to obtain the information in the statistical objects.

2.2 Database Audit

Context

Database security is critical to the database system. GaussDB writes all user operations on the database into audit logs. Database security administrators can use the audit logs to reproduce a series of events that cause faults in the database and identify unauthorized users, unauthorized operations, and the time when these operations are performed.

You need to know the following about the audit function:

- The overall audit parameter **audit_enabled** supports dynamic loading. If you change the value of this configuration when the database is running, the change takes effect immediately and you do not need to restart the database. The default value is **on**, indicating that the audit function is enabled.
- In addition to the overall audit configuration, each audit item can be independently configured. The function of each audit item takes effect only after the configuration is enabled.
- The configuration of each audit item supports dynamic loading. After changing the audit switch status when the database is running, the modification takes effect immediately without restarting the database.

Table 2-4 describes the audit items supported by GaussDB.

Table 2-4 Audit items

Parameter	Description
User login and logout audit	Parameter: audit_login_logout Its default value is 7 , which indicates that the function of user login and logout audit is enabled. 0 indicates that the function of user login and logout audit is disabled. Other values are not recommended.
Database startup, stop, recovery, and switchover audit	Parameter: audit_database_process Its default value is 1 , which indicates that the audit of database startup, stop, recovery, and switchover is enabled.
User locking and unlocking audit	Parameter: audit_user_locked Its default value is 1 , which indicates that the audit of user locking and unlocking is enabled.
Unauthorized access audit	Parameter: audit_user_violation Its default value is 0 , which indicates that the audit of unauthorized access is disabled.
Permission granting and revoking audit	Parameter: audit_grant_revoke Its default value is 1 , which indicates that the audit of permission granting and revoking is enabled.

Parameter	Description
Audit of CREATE, ALTER, and DROP operations on database objects	Parameter: audit_system_object Its default value is 67121159 , only the CREATE, ALTER, and DROP operations on databases, schemas, users, data sources, and SQL patches are audited.
Audit of INSERT, UPDATE, and DELETE operations on a specific table	Parameter: audit_dml_state Its default value is 0 , which indicates that the audit of DML operations (except SELECT) on a specific table is disabled.
SELECT operation audit	Parameter: audit_dml_state_select Its default value is 0 , which indicates that the audit of the SELECT operation is disabled.
COPY operation audit	Parameter: audit_copy_exec Its default value is 1 , which indicates that the audit of COPY operations is enabled.
Stored procedure and user-defined function execution audit	Parameter: audit_function_exec The default value is 0 , which indicates that no execution audit logs of stored procedures and user-defined functions are recorded.
SET operation audit	Parameter: audit_set_parameter Its default value is 0 , which indicates that the audit of the SET operation is disabled.
Transaction ID record	Parameter: audit_xid_info Its default value is 0 , which indicates that the function of recording transaction IDs in audit logs is disabled.

3 Database Quick Start

3.1 Connecting to a Database

Client tools used for connecting to a database include **gsql** and APIs (such as **ODBC** and **JDBC**).

- **gsql** is a client tool provided by GaussDB. As described in [Using gsql to Connect to a Database](#), **psql** is used to enter, edit, and run SQL statements in an interactive manner.
- As described in [APIs](#), standard databases, such as **ODBC** and **JDBC**, can be used to develop GaussDB-based applications.

3.1.1 Using gsql to Connect to a Database

gsql provided by GaussDB is a database connection tool running in the CLI. **gsql** provides basic and advanced functions of databases to facilitate user operations. This section describes how to use **gsql** to connect to a database. For details about how to use **gsql**, see "Client Tools > gsql" in *Tool Reference*.

Precautions

By default, when a client has been idle for the period specified by **session_timeout** after connecting to a database, the client automatically disconnects from the database. To disable the timeout setting, set **session_timeout** to 0.

Prerequisites

You have contacted the administrator for connection information.

Remotely Connecting to a Database

- Step 1** Contact an administrator to configure the remote connection.
- Step 2** On the host, upload the client tool package and configure environment variables for the **gsql** client. The openEuler environment is used as an example.

1. Log in to the client.
2. Run the following command to create the **/tmp/tools** directory:
`mkdir /tmp/tools`
3. Obtain **GaussDB-Kernel_VxxxRxxxCxx-XXXXX-64bit-gsql.tar.gz** from the software installation package and upload it to the **/tmp/tools** directory.

NOTE

- The software package is located where you put it before installation. Set it based on actual situation.
 - The tool package name may vary in different OSs. Select the tool package suitable for your OS.
4. Run the following commands to decompress the package:
`cd /tmp/tools`
`tar -zxvf GaussDB-Kernel_VxxxRxxxCxx-XXXXX-64bit-gsql.tar.gz`
 5. Log in to the server where the primary database node is located, and copy the **bin** directory in the database installation directory to **/tmp/tools** on the client host.
`scp -r /opt/huawei/install/app/bin root@10.10.0.30:/tmp/tools`
In the preceding command, **/opt/huawei/install/app** indicates the *{gaussdbAppPath}* path configured in the **clusterconfig.xml** file, and **10.10.0.30** indicates the IP address of the client host.
 6. Log in to the host where the client is installed and set environment variables.
Run the following command to open the **~/.bashrc** file:
`vi ~/.bashrc`
Enter the following content and run **:wq!** to save and exit:
`export PATH=/tmp/tools/bin:$PATH`
`export LD_LIBRARY_PATH=/tmp/tools/lib:$LD_LIBRARY_PATH`
 7. Run the following command to make the environment variables take effect:
`source ~/.bashrc`

Step 3 Connect to a database.

After the database is installed, a database named **postgres** is generated by default. When connecting to a database for the first time, you can connect to this database.

```
gsql -d postgres -h 10.10.0.11 -U jack -p 8000
Password for user jack:
```

postgres is the name of the database, **10.10.0.11** is the IP address of the server where the primary node of the database resides, **jack** is the user for connecting to the database, and **8000** is the port number of the primary database node.

In a centralized database instance, when you need to connect to the primary node and the IP addresses of three nodes in the database instance are 10.10.0.11, 10.10.0.12, and 10.10.0.13, you can use the **gsql -d postgres -h 10.10.0.11,10.10.0.12,10.10.0.13 -U jack -p 8000** command. **gsql** connects to the three IP addresses in sequence. If the current IP address is not the IP address of the primary node, **gsql** disconnects from the current IP address and attempts to connect to the next IP address, until the primary node is found.

 NOTE

- If a machine connected to GaussDB is not in the same network segment as GaussDB, the IP address specified by **-h** should be the value of **coo.cooListenIp2** (application access IP address) set in Manager.
- By default, the initial user of the database is not allowed to remotely connect to the database. If Kerberos authentication is enabled in an instance, the initial user is allowed to remotely connect to the database in the instance.

----End

3.1.2 APIs

You can use standard database APIs, such as **ODBC** and **JDBC**, to develop GaussDB-based applications.

Supported APIs

Each application is an independent GaussDB development project. APIs alleviate applications from directly operating in databases, and enhance the database portability, extensibility, and maintainability. [Table 3-1](#) lists the APIs supported by GaussDB and the download addresses.

Table 3-1 Database APIs

API	How to Obtain
ODBC	• Linux: Driver: GaussDB-Kernel_VxxxRxxxCxx-xxxxx-64bit-Odbc.tar.gz unixODBC source code package: https://sourceforge.net/projects/unixodbc/files/unixODBC/2.3.7/ • Windows: Driver: GaussDB-Kernel_VxxxRxxxCxx-Windows-Odbc.tar.gz
JDBC	• Driver: GaussDB-Kernel_VxxxRxxxCxx-xxxxx-64bit-Jdbc.tar.gz • Driver: com.huawei.opengauss.jdbc.Driver

The JDBC and ODBC APIs are used to connect to the database remotely. Therefore, you need to configure a remote connection in GaussDB. Contact an administrator for related operations.

For details about more APIs, see [Application Development Guide](#).

3.2 Operating a Database

This section explains how to use databases, including creating database accounts, databases, and tables, inserting data to tables, and querying data in tables.

3.2.1 Creating a Database Account

Only administrators that are created during database installation can access the initial database by default. You can also create other database accounts.

```
openGauss=# CREATE USER joe WITH PASSWORD 'xxxxxxxxx';
```

If the following information is displayed, the creation is successful:

```
CREATE ROLE
```

In this case, you have created a user account named **joe**, and the user password is **xxxxxxxxx**.

Run the following command to set user **joe** as the system administrator:

```
openGauss=# GRANT ALL PRIVILEGES TO joe;
```

Run the **GRANT** command to set related permissions. For details, see [GRANT](#).

NOTE

Note: For details about how to create database accounts, see [Users and Permissions](#).

3.2.2 Creating and Managing Databases

Prerequisites

Only the database system administrators or users granted with database creation permissions can create a database. For details about how to grant database creation permissions to a user, see [Users and Permissions](#).

Context

- GaussDB has two default template databases **template0** and **template1** and a default user database **postgres**. The default compatible database type of Postgres is O (that is, **DBCOMPATIBILITY** is set to **A**). In this compatible type, empty strings are processed as null values.
- CREATE DATABASE** creates a database by copying a template database (**template0** by default). Do not use a client or any other tools to connect to or to perform operations on the template databases.

NOTE

- The template database does not contain any user table. You can view the attributes of the template database in the **PG_DATABASE** system catalog.
- The **template0** template does not allow user connections. Only the initial user and system administrators can connect to **template1**.
- A maximum of 128 databases can be created in GaussDB.
- A database system consists of multiple databases. A client can connect to only one database at a time. Users cannot query data across databases. If one GaussDB contains multiple databases, set the **-d** parameter to specify the database instance to be connected.

Precautions

Assume that the database encoding is SQL_ASCII. (You can run the **show server_encoding;** command to query the encoding used for storing data in the current database.) If the database object name contains multi-byte characters (such as Chinese) or if the object name length exceeds the allowed maximum (63 bytes), the database truncates the last byte (not the last character) of the object name. In this case, half characters may appear.

To resolve this problem, you need to:

- Ensure that the name of the data object does not exceed the maximum length.
- Change the default database storage code set (**server_encoding**) to UTF-8.
- Exclude multi-byte characters from object names.
- Ensure that no more than 128 databases are created (recommended).
- If you fail to delete an object by specifying its name after truncation, specify its original name to delete it, or manually delete it from the system catalogs on each node.

Procedure

Step 1 Run the following command to create a database named **db_tpcc**:

```
openGauss=# CREATE DATABASE db_tpcc;  
CREATE DATABASE
```

NOTE

- Database names must comply with the general naming convention rules of SQL identifiers. The current role automatically becomes the owner of this new database.
- If a database system is used to support independent users and projects, store them in different databases.
- If the projects or users are associated with each other and share resources, store them in one database. However, you can divide them into different schemas. A schema is a logical structure, and the access permission for a schema is controlled by the permission system module.
- A database name contains a maximum of 63 bytes and the excessive bytes at the end of the name will be truncated by the server. You are advised to specify a database name no longer than 63 bytes when you create a database.

- New databases are created in the **pg_default** tablespace by default. To specify another tablespace, run the following statement:

```
openGauss=# CREATE DATABASE db_tpcc WITH TABLESPACE = hr_local;  
CREATE DATABASE
```

hr_local indicates the tablespace name. For details about how to create a tablespace, see [Creating and Managing Tablespace](#).

- After creating the **db_tpcc** database, you can perform other operations in the default **postgres** database. Alternatively, you can perform the following operations to exit the **postgres** database, connect to the **db_tpcc** database as a new user, and perform operations such as creating tables.

```
openGauss=# \q  
gsqll -d db_tpcc -p 8000 -U joe  
Password for user joe:  
gsqll ((GaussDB Kernel 503.1.XXX build f521c606) compiled at 2021-09-16 14:55:22 commit 2935  
last mr 6385 release)  
Non-SSL connection (SSL connection is recommended when requiring high-security)  
Type "help" for help.  
db_tpcc=>
```

Step 2 View databases.

- Run the **\l** meta-command to view the database list of the database system.
openGauss=# \l
- Run the following command to query the database list in the **pg_database** system catalog:

```
openGauss=# SELECT datname FROM pg_database;
```

Step 3 Modify the database.

You can modify database configuration such as the database owner, name, and default settings.

- Run the following command to set the default search path for the database:

```
openGauss=# ALTER DATABASE db_tpcc SET search_path TO pa_catalog,public;  
ALTER DATABASE
```
- Run the following command to rename the database:

```
openGauss=# ALTER DATABASE db_tpcc RENAME TO human_tpcds;  
ALTER DATABASE
```

Step 4 Delete the database.

You can run the **DROP DATABASE** command to delete a database. This command deletes the system catalog of the database and the database directory on the disk. Only the database owner or system administrators can delete a database. A database accessed by users cannot be deleted. You need to connect to another database before deleting this database.

Run the following command to delete the database:

```
openGauss=# DROP DATABASE human_tpcds;  
DROP DATABASE
```

----End

3.2.3 Creating and Managing Tablespaces

Background

The administrator can use tablespaces to control the layout of disks where a database is installed. This has the following advantages:

- If the disk partition or tablespace initially allocated to the database is full and the space cannot be logically extended, you can create and use tablespaces in other partitions until the system space is reconfigured.
- Tablespaces allow the administrator to distribute data based on the schema of database objects, improving system performance.
 - A frequently used index can be placed in a disk having stable performance and high computing speed, such as a solid device.
 - A table that stores archived data and is rarely used or has low performance requirements can be placed in a disk with a slow computing speed.
- The administrator can use tablespaces to set the maximum available disk space. In this way, when a partition is shared with other data, tablespaces will not occupy excessive space in the partition.
- You can use tablespaces to control the disk space occupied by data in a database. If the usage of a disk where a tablespace resides reaches 90%, the database switches to the read-only mode. It switches back to read/write mode when the disk usage becomes less than 90%.

You are advised to use the background monitoring program or Database Manager to monitor the disk space usage when using the database to prevent the database from switching to the read-only mode.

- Each tablespace corresponds to a file system directory. If *data directory / pg_location/mount1/path1* is an empty directory for which users have read

and write permissions, an absolute path tablespace can be created in this directory.

If the tablespace quota management is used, the performance may deteriorate by about 30%. **MAXSIZE** specifies the maximum quota for each database node. The deviation must be within 500 MB. Determine whether to set a tablespace to its maximum size as required.

GaussDB provides two tablespaces: **pg_default** and **pg_global**.

- Default tablespace **pg_default**: stores non-shared system tables, user tables, user table indexes, temporary tables, temporary table indexes, and internal temporary tables. The corresponding storage directory is the base directory in the instance data directory.
- Shared tablespace **pg_global**: stores shared system tables. The corresponding storage directory is the base directory in the global data directory.

Precautions

You are not advised to use user-defined tablespaces. This is because user-defined tablespaces are usually used with storage media other than the main storage (storage device where the default tablespace is located, such as a disk) to isolate I/O resources that can be used by different services. Storage devices use standard configurations and do not have other available storage media. If the user-defined tablespace is not properly used, the system cannot run stably for a long time and the overall performance is affected. Therefore, you are advised to use the default tablespace.

Procedure

- Create a tablespace.
 - a. Run the following command to create user **jack** and set the password to **XXXXXXXX**:
openGauss=# **CREATE USER jack IDENTIFIED BY 'XXXXXXXX';**
If the following information is displayed, the creation is successful:
CREATE ROLE
 - b. Run the following command to create a tablespace:
openGauss=# **CREATE TABLESPACE fastspace RELATIVE LOCATION 'tablespace/tablespace_1';**
If the following information is displayed, the creation is successful:
CREATE TABLESPACE
fastspace is the new tablespace, and *Database node data directory/pg_location/tablespace/tablespace_1* is an empty directory for which users have read and write permissions.
 - c. A database system administrator can run the following command to grant the permission of accessing the **fastspace** tablespace to user **jack**:
openGauss=# **GRANT CREATE ON TABLESPACE fastspace TO jack;**
If the following information is displayed, the grant operation is successful:
GRANT
- Create an object in a tablespace.

If you have the CREATE permission on the tablespace, you can create database objects in the tablespace, such as tables and indexes.

Take creating a table as an example:

- Method 1: Run the following command to create a table in a specified tablespace:

```
openGauss=# CREATE TABLE foo(i int) TABLESPACE fastspace;
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

- Method 2: Use **set default_tablespace** to set the default tablespace and then create a table:

```
openGauss=# SET default_tablespace = 'fastspace';
```

```
SET
```

```
openGauss=# CREATE TABLE foo2(i int);
```

```
CREATE TABLE
```

In this example, **fastspace** is the default tablespace, and **foo2** is the created table.

- Use one of the following methods to query a tablespace:

- Method 1: Check the **pg_tablespace** system catalog. Run the following command to view all the tablespaces defined by the system and users:

```
openGauss=# SELECT spcname FROM pg_tablespace;
```

- Method 2: Run the following meta-command of the **gsql** program to query the tablespaces:

```
openGauss=# \db
```

- Query the tablespace usage.

- a. Query the current usage of the tablespace.

```
openGauss=# SELECT PG_TABLESPACE_SIZE('example');
```

The following information is displayed:

```
pg_tablespace_size
-----
                2146304
(1 row)
```

2146304 is the size of the tablespace, and its unit is byte.

- b. Calculate the tablespace usage.

Tablespace usage rate = **PG_TABLESPACE_SIZE**/Size of the disk where the tablespace resides

- Modify a tablespace.

Run the following command to rename tablespace **fastspace** to **fspace**:

```
openGauss=# ALTER TABLESPACE fastspace RENAME TO fspace;
```

```
ALTER TABLESPACE
```

- Delete a tablespace and related data.

- Run the following command to delete user **jack**:

```
openGauss=# DROP USER jack CASCADE;
```

```
DROP ROLE
```

- Run the following commands to delete tables **foo** and **foo2**:

```
openGauss=# DROP TABLE foo;
```

```
openGauss=# DROP TABLE foo2;
```

If the following information is displayed, the deletion is successful:

```
DROP TABLE
```

- Run the following command to delete tablespace **fspace**:

```
openGauss=# DROP TABLESPACE fspace;
```

```
DROP TABLESPACE
```

NOTE

Only the tablespace owner or system administrator can delete a tablespace.

3.2.4 Creating and Managing Tables

3.2.4.1 Creating Tables

Context

A table is created in a database and can be saved in different databases. Tables under different schemas in a database can have the same name.

Procedure

Run the following statement to create a table:

```
openGauss=# CREATE TABLE customer_t1
(
  c_customer_sk      integer,
  c_customer_id      char(5),
  c_first_name       char(6),
  c_last_name        char(8)
);
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

c_customer_sk, **c_customer_id**, **c_first_name**, and **c_last_name** are the column names of the table. **integer**, **char(5)**, **char(6)**, and **char(8)** are column name types.

NOTE

- By default, new database objects, such as the **mytable** table, are created in the **\$user** schema. For more details about schemas, see [Creating and Managing Schemas](#).
- In addition to the created tables, a database contains many system catalogs. These system catalogs contain database installation information and information about various queries and processes in GaussDB. You can collect information about the database by querying system catalogs. For details, see [Querying a System Catalog](#).

3.2.4.2 Inserting Data to Tables

A new table contains no data. You need to insert data to the table before using it. This section describes how to insert a row or multiple rows of data by running the **INSERT** command and to insert data from a specified table. Contact an administrator if a large amount of data needs to be imported.

Context

The length of a character on the server and client may vary by the character sets they use. A string entered on the client will be processed based on the server's character set, so the result may differ from expected.

Table 3-2 Comparison of character set output between the client and server

Procedure	Server and Client Use Same Encoding	Server and Client Use Different Encoding
No operations are performed to the string while it is saved and read.	Your expected result is returned.	If the encoding for input and output on the client is the same, your expected result is returned.
Operations (such as executing string functions) are performed to the string while it is saved and read.	Your expected result is returned.	The result may differ from expected, depending on the operations performed to the string.
A long string is truncated while it is saved.	Your expected result is returned.	If the character sets used on the client and server have different character length, the result may differ from expected.

More than one of the preceding operations can be performed to a string. For example, if the character sets of the client and server are different, a string may be processed and then truncated. In this case, the result will also be unexpected.

NOTE

Long strings are truncated only if **DBCOMPATIBILITY** is set to **TD** (compatible with Teradata) and **td_compatible_truncation** is set to **on**.

Procedure

You need to create a table before inserting data to it. For details about how to create a table, see [Creating and Managing Tables](#).

- Insert a row to table **customer_t1**.

Data values are arranged in the same order as the columns in the table and are separated by commas (,). Generally, column values are text values (constants). But column values can also be scalar expressions.

```
openGauss=# INSERT INTO customer_t1(c_customer_sk, c_customer_id, c_first_name) VALUES (3769, 'hello', 'Grace');
```

If you know the sequence of the columns in the table, you can obtain the same result without listing these columns. For example, the following command generates the same result as the preceding command:

```
openGauss=# INSERT INTO customer_t1 VALUES (3769, 'hello', 'Grace');
```

If you do not know some of the column values, you can omit them. In the INSERT statement, if the column name of the specified target table is not displayed, the values to be inserted in the VALUES clause correspond to the columns of the target table based on the column number. That is, the first value in the VALUES clause corresponds to the first column of the target table, the second value of the VALUES clause corresponds to the second

column of the target table, and so on. Columns that do not have corresponding values in the VALUES clause are automatically filled with default values or NULL. Example:

```
openGauss=# INSERT INTO customer_t1 (c_customer_sk, c_first_name) VALUES (3769, 'Grace');
```

```
openGauss=# INSERT INTO customer_t1 VALUES (3769, 'hello');
```

You can also specify the default value of a column or row:

```
openGauss=# INSERT INTO customer_t1 (c_customer_sk, c_customer_id, c_first_name) VALUES (3769, 'hello', DEFAULT);
```

```
openGauss=# INSERT INTO customer_t1 DEFAULT VALUES;
```

- To insert multiple rows to a table, run the following command:

```
openGauss=# INSERT INTO customer_t1 (c_customer_sk, c_customer_id, c_first_name) VALUES (6885, 'maps', 'Joes'), (4321, 'tpcds', 'Lily'), (9527, 'world', 'James');
```

You can also insert multiple rows by running the command for inserting one row for multiple times. However, you are advised to run this command to improve efficiency.

- Assume that you have created a backup table **customer_t2** for table **customer_t1**. To insert data from **customer_t1** to **customer_t2**, run the following statements:

```
openGauss=# CREATE TABLE customer_t2 (
  c_customer_sk      integer,
  c_customer_id      char(5),
  c_first_name       char(6),
  c_last_name        char(8)
);
```

```
openGauss=# INSERT INTO customer_t2 SELECT * FROM customer_t1;
```

NOTE

If implicit conversion is not implemented between the column data types of the specified table and those of the current table, the two tables must have the same column data types when data is inserted from the specified table to the current table.

- To delete a backup table, run the following command:

```
openGauss=# DROP TABLE customer_t2 CASCADE;
```

NOTE

If the table to be deleted is in dependent relationship with other tables, you need to delete its dependent tables first.

3.2.4.3 Updating Data in a Table

Existing data in a database can be updated. You can update one row, all rows, or specified rows of data, or update data in a single column without affecting the data in the other columns.

The following types of information are required when the **UPDATE** statement is used to update a row:

- Table name and column name of the data to be updated
- New column value
- Rows of the data to be updated

Generally, the SQL language does not provide a unique ID for a row of data. Therefore, it is impossible to directly specify the rows of the data to be updated.

However, you can specify the rows by declaring the conditions that must be met by the updated row. If a table contains primary keys, you can specify a row using the primary keys.

For details about how to create a table and insert data to it, see [Creating Tables](#) and [Inserting Data to Tables](#).

c_customer_sk in the table **customer_t1** must be changed from **9527** to **9876**:

```
openGauss=# UPDATE customer_t1 SET c_customer_sk = 9876 WHERE c_customer_sk = 9527;
```

You can use a schema to modify the table name. If no such modifier is specified, the table is located based on the default schema path. In the statement, **SET** is followed by the target column and the new column value. The new value can be a constant or an expression.

For example, run the following statement to increase all the values in the **c_customer_sk** column by 100:

```
openGauss=# UPDATE customer_t1 SET c_customer_sk = c_customer_sk + 100;
```

This statement does not include the **WHERE** clause, so all rows are updated. If the statement includes the **WHERE** clause, only the rows matching the clause are updated.

In the **SET** clause, the equal sign (=) indicates value setting. In the **WHERE** clause, the equal sign indicates comparison. **WHERE** may not represent an equation and can be replaced by other operators.

You can run an **UPDATE** statement to update multiple columns by specifying multiple values in the **SET** clause. For example:

```
openGauss=# UPDATE customer_t1 SET c_customer_id = 'Admin', c_first_name = 'Local' WHERE c_customer_sk = 4421;
```

After data has been updated or deleted in batches, a large number of deletion markers are generated in the data file. During query, data with these deletion markers needs to be scanned as well. In this case, a large amount of data with deletion marks can greatly affect the query performance after batch updates or deletions. If data needs to be updated or deleted in batches frequently, you are advised to periodically run the **VACUUM FULL** statement to maintain the query performance.

3.2.4.4 Viewing Data

- Run the following command to query information about all tables in a database in the system catalog **pg_tables**:

```
openGauss=# SELECT * FROM pg_tables;
```
- Run the **\d+** command of the **gsql** tool to query table attributes:

```
openGauss=# \d+ customer_t1;
```
- Run the following command to query the data volume of table **customer_t1**:

```
openGauss=# SELECT count(*) FROM customer_t1;
```
- Run the following command to query all data in the table **customer_t1**:

```
openGauss=# SELECT * FROM customer_t1;
```
- Run the following command to query only the data in the column **c_customer_sk**:

```
openGauss=# SELECT c_customer_sk FROM customer_t1;
```
- Run the following command to filter repeated data in the column **c_customer_sk**:

```
openGauss=# SELECT DISTINCT( c_customer_sk ) FROM customer_t1;
```

- Run the following command to query all data whose column **c_customer_sk** is **3869**:

```
openGauss=# SELECT * FROM customer_t1 WHERE c_customer_sk = 3869;
```

- Run the following command to collate data based on the column **c_customer_sk**:

```
openGauss=# SELECT * FROM customer_t1 ORDER BY c_customer_sk;
```

3.2.4.5 Deleting Data from a Table

Outdated data may need to be deleted when tables are used. Data can be deleted from tables only by row.

SQL statements can only access and delete an independent row by declaring conditions that match the row. If a table has a primary key, you can use it to specify a row. You can delete several rows that match the specified condition or delete all the rows from a table.

For example, to delete all the rows whose **c_customer_sk** column is **3869** from the table **customer_t1**, run the following command:

```
openGauss=# DELETE FROM customer_t1 WHERE c_customer_sk = 3869;
```

To delete all rows from the table, run either of the following commands:

```
openGauss=# DELETE FROM customer_t1;
```

Or

```
openGauss=# TRUNCATE TABLE customer_t1;
```

NOTE

If you need to delete an entire table, you are advised to use the **TRUNCATE** statement rather than **DELETE**.

To delete a table, run the following command:

```
openGauss=# DROP TABLE customer_t1;
```

3.2.5 Querying a System Catalog

In addition to the created tables, a database contains many system catalogs. These system catalogs contain database installation information and information about various queries and processes in GaussDB. You can collect information about the database by querying system catalogs.

In [System Catalogs and System Views](#), the description about each table specifies whether the table is visible to all users or only the initial user. To query tables that are visible only to the initial user, log in as the user.

GaussDB provides the following types of system catalogs and system views:

- PG-compatible system catalogs and views
These system catalogs and views have the prefix **PG**.
- New system catalogs and system views of GaussDB
These system catalogs and views have the prefix **GS**.

Querying Database Tables

Create the following tables in the public schema:

```
openGauss=# CREATE TABLE public.search_table_t1(a int);
CREATE TABLE
openGauss=# CREATE TABLE public.search_table_t2(b int);
CREATE TABLE
openGauss=# CREATE TABLE public.search_table_t3(c int);
CREATE TABLE
openGauss=# CREATE TABLE public.search_table_t4(d int);
CREATE TABLE
openGauss=# CREATE TABLE public.search_table_t5(e int);
CREATE TABLE
```

In the PG_TABLES system catalog, view the tables prefixed with search_table in the public schema.

```
openGauss=# SELECT distinct(tablename) FROM pg_tables WHERE SCHEMANAME = 'public' AND
TABLENAME LIKE 'search_table%';
```

Information similar to the following is displayed:

```
tablename
-----
search_table_t1
search_table_t2
search_table_t3
search_table_t4
search_table_t5
(5 rows)
```

Viewing Database Users

You can run the **PG_USER** command to view the list of all users in the database, and view the user ID (**USESYSID**) and permissions.

```
SELECT * FROM pg_user;
username | usesysid | usecreatedb | usesuper | usecatupd | userepl | passwd | valbegin | valuntil |
respool  | parent  | spacelimit | useconfig | no        | degroup | tempspacelimit | spillspacelimit | usemonitoradmin | useoperatoradmin | usepolicyadmin
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
omm      | 10      | t          | t        | t        | t       | ***** |               |               | default_pool | 0
|         |          |          |          |          |          |          |               |               |               |
|         |          |          | t        | t        | t       |          |               |               |               |
```

Viewing and Stopping the Running Query Statements

You can view the running query statements in the **PG_STAT_ACTIVITY** view. You can use the following methods:

Step 1 Set the parameter **track_activities** to **on**.

```
SET track_activities = on;
```

The database collects the running information about active queries only if the parameter is set to **on**.

Step 2 View the running query statements. Run the following command to view the database names, users performing queries, query status, and the corresponding PIDs which are connected to the running query statements:

```
SELECT datname, username, state, pid FROM pg_stat_activity;
datname | username | state | pid
-----+-----+-----+-----
testdb  | Ruby    | active | 140298793514752
```

```
testdb | Ruby | active | 140298718004992
testdb | Ruby | idle | 140298650908416
testdb | Ruby | idle | 140298625742592
testdb | omm | active | 140298575406848
(5 rows)
```

If the **state** column is **idle**, the connection is idle and requires a user to enter a command.

To identify only active query statements, run the following command:

```
SELECT datname, username, state, pid FROM pg_stat_activity WHERE state != 'idle';
```

Step 3 To cancel queries that have been running for a long time, use the **PG_TERMINATE_BACKEND** function to end sessions based on the thread ID.

```
SELECT PG_TERMINATE_BACKEND(139834759993104);
```

If information similar to the following is displayed, the session is successfully terminated:

```
PG_TERMINATE_BACKEND
-----
t
(1 row)
```

If information similar to the following is displayed, a user has terminated the current session:

```
FATAL: terminating connection due to administrator command
FATAL: terminating connection due to administrator command
```

NOTE

If the **PG_TERMINATE_BACKEND** function is used to terminate the backend threads of the current session, the **gsql** client will be reconnected automatically rather than be logged out. The message "The connection to the server was lost. Attempting reset: Succeeded." is returned.

```
FATAL: terminating connection due to administrator command
FATAL: terminating connection due to administrator command
The connection to the server was lost. Attempting reset: Succeeded.
```

----End

3.2.6 Other Operations

3.2.6.1 Creating and Managing Schemas

Context

Schemas function as models. Schema management allows multiple users to use the same database without mutual impacts, to organize database objects as manageable logical groups, and to add third-party applications to the same schema without causing conflicts. Schema management involves creating a schema, using a schema, deleting a schema, setting a search path for a schema, and setting schema permissions.

Precautions

- GaussDB contains one or more named databases. Users and user groups are shared within the database, but their data is not shared. Any user who has

connected to a server can access only the database specified in the connection request.

- A database can have one or more schemas, and a schema can contain tables and other data objects, such as data types, functions, and operators. One object name can be used in different schemas. For example, both **schema1** and **schema2** can have a table named **mytable**.
- Different from databases, schemas are not isolated. You can access the objects in a schema of the connected database if you have schema permissions. To manage schema permissions, you need to have knowledge about database permissions.
- A schema named with the **PG_** prefix cannot be created because this type of schema is reserved for the database system.
- Each time a new user is created, the system creates a schema with the same name for the new user in the current database. In other databases, such a schema needs to be manually created.
- To reference a table that is not modified with a schema name, the system uses **search_path** to find the schema that the table belongs to. **pg_temp** and **pg_catalog** are always the first two schemas to be searched no matter whether or how they are specified in **search_path**. **search_path** is a schema name list, and the first table detected in it is the target table. If no target table is found, an error will be reported. (If a table exists but the schema it belongs to is not listed in **search_path**, the search fails as well.) The first schema in **search_path** is called "current schema". This schema is the first one to be searched. If no schema name is declared, newly created database objects are saved in this schema by default.
- Each database has a **pg_catalog** schema, which contains system catalogs and all embedded data types, functions, and operators. **pg_catalog** is a part of the search path and has the second highest search priority. It is searched after the schema of temporary tables and before other schemas specified in **search_path**. This search order ensures that database built-in objects can be found. To use a custom object that has the same name as a built-in object, you can specify the schema of the custom object.

Procedure

- Create a management user and permission schema.
 - Run the following command to create a schema:

```
openGauss=# CREATE SCHEMA myschema;
```

If the following information is displayed, the schema named **myschema** is successfully created:

```
CREATE SCHEMA
```

To create or access an object in the schema, the object name in the command should consist of the schema name and the object name, which are separated by a dot (.), for example, **myschema.table**.
 - Run the following command to create a schema and specify the owner:

```
openGauss=# CREATE SCHEMA myschema AUTHORIZATION omm;
```

If the following information is displayed, the **myschema** schema that belongs to **omm** is created successfully:

```
CREATE SCHEMA
```
- Use a schema.

If you want to create or access an object in a specified schema, the object name must contain the schema name. To be specific, the name consists of a schema name and an object name, which are separated by a dot (.).

- Run the following command to create the **mytable** table in **myschema**:

```
openGauss=# CREATE TABLE myschema.mytable(id int, name varchar(20));
CREATE TABLE
```

To specify the location of an object, the object name must contain the schema name.

- Run the following command to query all data of the **mytable** table in **myschema**:

```
openGauss=# SELECT * FROM myschema.mytable;
id | name
----+-----
(0 rows)
```

- View the search path of a schema.

You can set **search_path** to specify the sequence of schemas in which objects are searched. The first schema listed in the search path will become the default schema. If no schema is specified during object creation, the object will be created in the default schema.

- Run the following command to view the search path:

```
openGauss=# SHOW SEARCH_PATH;
search_path
-----
"$user",public
(1 row)
```

- Run the following command to set the search path to **myschema,public** (**myschema** will be searched before **public**):

```
openGauss=# SET SEARCH_PATH TO myschema, public;
SET
```

- Set permissions for a schema.

By default, a user can only access database objects in their own schema. Only after a user is granted with the usage permission for a schema by the schema owner, the user can access the objects in the schema.

By granting the **CREATE** permission for a schema to a user, the user can create objects in this schema. By default, all roles have the **USAGE** permission in the **public** schema, but common users do not have the **CREATE** permission in the **public** schema. It is insecure for a common user to connect to a specified database and create objects in its **public** schema. If the common user has the **CREATE** permission on the **public** schema, it is advised to:

- Run the following command to revoke **PUBLIC**'s permission to create objects in the **public** schema. **public** indicates the schema and **PUBLIC** indicates all roles.

```
openGauss=# REVOKE CREATE ON SCHEMA public FROM PUBLIC;
REVOKE
```

- Run the following command to view the current schema:

```
openGauss=# SELECT current_schema();
current_schema
-----
myschema
(1 row)
```

- Run the following commands to create user **jack** and grant the **usage** permission for **myschema** to the user:

```
openGauss=# CREATE USER jack IDENTIFIED BY '*****';
CREATE ROLE
```

```
openGauss=# GRANT USAGE ON schema myschema TO jack;  
GRANT
```

- Run the following command to revoke the usage permission for **myschema** from **jack**:

```
openGauss=# REVOKE USAGE ON schema myschema FROM jack;  
REVOKE
```

- Delete a schema.

- If a schema is empty, that is, it contains no database objects, you can execute the **DROP SCHEMA** command to delete it. For example, run the following command to delete an empty schema named **nullschema**:

```
openGauss=# DROP SCHEMA IF EXISTS nullschema;  
DROP SCHEMA
```

- To delete a schema that is not null, use the keyword **CASCADE** to delete it and all its objects. For example, run the following command to delete **myschema** and all its objects in it:

```
openGauss=# DROP SCHEMA myschema CASCADE;  
DROP SCHEMA
```

- Run the following command to delete user **jack**:

```
openGauss=# DROP USER jack;  
DROP ROLE
```

3.2.6.2 Creating and Managing Partitioned Tables

Context

GaussDB supports range partitioned tables, interval partitioned tables, list partitioned tables, and hash partitioned tables.

- Range partitioned table: Data within a certain range is mapped to each partition. The range is determined by the partition key specified when the partitioned table is created. This partitioning mode is most commonly used. The partition key is usually a date. For example, sales data is partitioned by month.
- Interval partitioned table: A special type of range partitioned tables. Compared with range partitioned tables, interval value definition is added. When no matching partition can be found for an inserted record, a partition can be automatically created based on the interval value.
- List partitioned table: Key values contained in the data are stored in different partitions, and the data is mapped to each partition in sequence. The key values contained in the partitions are specified when the partitioned table is created.
- Hash partitioned table: Data is mapped to each partition based on the internal hash algorithm. The number of partitions is specified when the partitioned table is created.

A partitioned table has the following advantages over an ordinary table:

- High query performance: You can specify partitions when querying partitioned tables, improving query efficiency.
- High availability: If a certain partition in a partitioned table is faulty, data in the other partitions is still available.
- Easy maintenance: To fix a partitioned table having a faulty partition, you only need to fix the partition.

To convert an ordinary table to a partitioned table, you need to create a partitioned table and import data to it from the ordinary table. When you design tables, plan whether to use partitioned tables based on service requirements.

Procedure

Example 1: using the default tablespace

- Create a partitioned table (assuming that the **tpcds** schema has been created).

```
openGauss=# CREATE TABLE tpcds.customer_address
(
  ca_address_sk      integer          NOT NULL ,
  ca_address_id      character(16)    NOT NULL ,
  ca_street_number   character(10)    ,
  ca_street_name     character varying(60) ,
  ca_street_type     character(15)    ,
  ca_suite_number    character(10)    ,
  ca_city            character varying(60) ,
  ca_county          character varying(30) ,
  ca_state           character(2)     ,
  ca_zip            character(10)     ,
  ca_country         character varying(20) ,
  ca_gmt_offset      numeric(5,2)    ,
  ca_location_type   character(20)
)
PARTITION BY RANGE (ca_address_sk)
(
  PARTITION P1 VALUES LESS THAN(5000),
  PARTITION P2 VALUES LESS THAN(10000),
  PARTITION P3 VALUES LESS THAN(15000),
  PARTITION P4 VALUES LESS THAN(20000),
  PARTITION P5 VALUES LESS THAN(25000),
  PARTITION P6 VALUES LESS THAN(30000),
  PARTITION P7 VALUES LESS THAN(40000),
  PARTITION P8 VALUES LESS THAN(MAXVALUE)
)
ENABLE ROW MOVEMENT;
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

- Insert data.

Insert data from the **tpcds.customer_address** table to the **tpcds.web_returns_p2** table.

Suppose the backup table **tpcds.web_returns_p2** of the **tpcds.customer_address** table has been created in the database. You can run the following command to insert the data of the **tpcds.customer_address** table into the backup table **tpcds.web_returns_p2**:

```
openGauss=# CREATE TABLE tpcds.web_returns_p2
(
  ca_address_sk      integer          NOT NULL ,
  ca_address_id      character(16)    NOT NULL ,
  ca_street_number   character(10)    ,
  ca_street_name     character varying(60) ,
  ca_street_type     character(15)    ,
  ca_suite_number    character(10)    ,
  ca_city            character varying(60) ,
  ca_county          character varying(30) ,
  ca_state           character(2)     ,
  ca_zip            character(10)     ,
  ca_country         character varying(20) ,
  ca_gmt_offset      numeric(5,2)    ,
  ca_location_type   character(20)
)
```

```
)  
  
PARTITION BY RANGE (ca_address_sk)  
(  
    PARTITION P1 VALUES LESS THAN(5000),  
    PARTITION P2 VALUES LESS THAN(10000),  
    PARTITION P3 VALUES LESS THAN(15000),  
    PARTITION P4 VALUES LESS THAN(20000),  
    PARTITION P5 VALUES LESS THAN(25000),  
    PARTITION P6 VALUES LESS THAN(30000),  
    PARTITION P7 VALUES LESS THAN(40000),  
    PARTITION P8 VALUES LESS THAN(MAXVALUE)  
)  
  
ENABLE ROW MOVEMENT;  
CREATE TABLE  
openGauss=# INSERT INTO tpcds.web_returns_p2 SELECT * FROM tpcds.customer_address,  
INSERT 0 0
```

- Modify the row movement attributes of the partitioned table.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 DISABLE ROW MOVEMENT;  
ALTER TABLE
```

- Delete a partition.

Delete partition **P8**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 DROP PARTITION P8;  
ALTER TABLE
```

- Add a partition.

Add partition **P8** and set its range to [40000,MAXVALUE].

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 ADD PARTITION P8 VALUES LESS THAN  
(MAXVALUE);  
ALTER TABLE
```

- Rename a partition.

- Rename partition **P8** to **P_9**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 RENAME PARTITION P8 TO P_9;  
ALTER TABLE
```

- Rename partition **P_9** to **P8**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 RENAME PARTITION FOR (40000) TO P8;  
ALTER TABLE
```

- Query a partition.

Query partition **P6**.

```
openGauss=# SELECT * FROM tpcds.web_returns_p2 PARTITION (P6);  
openGauss=# SELECT * FROM tpcds.web_returns_p2 PARTITION FOR (35888);
```

- Delete a partitioned table and its tablespaces.

```
openGauss=# DROP TABLE tpcds.customer_address;  
DROP TABLE  
openGauss=# DROP TABLE tpcds.web_returns_p2;  
DROP TABLE
```

Example 2: using a user-defined tablespace

Perform the following operations on range partitioned tables.

- Creating tablespaces

```
openGauss=# CREATE TABLESPACE example1 RELATIVE LOCATION 'tablespace1/tablespace_1';  
openGauss=# CREATE TABLESPACE example2 RELATIVE LOCATION 'tablespace2/tablespace_2';  
openGauss=# CREATE TABLESPACE example3 RELATIVE LOCATION 'tablespace3/tablespace_3';  
openGauss=# CREATE TABLESPACE example4 RELATIVE LOCATION 'tablespace4/tablespace_4';
```

If the following information is displayed, the creation is successful:

```
CREATE TABLESPACE
```

- Create a partitioned table.

```
openGauss=# CREATE TABLE tpcds.customer_address  
(
```

```
ca_address_sk integer NOT NULL ,
ca_address_id character(16) NOT NULL ,
ca_street_number character(10) ,
ca_street_name character varying(60) ,
ca_street_type character(15) ,
ca_suite_number character(10) ,
ca_city character varying(60) ,
ca_county character varying(30) ,
ca_state character(2) ,
ca_zip character(10) ,
ca_country character varying(20) ,
ca_gmt_offset numeric(5,2) ,
ca_location_type character(20)
)
TABLESPACE example1
PARTITION BY RANGE (ca_address_sk)
(
PARTITION P1 VALUES LESS THAN(5000),
PARTITION P2 VALUES LESS THAN(10000),
PARTITION P3 VALUES LESS THAN(15000),
PARTITION P4 VALUES LESS THAN(20000),
PARTITION P5 VALUES LESS THAN(25000),
PARTITION P6 VALUES LESS THAN(30000),
PARTITION P7 VALUES LESS THAN(40000),
PARTITION P8 VALUES LESS THAN(MAXVALUE) TABLESPACE example2
)
ENABLE ROW MOVEMENT;
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

- Insert data.

Insert data from the **tpcds.customer_address** table to the **tpcds.web_returns_p2** table.

Suppose the backup table **tpcds.web_returns_p2** of the **tpcds.customer_address** table has been created in the database. You can run the following command to insert the data of the **tpcds.customer_address** table into the backup table **tpcds.web_returns_p2**:

```
openGauss=# CREATE TABLE tpcds.web_returns_p2
(
ca_address_sk integer NOT NULL ,
ca_address_id character(16) NOT NULL ,
ca_street_number character(10) ,
ca_street_name character varying(60) ,
ca_street_type character(15) ,
ca_suite_number character(10) ,
ca_city character varying(60) ,
ca_county character varying(30) ,
ca_state character(2) ,
ca_zip character(10) ,
ca_country character varying(20) ,
ca_gmt_offset numeric(5,2) ,
ca_location_type character(20)
)
TABLESPACE example1
PARTITION BY RANGE (ca_address_sk)
(
PARTITION P1 VALUES LESS THAN(5000),
PARTITION P2 VALUES LESS THAN(10000),
PARTITION P3 VALUES LESS THAN(15000),
PARTITION P4 VALUES LESS THAN(20000),
PARTITION P5 VALUES LESS THAN(25000),
PARTITION P6 VALUES LESS THAN(30000),
PARTITION P7 VALUES LESS THAN(40000),
PARTITION P8 VALUES LESS THAN(MAXVALUE) TABLESPACE example2
)
```

ENABLE ROW MOVEMENT;

```
CREATE TABLE  
openGauss=# INSERT INTO tpcds.web_returns_p2 SELECT * FROM tpcds.customer_address;  
INSERT 0 0
```

- Modify the row movement attributes of the partitioned table.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 DISABLE ROW MOVEMENT;  
ALTER TABLE
```

- Delete a partition.

Delete partition **P8**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 DROP PARTITION P8;  
ALTER TABLE
```

- Add a partition.

Add partition **P8** and set its range to [40000,MAXVALUE].

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 ADD PARTITION P8 VALUES LESS THAN  
(MAXVALUE);  
ALTER TABLE
```

- Rename a partition.

- Rename partition **P8** to **P_9**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 RENAME PARTITION P8 TO P_9;  
ALTER TABLE
```

- Rename partition **P_9** to **P8**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 RENAME PARTITION FOR (40000) TO P8;  
ALTER TABLE
```

- Modify the tablespace of a partition.

- Change the tablespace of partition **P6** to **example3**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 MOVE PARTITION P6 TABLESPACE  
example3;  
ALTER TABLE
```

- Change the tablespace of partition **P4** to **example4**:

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 MOVE PARTITION P4 TABLESPACE  
example4;  
ALTER TABLE
```

- Query a partition.

Query partition **P6**.

```
openGauss=# SELECT * FROM tpcds.web_returns_p2 PARTITION (P6);  
openGauss=# SELECT * FROM tpcds.web_returns_p2 PARTITION FOR (35888);
```

- Delete a partitioned table and its tablespaces.

```
openGauss=# DROP TABLE tpcds.web_returns_p2;  
DROP TABLE  
openGauss=# DROP TABLESPACE example1;  
openGauss=# DROP TABLESPACE example2;  
openGauss=# DROP TABLESPACE example3;  
openGauss=# DROP TABLESPACE example4;  
DROP TABLESPACE
```

3.2.6.3 Creating and Managing Indexes

Background

Indexes accelerate data access but increase the processing time of insertion, update, and deletion operations. Therefore, before creating an index, consider whether it is necessary and select the columns where indexes are to be created. You can determine whether to create an index for a table by analyzing the service processing and data use of applications, as well as columns that are frequently used as search criteria or need to be collated.

Indexes are created based on columns in database tables. Therefore, you must correctly identify which columns require indexes. You are advised to create indexes for any of the following columns:

- Columns that are often searched and queried. This speeds up searches.
- Columns that function as primary keys. This enforces the uniqueness of the columns and the data collation structures in organized tables.
- Columns that are often searched by range. The index helps collate data, and therefore the specified ranges are contiguous.
- Columns that often need to be collated. The index helps collate data, reducing the time for a collation query.
- Columns where the **WHERE** clause is executed frequently. This speeds up condition judgment.
- Columns that often appear after the keywords **ORDER BY**, **GROUP BY**, and **DISTINCT**.

NOTE

- After an index is created, the system automatically determines when to reference it. If the system determines that indexing is faster than sequential scan, the index will be used.
- After an index is created, it must be synchronized with the associated table to ensure that new data can be accurately located, which increases the data operation load. Therefore, delete unnecessary indexes periodically.
- Partitioned table indexes are classified into local indexes and global indexes. A local index corresponds to a specific partition, and a global index corresponds to the entire partitioned table.
- When logical replication is enabled, if you need to create a primary key index that contains system columns, you must set the **REPLICA IDENTITY** attribute of the table to **FULL** or use USING INDEX to specify a unique, non-local, non-deferrable index that does not contain system columns and contains only columns marked **NOT NULL**.

Procedure

For details about how to create a partitioned table, see [Creating and Managing Partitioned Tables](#).

- Create an index.
 - Create the local index **tpcds_web_returns_p2_index1** without specifying the partition name.

```
openGauss=# CREATE INDEX tpcds_web_returns_p2_index1 ON tpcds.web_returns_p2
(ca_address_id) LOCAL;
```

If the following information is displayed, the creation is successful:

```
CREATE INDEX
```

- Create the local index **tpcds_web_returns_p2_index2** with the partition name specified.

```
openGauss=# CREATE INDEX tpcds_web_returns_p2_index2 ON tpcds.web_returns_p2
(ca_address_sk) LOCAL
```

```
(
```

```
  PARTITION web_returns_p2_P1_index,
  PARTITION web_returns_p2_P2_index TABLESPACE example3,
  PARTITION web_returns_p2_P3_index TABLESPACE example4,
  PARTITION web_returns_p2_P4_index,
```

```
PARTITION web_returns_p2_P5_index,  
PARTITION web_returns_p2_P6_index,  
PARTITION web_returns_p2_P7_index,  
PARTITION web_returns_p2_P8_index  
) TABLESPACE example2;
```

If the following information is displayed, the creation is successful:

```
CREATE INDEX
```

- Create the global index **tpcds_web_returns_p2_global_index** for a partitioned table.

```
CREATE INDEX tpcds_web_returns_p2_global_index ON tpcds.web_returns_p2  
(ca_street_number) GLOBAL;
```

- Modify the tablespace of an index partition.

- Change the tablespace of index partition **web_returns_p2_P2_index** to **example1**.

```
openGauss=# ALTER INDEX tpcds.tpcds_web_returns_p2_index2 MOVE PARTITION  
web_returns_p2_P2_index TABLESPACE example1;
```

If the following information is displayed, the modification is successful:

```
ALTER INDEX
```

- Change the tablespace of index partition **web_returns_p2_P3_index** to **example2**.

```
openGauss=# ALTER INDEX tpcds.tpcds_web_returns_p2_index2 MOVE PARTITION  
web_returns_p2_P3_index TABLESPACE example2;
```

If the following information is displayed, the modification is successful:

```
ALTER INDEX
```

- Rename an index partition.

Rename the name of index partition **web_returns_p2_P8_index** to **web_returns_p2_P8_index_new**.

```
openGauss=# ALTER INDEX tpcds.tpcds_web_returns_p2_index2 RENAME PARTITION  
web_returns_p2_P8_index TO web_returns_p2_P8_index_new;
```

If the following information is displayed, the rename operation is successful:

```
ALTER INDEX
```

- Query indexes.

- Run the following command to query all indexes defined by the system and users:

```
openGauss=# SELECT RELNAME FROM PG_CLASS WHERE RELKIND='i' or RELKIND='I';
```

- Run the following command to query information about a specified index:

```
openGauss=# \di+ tpcds.tpcds_web_returns_p2_index2
```

- Delete indexes.

```
openGauss=# DROP INDEX tpcds.tpcds_web_returns_p2_index1;  
openGauss=# DROP INDEX tpcds.tpcds_web_returns_p2_index2;
```

If the following information is displayed, the deletion is successful:

```
DROP INDEX
```

GaussDB supports four methods for creating indexes. For details, see [Table 3-3](#).

NOTE

- After an index is created, the system automatically determines when to reference it. If the system determines that indexing is faster than sequential scan, the index will be used.
- After an index is successfully created, it must be synchronized with the associated table to ensure new data can be accurately located, which increases the data operation load. Therefore, delete unnecessary indexes periodically.

Table 3-3 Indexing method

Indexing Method	Description
Unique index	An index that requires the uniqueness of an index attribute or an attribute group. If a table declares unique constraints or primary keys, GaussDB automatically creates unique indexes (or composite indexes) for columns that form the primary keys or unique constraints. Currently, unique indexes can be created only for the B-tree and UB-tree in GaussDB.
Composite index	An index that can be defined for multiple attributes of a table. Currently, the B-tree in GaussDB supports composite indexes.
Partial index	An index that can be created for subsets of a table. This indexing method contains only tuples that meet condition expressions.
Expression index	An index that is built on a function or expression calculated based on one or more attributes of a table. An expression index works only when the queried expression is the same as the created expression.

- Create an ordinary table.

```
openGauss=# CREATE TABLE tpcds.customer_address_bak AS TABLE tpcds.customer_address;  
INSERT 0 0
```

- Create an ordinary index.

For the **tpcds.customer_address_bak** table, you need to perform the following operations frequently:

```
openGauss=# SELECT ca_address_sk FROM tpcds.customer_address_bak WHERE  
ca_address_sk=14888;
```

Generally, the database system needs to scan the **tpcds.customer_address_bak** table row by row to find all matched tuples. If the size of the **tpcds.customer_address_bak** table is large but only a few (possibly zero or one) of the **WHERE** conditions are met, the performance of this sequential scan is poor. If the database system uses an index to maintain the **ca_address_sk** attribute, the database system only needs to search a few tree layers for the matched tuples. This greatly improves data query performance. Furthermore, indexes can improve the update and deletion operation performance in the database.

Run the following command to create an index:

```
openGauss=# CREATE INDEX index_wr_returned_date_sk ON tpcds.customer_address_bak  
(ca_address_sk);  
CREATE INDEX
```

- Create a unique index.

Create a unique index on the **SM_SHIP_MODE_SK** column in the **tpcds.ship_mode_t1** table.

```
openGauss=# CREATE UNIQUE INDEX ds_ship_mode_t1_index1 ON  
tpcds.ship_mode_t1(SM_SHIP_MODE_SK);
```

- Create a composite index.

Assume you need to frequently query records with **ca_address_sk** being **5050** and **ca_street_number** smaller than **1000** in the

tpcds.customer_address_bak table. Run the following commands:

```
openGauss=# SELECT ca_address_sk,ca_address_id FROM tpcds.customer_address_bak WHERE  
ca_address_sk = 5050 AND ca_street_number < 1000;
```

Run the following command to define a composite index on **ca_address_sk** and **ca_street_number** columns:

```
openGauss=# CREATE INDEX more_column_index ON  
tpcds.customer_address_bak(ca_address_sk ,ca_street_number);  
CREATE INDEX
```

- Create a partial index.

If you only want to find records whose **ca_address_sk** is **5050**, you can create a partial index to facilitate your query.

```
openGauss=# CREATE INDEX part_index ON tpcds.customer_address_bak(ca_address_sk) WHERE  
ca_address_sk = 5050;  
CREATE INDEX
```

- Create an expression index.

Assume that you need to frequently query records with **ca_street_number** smaller than **1000**, run the following command:

```
openGauss=# SELECT * FROM tpcds.customer_address_bak WHERE trunc(ca_street_number) < 1000;
```

The following expression index can be created for this query task:

```
openGauss=# CREATE INDEX para_index ON tpcds.customer_address_bak (trunc(ca_street_number));  
CREATE INDEX
```

- Delete the **tpcds.customer_address_bak** table.

```
openGauss=# DROP TABLE tpcds.customer_address_bak;  
DROP TABLE
```

3.2.6.4 Creating and Managing Views

Context

If some columns in one or more tables in a database are frequently searched for, an administrator can define a view for these columns, and then users can directly access these columns in the view without entering search criteria.

A view is different from a base table. It is only a virtual object rather than a physical one. Only view definition is stored in the database and view data is not. The data is stored in a base table. If data in the base table changes, the data in the view changes accordingly. In this sense, a view is like a window through which users can know their interested data and data changes in the database. A view is triggered every time it is referenced.

Managing Views

- Create a view.

Create the **MyView** view. In the command, **tpcds.web_returns** indicates the created user table that contains the **wr_refunded_cash** integer column.

```
openGauss=# CREATE OR REPLACE VIEW MyView AS SELECT * FROM tpcds.web_returns WHERE  
trunc(wr_refunded_cash) > 10000;  
CREATE VIEW
```

NOTE

The **OR REPLACE** parameter in this command is optional. It indicates that if the view exists, the new view will replace the existing view.

- Query a view.

Run the following command to query **MyView**:

```
openGauss=# SELECT * FROM MyView;
```

- View details about a specified view.

Run the following command to view details about **MyView**:

```
openGauss=# \d+ MyView
```

```
View "PG_CATALOG.MyView"
Column |      Type      | Modifiers | Storage | Description
-----+-----+-----+-----+-----
USERNAME | CHARACTER VARYING(64) |          | extended |
View definition:
SELECT PG_AUTHID.ROLNAME::CHARACTER VARYING(64) AS USERNAME
FROM PG_AUTHID;
```

- Delete a view.

Run the following command to delete **MyView**:

```
openGauss=# DROP VIEW MyView;
DROP VIEW
```

3.2.6.5 Creating and Managing Sequences

Background

A sequence is a database object that generates unique integers. Sequence numbers are generated according to a certain rule. Sequences are unique because they increase automatically. This is why they are often used as primary keys.

You can create a sequence for a column in either of the following methods:

- Set the data type of a column to **sequence integer**. A sequence will be automatically created by the database for this column.
- Run the **CREATE SEQUENCE** statement to create a sequence. Set the initial value of the **nextval('sequence_name')** function to the default value of a column.

Procedure

Method 1: Set the data type of a column to a sequence integer. For example:

```
openGauss=# CREATE TABLE T1
(
    id serial,
    name text
);
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

Method 2: Create a sequence and set the initial value of the **nextval('sequence_name')** function to the default value of a column.

1. Create a sequence.

```
openGauss=# CREATE SEQUENCE seq1 cache 100;
```

If the following information is displayed, the creation is successful:

```
CREATE SEQUENCE
```

2. Set the default value of a column so that the column has a unique identification attribute.

```
openGauss=# CREATE TABLE T2
(
```

```
id int not null default nextval('seq1'),  
name text  
);
```

If the following information is displayed, the default value has been specified:

```
CREATE TABLE
```

3. Associate the sequence with a column.

Associates a sequence with a specified column included in a table. In this way, the sequence will be deleted when you delete its associated column or the table where the column belongs to.

```
openGauss=# ALTER SEQUENCE seq1 OWNED BY T2.id;
```

If the following information is displayed, the configuration is successful:

```
ALTER SEQUENCE
```

NOTE

The preceding methods are similar, except that the second method specifies cache for the sequence. A sequence having cache defined has inconsecutive values (such as 1, 4, and 5) and cannot maintain the order of its values. After the dependent column of a sequence has been specified, once the sequence is deleted, the sequence of the dependent will be deleted. A sequence shared by multiple columns is not forbidden in a database, but you are not advised to do that.

In the current version, you can specify the auto-increment column or set the default value of a column to **nextval('seqname')** when defining a table. You cannot add an auto-increment column or a column whose default value is **nextval('seqname')** to an existing table.

3.2.6.6 Creating and Managing Scheduled Jobs

Context

Time-consuming jobs, such as summarizing statistics or synchronizing data from another database, affect service performance if they are performed during the daytime and incur overtime hours if performed at night. To solve this problem, GaussDB Kernel is compatible with the scheduled job function in the A database. You can create scheduled jobs that are automatically triggered to reduce O&M workload.

This function calls APIs provided by the DBE_TASK package to create scheduled jobs, execute jobs automatically, delete jobs, modify job attributes (including job ID, the enabled/disabled status of a job, job triggering time, triggering interval, and job content), and implement other functions.

Managing Scheduled Jobs

Step 1 Create a test table.

```
openGauss=# CREATE TABLE test(id int, time date);
```

If the following information is displayed, the creation is successful:

```
CREATE TABLE
```

Step 2 Create a user-defined stored procedure.

```
openGauss=# CREATE OR REPLACE PROCEDURE PRC_JOB_1()  
AS  
N_NUM integer :=1;  
BEGIN
```

```
FOR I IN 1..1000 LOOP
INSERT INTO test VALUES(I,SYSDATE);
END LOOP;
END;
/
```

If the following information is displayed, the creation is successful:

```
CREATE PROCEDURE
```

Step 3 Create a job.

- Create a job with unspecified **job_id** and execute the **PRC_JOB_1** stored procedure every minute.

```
openGauss=# call db_task.submit('call public.prc_job_1(); ', sysdate, 'interval "1 minute"', :a);
id
-----
1
(1 row)
```

- Specify **job_id** to create a job. The value of **job_id** ranges from 1 to 32767.

```
openGauss=# call db_task.id_submit(1,'call public.prc_job_1(); ', sysdate, 'interval "1 minute"');
id_submit
-----
(1 row)
```

Step 4 View details of jobs created by the current user.

```
openGauss=# select job,dbname,start_date,last_date,this_date,next_date,broken,status,interval,failures,what
from my_jobs;
 job | dbname | start_date | last_date | this_date | next_date | broken | status | interval | failures | what
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | testdb | 2017-07-18 11:38:03 | 2017-07-18 13:53:03.607838 | 2017-07-18 13:53:03.607838 | 2017-07-18 13:54:03 | n | s | interval '1 minute' | 0 | call public.prc_job_1();
(1 row)
```

Step 5 Stop the job.

```
openGauss=# call db_task.finish(1,true);
finish
-----
(1 row)
```

Step 6 Start a job.

```
openGauss=# call db_task.finish(1,false);
finish
-----
(1 row)
```

Step 7 Modify job attributes.

- Modify the **next_time** parameter information about the job.

-- Set **next_time** of **Job1** to 1 hour so that **Job1** will be executed in one hour.

```
openGauss=# call db_task.next_time(1, sysdate+1.0/24);
next_time
-----
(1 row)
```

- Modify the **Interval** parameter of the job.

-- Set **Interval** of **Job1** to 1 hour so that **Job1** will be executed every one hour.

```
openGauss=# call db_task.interval(1,'sysdate + 1.0/24');
interval
-----
(1 row)
```

- Modify the **What** parameter of the job.
-- Set **What** to **insert into public.test values(333, sysdate+5);** for **Job1**.
openGauss=# call db_task.content(1,'insert into public.test values(333, sysdate+5);');
content

(1 row)
- Modify **Next_date**, **Interval**, and **What** parameters of the job.
openGauss=# call db_task.update(1, 'call public.prc_job_1();', sysdate, 'interval "1 minute"');
update

(1 row)

Step 8 Remove the job.

```
openGauss=# call db_task.cancel(1);
cancel
-----
(1 row)
```

Step 9 View the job execution status.

If a job fails to be automatically executed (that is, the value of **job_status** is 'f'), contact the administrator to view the gs_log run log to view the failure information of the job.

From **detail error msg**, you can see the failure causes.

```
LOG: Execute Job Detail:
     job_id: 1
    what: call public.test();
 start_date: 2017-07-19 23:30:47.401818
 job_status: failed
detail error msg: relation "test" does not exist
 end_date: 2017-07-19 23:30:47.401818
next_run_date: 2017-07-19 23:30:56.855827
```

Step 10 Set job permissions.

- During the creation of a job, the job is bound to the user and database that created the job. Accordingly, the user and database are added to **dbname** and **log_user** columns in the pg_job system catalog, respectively.
- If the current user is a database administrator, system administrator, or the user (**log_user** in pg_job) who created the job, the user has permissions to delete or modify job parameters using the **remove**, **change**, **next_data**, **what**, or **interval** parameter. Otherwise, the system displays a message indicating that the user has no permissions to perform operations on this job.
- If the current database is the one that created a job (that is, **dbname** in pg_job), you can delete or modify parameter settings of the job using the **cancel**, **update**, **next_data**, **content**, or **interval** parameter.
- When deleting the database that created a job (that is, **dbname** in pg_job), the system automatically deletes the job records of the database.
- When deleting the user who created a job (that is, **log_user** in pg_job), the system automatically deletes the job records of the user.

Step 11 Manage job concurrency.

You can configure the GUC parameter **job_queue_processes** to adjust the number of jobs running at the same time.

- Setting **job_queue_processes** to **0** indicates that the scheduled job function is disabled and no job will be executed.
- Setting **job_queue_processes** to a value that is greater than 0 indicates that the scheduled job function is enabled and this value is the maximum number of jobs that can be concurrently executed.

Too many concurrent jobs consume many system resources, so you need to set the number of concurrent jobs to be processed. If the current number of concurrent jobs reaches the value of **job_queue_processes** and some of them expire, these jobs will be postponed to the next polling period. Therefore, you are advised to set the polling interval (the **Interval** parameter of the submit API) based on the execution duration of each job to avoid the problem that jobs in the next polling period cannot be properly processed because of overlong job execution time.

Note: For database instances that do not use jobs, set **job_queue_processes** to **0** to disable the job function to reduce the resource consumption.

----End

4 Development and Design Proposal

4.1 Overview

Introduction

This document describes the database-related design and development specifications in the product design and development process based on the product lifecycle.

This document aims to improve the readability and code quality, and emphasizes the practicability and operability. It specifies the issues that may occur during the database development. The specifications are as follows:

Design specifications in terms of database and performance

Programming specifications in terms of typesetting, naming, comments, syntax, scripts, database programming, installation and deployment, and security

In addition, detailed rules and specific examples are provided for some specifications if necessary.

The development specifications are for reference only. The specific specifications should be comprehensively evaluated and implemented by database users based on the application technical architecture and planning.

Terms

This document uses the following terms for description:

- **Specification:** database specification, which must be complied with during programming and design. Otherwise, the database reports an error.
- **Rule:** a convention that must be followed during programming and design.
- **Recommendation:** a convention that must be considered during programming and design.
- **Description:** explanation of the rule or recommendation in question.
- **Example:** a positive or negative example of the rule or recommendation.

Applicability

This document is intended for designers, developers, development database administrators (DBAs), operation and maintenance (O&M) DBAs, and O&M personnel.

4.2 Database Design Specifications

4.2.1 General Specifications

- [Rule] Database valuable features

Table 4-1 Recommended database valuable features

Category	Feature List	Description
Table type	Partitioned table	Data partitioning.
Storage engine	Row store	Stores tables by row, which is recommended in the point query and a scenario with many add, delete, and alter operations.
Transaction	Transaction block	Explicitly starts a transaction.
	Single-statement transaction	Does not explicitly start a transaction. A single statement is a transaction.
Deployment	One primary and multiple standby DN	One primary DN and multiple standby DNs, ensuring data backup and disaster recovery.
Security	Transparent data encryption	Supports database-level storage encryption. Upper-layer services are unaware of the encryption.
Data type	Integer	TINYINT, SMALLINT, INTEGER, and BIGINT
	Arbitrary-precision	NUMERIC/DECIMAL
	Floating-point	REAL/FLOAT4, DOUBLE PRECISION/FLOAT8, and FLOAT
	Boolean	Boolean
	Fixed-length character	CHAR(n)
	Variable-length character	VARCHAR(n), NVARCHAR2(n), and TEXT
	Time	DATE, TIME, TIMETZ, TIMESTAMP, TIMESTAMPTZ, SMALLDATETIME, INTERVAL, and REALTIME

Category	Feature List	Description
	Binary	BYTEA (variable-length binary)
	Bit string	BIT(n) and VARBIT(n)
Function	String function	String data type processing function.
	Binary string function	Binary string type processing function.
	Numeric operation function	Numeric type processing function.
	Time and date processing function	Time and date type processing function.
Index	Primary key or unique index	Single-column or multi-column primary key or unique index.
	B-tree index	Index type.

- [Rule] Recommended best practices for database usage

Table 4-2 Recommended best practices for database usage

No.	Item	Recommended Value
1	Optimal number of long connections to the database	For details, see the default configuration of GUC parameter max_connections in the corresponding hardware specifications.
2	Maximum data volume on a single physical node	16 TB (Determine the capacity based on the backup and restoration specifications.)
3	Number of active databases	1
4	Total number of database tables and indexes	10000 (It is recommended that the number of tables in a single schema be less than or equal to 200.) NOTE The value is only a recommended value. You need to plan the value based on the disk capacity and customer services. Theoretically, the maximum value is 2^{32} .

No.	Item	Recommended Value
5	Optimal number of columns in a single table	< 50
6	Optimal number of indexes in a table	< 5
7	Optimal number of composite indexes in a table	< 3
8	Optimal number of columns in a single composite index	< 5
9	Optimal width of a single row	< 2000
10	Optimal capacity of a single column	< 10 MB
11	Optimal SQL statement length	< 5000
12	Percentage of available disk space	85% (If the disk capacity is greater than 24 TB, 90% is recommended.)

4.2.2 Deployment Specifications

[Rule] After a database is installed, DBA must configure the database GUC parameters properly based on the environment features.

4.2.3 Database Object Naming Specifications

- [Specification] Do not use reserved or non-reserved keywords to name database objects.

You can run the following command to query the database keyword, or view it in the [Keywords](#) section:

```
SELECT * FROM pg_get_keywords();
```

- [Specification] The database object name contains a maximum of 63 bytes.

Table 4-3 Restriction on the length of database object names

Object Type	Length Restriction (Unit: Byte)	Remarks
Database name	≤ 63	-
Table name	≤ 63	-
Table object	≤ 63	If the length of a table name exceeds this value, the kernel truncates the table name. As a result, the actual name is inconsistent with the configured value. In addition, characters may be truncated in different character sets and unexpected characters may appear.
Temporary table not used by services	≤ 63	For example, an intermediate table used by operation personnel for temporary backup or data collection. Naming rule: tmp_Table name abbreviation_Creator account abbreviation_Creation date , for example, tmp_user_ytw_160505 .
Table column	≤ 63	Meaningful English words or abbreviations are used to name the column. For example, a column of the bool type is named in the format of is_Description . A column whose member status is enabled in the member table is named is_enabled .
View name and column name	≤ 63	-
Function name and parameter name	≤ 63	-

- [Rule] Do not use a string enclosed in double quotation marks (") to define the database object name, unless you must specify its capitalization.
 - By default, GaussDB is case-insensitive for object names in SQL statements. If two different tables (**t_Table** and **t_table**) exist in the same database, use double quotation marks ("). Otherwise, avoid using double quotation marks (").
 - Case sensitivity of database object names makes fault locating difficult.
- [Rule] Identifiers must meet the following requirements:
 - An identifier starts with a letter or underscore (_) and can contain letters, digits, underscores (_), dollar signs (\$), and number signs (#).

- If an identifier is enclosed in backquotes (``) in B-compatible mode or double quotation marks (""), any combination of valid characters can be used, for example, "123gs_column".
 - Identifiers are case-insensitive. They are case-sensitive only when they are enclosed in backquotes (``) in B-compatible mode or double quotation marks ("").
- [Recommendation] Do not use a string enclosed in double quotation marks (") to define the database object name, unless you need to specify its capitalization. Case sensitivity of database object names makes fault locating difficult.
- [Recommendation] Use the same naming style for database objects.
 - In a system undergoing incremental development or service migration, you are advised to comply with its historical naming conventions.
 - You are advised to use multiple words separated with underscores (_).
 - You are advised to use intelligible names and common acronyms or abbreviations for database objects. For example, you can use English words or Chinese pinyin indicating actual business terms. The naming format should be consistent within a database instance.
 - A variable name must be descriptive and meaningful. It must have a prefix indicating its type.
- [Recommendation] The name of a table object should indicate its main characteristics, for example, whether it is an ordinary, temporary, or unlogged table.
 - An ordinary table name should indicate the business relevant to a data set.
 - Temporary tables are named in the format of **tmp_Suffix**.
 - Unlogged tables are named in the format of **ul_Suffix**.
 - Do not create database objects whose names start with **redis_**.
 - Do not create database objects whose names start with **mlog_** or **matviewmap_**.

4.2.4 Database and Schema Design Specifications

- [Description] When you create a database, exercise caution when you set **ENCODING** and **DBCOMPATIBILITY** configuration items. GaussDB supports the A, B, C, and PG compatibility modes, which are compatible with the syntax of Oracle Database, MySQL, Teradata, and Postgres, respectively. The syntax behavior varies according to the compatibility mode. By default, the A compatibility mode is used.
- [Description] By default, a database owner has all permissions for all objects in the database, including the deletion permission. Exercise caution when deleting a permission.
- [Description] To let a user access an object in a schema, assign the usage permission and the permissions for the object to the user, unless the user has the SYSADMIN permission or is the schema owner.
- [Description] To let a user create an object in the schema, grant the CREATE permission for the schema to the user.

- [Description] By default, a schema owner has all permissions for all objects in the schema, including the deletion permission. Exercise caution when deleting a permission.
- [Specification] The database name must be specified when the JDBC client is used to connect to the database. The format is as follows:
`jdbc:postgresql://ip:port/database_name`

NOTICE

Once a JDBC instance is created, the database cannot be switched.

-
- [Specification] The database does not support case-insensitive collation.
 - [Rule] Before using services, a system administrator must create databases, schemas, and users for services, and then grant object permissions to corresponding users.
 - [Rule] Create a service database before using a service. Do not use the Postgres database created by default after the database is installed to store service data. You are advised to create your own database based on service requirements.
 - [Rule] When creating a database, you must set the character set to **UTF8** and must select the character set that is the same as that of the client.

To meet globalization requirements, database encoding should store and identify most characters. Therefore, UTF8 is recommended. The UTF-8 character set is equivalent to the utf8mb4 character set in MySQL databases and supports emoji.

If the encoding mode of the client is different from that of the database, transcoding is required, affecting performance. In addition, kernel optimization for the same encoding cannot be triggered, affecting the query efficiency.

To change the character set of the client, perform the following steps:

- Set client connection parameters. Take JDBC as an example. Add the parameters **characterEncoding** and **allowEncodingChanges** to the URL.
`jdbc:postgresql://ip:port/database_name?characterEncoding=utf8&allowEncodingChanges=true`
- Modify the database GUC parameter.
`SET client_encoding = 'UTF8';`

Set the database encoding when executing **CREATE DATABASE**.

`CREATE DATABASE tester WITH ENCODING = 'UTF8';`

NOTICE

The character set cannot be changed once the database is created.

-
- [Recommendation] In a database instance, it is recommended that the number of user-defined databases be 3 and no more than 10. If there are too many user-defined databases, O&M operations, such as upgrade and backup, will be inefficient.
 - [Recommendation] It is recommended that the number of schemas in the actual user environment be no more than 100. If there are too many schemas

in a database, operations that depend on the number of schemas, such as `gs_dump`, becomes slow.

- [Recommendation] You are advised to use schemas to isolate services for convenience and resource sharing.

Both database and schema can be used to isolate services . Differences are as follows:

- Databases share few resources. Connections to and permissions on them are also isolated. However, the databases cannot access each other. The database must be specified during JDBC connection establishment. After the connection is established, the database cannot be switched.
 - Schemas share more resources than databases do. User permissions on schemas and subordinate objects can be controlled over the GRANT and REVOKE syntax.
- [Recommendation] You are advised to specify **LC_CTYPE** and **LC_COLLATE** when creating a database. These parameters will affect the data collation sequence.

Example:

```
CREATE DATABASE SAMPLE_DB WITH LC_CTYPE = 'zh_CN.gbk' LC_COLLATE = 'zh_CN.gbk';
```

4.2.5 Permission Design Specifications

- [Rule] The initial database user can only connect and access a database for DBA management. Services are not allowed to directly log in as this user to connect to and access the database.

NOTE

Initial database user, that is, the database installation user, which has the same name as the OS user to which the database belongs.

- [Rule] The initial database user creates a database and users for services. Services use the created users to log in to and access the database.

- [Rule] Assign permissions to roles and users based on the least privilege principle.

- [Recommendation] Manage permissions by roles instead of users.

Use roles to manage permissions, that is, configure permissions for roles and grant the roles to users.

Roles facilitate permission management in scenarios such as multiple users and user changes. Example:

- Roles and users are in many-to-many relationship. A role can be granted to multiple users. If permissions of a role are modified, the permissions of the granted users can be updated at the same time.
 - Deleting a user does not affect the role.
 - A new user can quickly obtain required permissions by granting a role to the user.
- [Recommendation] When deleting a specified database, revoke the CONNECT permission of users on the database to prevent deletion failures caused by active database connections.

4.2.6 Character Set Design Specifications

- [Specification] Currently, character sets can be defined only for databases.
- [Recommendation] The database and client use the same UTF-8 character set.
 - Clients include ODBC, JDBC, and GSQL.
 - For reasons, see [Database and Schema Design Specifications](#).

4.2.7 Table Design Specifications

- [Rule] View nesting is not recommended.
 - On the one hand, if a wildcard is used when a view is written, an error occurs in the view when a column is added to or deleted from the called view.
 - On the other hand, view nesting may be inefficient because indexes cannot be used. You are advised to use base tables with indexes instead of performing join operations on views.
- [Recommendation] Try to avoid collation operations when defining a view. ORDER BY is invalid in the top-level view. If you must collate the output data, use ORDER BY in a called view.
- [Recommendation] Minimize random I/Os. Through clustering, you can sequentially store hot data, converting random I/O to sequential I/O to reduce the cost of I/O scanning.

Selecting a Storage Model

[Recommendation] Selecting a storage model is the first step in defining a table. The storage model mainly depends on the customer's service type. For details, see [Table 4-4](#).

Table 4-4 Table storage models and scenarios

Storage Model	Application Scenario
Row store	• Point queries (simple index-based queries that only return a few records). • Scenarios requiring frequent addition, deletion, and modification

Selecting a Partitioning Mode

Comply with the following rules to partition a table containing a large amount of data:

- [Rule] Select a proper partition scheme for large data tables.
Currently, RANGE PARTITION, LIST PARTITION, and HASH PARTITION partitioning methods are provided. Note the following principles during partitioning:

- You are advised to use interval columns, such as date and area, for partitioning.
- The upper boundary value of a range partition must be set to **MAXVALUE** to prevent data overflow.
- The partition name must reflect the data features of the partition. For example, the partition name can be in the format of *keyword+interval feature*.
- [Recommendation] It is recommended that the number of partitions be less than or equal to 100, the data volume of a single partition be less than or equal to 50 million, and the data capacity of a single partition be less than or equal to 50 GB.
- [Description] Reduce the amount of data to be scanned. You can use the pruning mechanism of a partitioned table.
- [Recommendation] Create partitions on columns that indicate certain ranges, such as dates and regions.
- [Recommendation] A partition name should show the data characteristics of a partition. For example, its format can be *Keyword+Range* characteristics.
- [Recommendation] Set the upper limit of a partition to **MAXVALUE** to prevent data overflow.

Table 4-5 Table partitioning modes and scenarios

Partitioning Mode	Description
Range	Table data is partitioned by range.
Interval	Table data is partitioned by range. If the data exceeds the range, a new partition is automatically created based on the interval.
List	Table data is partitioned by a specified column based on a specific value.
Hash	Table data is partitioned by hash.

The example of defining a partitioned table is as follows:

```
-- Create a range partitioned table.
CREATE TABLE staffS_p1
(
  staff_ID    NUMBER(6) not null,
  FIRST_NAME  VARCHAR2(20),
  LAST_NAME   VARCHAR2(25),
  EMAIL       VARCHAR2(25),
  PHONE_NUMBER VARCHAR2(20),
  HIRE_DATE   DATE,
  employment_ID VARCHAR2(10),
  SALARY      NUMBER(8,2),
  COMMISSION_PCT NUMBER(4,2),
  MANAGER_ID  NUMBER(6),
  section_ID  NUMBER(4)
)
PARTITION BY RANGE (HIRE_DATE)
(
  PARTITION HIRE_19950501 VALUES LESS THAN ('1995-05-01 00:00:00'),
  PARTITION HIRE_19950502 VALUES LESS THAN ('1995-05-02 00:00:00'),
```

```
PARTITION HIRE_maxvalue VALUES LESS THAN (MAXVALUE)
);

-- Create an interval partitioned table. The table has two initial partitions. When data that is not in
the partition range is inserted, another partition is automatically added.
CREATE TABLE sales
(prod_id NUMBER(6),
 cust_id NUMBER,
 time_id DATE,
 channel_id CHAR(1),
 promo_id NUMBER(6),
 quantity_sold NUMBER(3),
 amount_sold NUMBER(10,2)
)
PARTITION BY RANGE (time_id)
INTERVAL('1 day')
( PARTITION p1 VALUES LESS THAN ('2019-02-01 00:00:00'),
  PARTITION p2 VALUES LESS THAN ('2019-02-02 00:00:00')
);

-- Create a list partitioned table.
CREATE TABLE test_list (col1 int, col2 int)
partition by list(col1)
(
 partition p1 values (2000),
 partition p2 values (3000),
 partition p3 values (4000),
 partition p4 values (5000)
);

-- Create a hash partitioned table.
CREATE TABLE test_hash (col1 int, col2 int)
partition by hash(col1)
(
 partition p1,
 partition p2
);
```

For details about the table partition syntax, see [CREATE TABLE PARTITION](#).

4.2.8 Column Design Specifications

- [Rule] Use recommended data types for column design.
Recommended data types must be used for column design. If you want to use prohibited or unrecommended data types, contact database experts for evaluation. Some data types are not recommended because they apply to limited service scenarios and are not used on a large scale for commercial purposes.

Table 4-6 Best practices for database data types

Data Type	Description	Recommended or Not
UUID	Different databases may generate the same UUID.	Prohibited
Serial integer	Auto-increment column, including SMALLSERIAL, SERIAL, and BIGSERIAL.	Prohibited
Integer	TINYINT, SMALLINT, INTEGER, and BIGINT	Recommended

Data Type	Description	Recommended or Not
Arbitrary-precision	NUMERIC/DECIMAL	Recommended
Floating-point	REAL/FLOAT4, DOUBLE PRECISION/FLOAT8, and FLOAT	Recommended
Boolean	Boolean	Recommended
Fixed-length character	CHAR(<i>n</i>)	Recommended
Variable-length character	VARCHAR(<i>n</i>) and NVARCHAR2(<i>n</i>)	Recommended
	TEXT and CLOB (character large object)	Not recommended
Time	DATE, TIME, TIMESTAMP, SMALLDATETIME, INTERVAL, and REALTIME	Recommended
	TIMETZ and TIMESTAMPTZ	Not recommended
Binary	BYTEA (variable-length binary)	Recommended
	BLOB (binary large object), and RAW (variable-length hexadecimal string)	Not recommended
Bit string	BIT(<i>n</i>) and VARBIT(<i>n</i>)	Recommended
Special character	NAME and "CHAR" are usually used within the database system.	Not recommended
JSON	JSON data does not support operators.	Not recommended
User-defined	Defines enumerated types.	Not recommended
HLL	You are advised to use the HLL functions to reduce the impact on performance.	Not recommended
Currency	The MONEY type stores a currency amount with fixed fractional precision.	Not recommended
Geometric	POINT, LSEG, BOX, PATH, POLYGON, and CIRCLE	Not recommended

Data Type	Description	Recommended or Not
Network address	Stores IPv4 and MAC addresses.	Not recommended

- [Rule] Use the most specific numeric data types. If all of the following numeric types provide the required service precision, they are recommended in descending order of priority: integer, floating-point number, and numeric.
- [Rule] Properly set the data type of a numeric column based on the value range, and use the NUMERIC or DECIMAL type as less as possible. NUMERIC and DECIMAL are equivalent. NUMERIC or DECIMAL data operations consume great CPU resources.

Table 4-7 Storage space and value range of numeric data types

Type	Storage Size (Unit: Byte)	Minimum Value	Maximum Value
TINYINT	1	0	255
SMALLINT	2	-32768	32767
INTEGER	4	-2147483648	2147483647
BIGINT	8	-9223372036854775808	9223372036854775807
REAL/FLOAT4	4	6-bit decimal digits	
DOUBLE PRECISION/FLOAT8	8	15-bit decimal digits	

- [Rule] Select a proper string type. If the value of a column must be a fixed-length character, use fixed-length character types or automatically add spaces. Otherwise, use the variable-length character type VARCHAR.
For a typical fixed-length column, for example, **gender**, you can enter only **f** or **m** that occupies a byte. You are advised to use the fixed-length data type (for example, CHAR(*n*)) for this type of columns.
If such requirement does not exist or longer characters may be required for future expansion, use variable-length character types (such as VARCHAR and TEXT) preferentially. You are advised not to specify the length of variable-length characters.
The reasons are as follows:
 - For fixed-length columns, the input data that is shorter than the fixed length will be padded with space characters and then be saved to the database. This wastes the storage space in the database.

- For fixed-length character types, the entire table needs to be scanned and rewritten if the length needs to be extended later. This causes high performance overhead and affects online services.
- For a variable-length column with a fixed length, the system checks whether the length exceeds the limit each time upon data insertion. This causes performance overhead.
- [Rule] Do not store data of the numeric type in columns of the character type.
If numeric calculation or comparison (for example, adding a filter condition) is performed on data stored in columns of the character type, unnecessary overhead will be caused due to data type conversion, and the column indexes may become invalid, affecting query performance.
- [Rule] Do not store data of the time or date type in columns of the character type.
If calculation or comparison (for example, adding a filter condition) with data of the time or date type is performed on data stored in columns of the character type, unnecessary overhead will be caused by data type conversion, and the column indexes may become invalid, affecting query performance.
- [Rule] Add NOT NULL constraints to columns that never have NULL values.
In certain scenarios, the optimizer may specially optimize non-**NULL** columns to improve query performance.
- [Rule] Use the same data type for joined columns.
If the column types are inconsistent during a join operation, overhead will be caused by data type conversion.
- [Recommendation] In tables that are logically related, columns having the same meaning should use the same data type.
- [Recommendation] For string data, you are advised to use variable-length strings and specify the maximum length. To avoid truncation, ensure that the specified maximum length is greater than the maximum number of characters to be stored. You are advised not to use CHAR(n), BPCHAR(n), NCHAR(n), or CHARACTER(n), unless you know that the string length is fixed.
- [Rule] The number of long columns, such as VARCHAR(1000) and VARCHAR(4000), does not exceed 8.
- [Recommendation] When defining a column, you are advised to create a comment for the column to facilitate subsequent maintenance.
For details about the description, value range, and usage of different types of fields, see [Data Type](#).
- [Recommendation] Add NOT NULL constraints to columns that are used for WHERE filtering and join operations.
In certain scenarios, the optimizer may specially optimize non-**NULL** columns to greatly improve query performance.
- [Recommendation] Do not reserve columns for a table. In most cases, you can quickly add or delete table columns, or change the default values of columns.
An added column must meet the following requirements. Otherwise, the entire table is updated, leading to additional overheads and affecting online services.

- The data type is BOOL, BYTEA, SMALLINT, BIGINT, SMALLINT, INTEGER, NUMERIC, FLOAT, DOUBLE PRECISION, CHAR, VARCHAR, TEXT, TIMESTAMPTZ, TIMESTAMP, DATE, TIME, TIMETZ, or INTERVAL.
- The length of the default value cannot exceed 128 bytes.
- The default value of the added column does not contain the volatile function.
- The default value is required and cannot be **NULL**.

If you are not sure whether the third condition is met, contact GaussDB technical support for evaluation.

4.2.9 Index Design Specifications

- [Specification] Use the index types in the recommended database index practices.

NOTE

Recommended types must be used for index design. If you need to use prohibited, unrecommended, or restricted index types, contact database experts for advice.

Table 4-8 Recommended database index practices

Index Type	Description	Recommended or Not
Primary key or unique index	Single-column or multi-column primary key or unique index.	Recommended
Global index	Index organization method.	Recommended for some data types
Expression index	An index column is a function or scalar expression calculated from one or multiple columns in a table.	Restricted
B-tree index	Index building type.	Recommended for some data types

- [Rule] Properly design composite indexes to avoid redundancy.
For example, if an index has been created for **(a, b, c)**, you shall not create an index for **(a)**, **(b)**, **(c)**, **(a, b)**, or **(b, c)** independently. If only the filtering condition of the **a** column is used for a query, composite indexes are used for the query as well.
- [Rule] A composite index contains a maximum of 5 columns.
- [Rule] A composite index contains a maximum of 200 characters.
- [Rule] Index (including single-column and composite indexes) columns must be **NOT NULL**.
- [Rule] Do not create multiple unique indexes for a single table.

Maintaining multiple unique indexes at the same time generates more overheads than maintaining a multi-column unique index. If multiple unique indexes are equivalent to a multi-column unique index in service logic, a multi-column unique index shall be preferred.

- [Recommendation] The efficiency of maintaining indexes created for the same column is different. Columns of the number type are better than those of the character type and other data types. Therefore, it is recommended that columns such as IDs and time for creating indexes be stored as data of the number type.
- [Recommendation] Create indexes on joined columns.

Hash join is supported, whereas nested loop join may be used for join operations if the rescan cost is low (for example, the inner table is small). If the **NESTLOOP JOIN** plan can be viewed by executing **EXPLAIN**, you can create indexes on joined columns to improve the efficiency of executing **NESTLOOP JOIN**.

4.2.10 Function/Stored Procedure Design Specifications

- [Rule] Do not use stored procedures or triggers to implement service logic. Instead, process the logic on the service server to prevent logical dependency on the database.
- [Rule] Do not use stored procedures to implement the upgrade logic in the database upgrade scripts of a service.
- [Rule] Create functions that have fixed return values for fixed input parameters and set the function type to **IMMUTABLE** or **SHIPPABLE**.

Currently, the database supports three types of functions: **IMMUTABLE**, **STABLE**, and **VOLATILE**.

If the attribute of an **IMMUTABLE** function is set to **SHIPPABLE**, the function can be executed on DNs. In most scenarios, such functions perform efficiently. However, such functions require fixed return values for fixed input parameters to ensure that the functions are correctly executed on DNs. If the function execution result depends on the table scan result (for example, maximum value of a column in the table) or scan time (for example, current time), set the function type to **STABLE** or **VOLATILE**, and set the function attribute to **NOT SHIPPABLE** to ensure that the function is correctly executed.

4.2.11 Constraint Design

DEFAULT and NULL Constraints

- [Recommendation] If all the column values can be obtained from services, you are advised not to use the **DEFAULT** constraint. Otherwise, unexpected results will be generated during data loading.
- [Recommendation] Add **NOT NULL** constraints to columns that never have **NULL** values. The optimizer automatically optimizes the columns in certain scenarios.
- [Recommendation] Explicitly name all constraints excluding **NOT NULL** and **DEFAULT**.

Unique Constraints

- [Description] Row-store tables support UNIQUE constraints.
- [Recommendation] The constraint name should indicate that it is a unique constraint, for example, **UNI***Included columns*.

Primary Key Constraints

- [Description] Row-store tables support PRIMARY KEY constraints.
- [Recommendation] The constraint name should indicate that it is a primary key constraint, for example, **PK***Column name*.

Check Constraints

- [Description] Row-store tables support CHECK constraints.
- [Recommendation] The constraint name should indicate that it is a check constraint, for example, **CK***Column name*.

4.2.12 View and Joined Table Design

View Design

- [Proposal] Do not nest views unless they have strong dependency on each other.
- [Proposal] Try to avoid collation operations in a view definition.

Joined Table Design

- [Proposal] Minimize joined columns across tables.
- [Proposal] Use the same data type for joined columns.
- [Proposal] The names of joined columns should indicate their relationship. For example, they can use the same name.

4.3 Database Design Specifications

4.3.1 GUC Parameter Programming Specifications

- [Rule] Clients (such as JDBC) should use default parameters for query. If session-level GUC parameters need to be modified, exercise caution when performing this operation.

When modifying GUC parameters through ODBC or JDBC, note that the GUC parameters take effect only in the current connection. Especially in the connection pool scenario, problems may occur and cause difficulty in locating problems.

If the GUC parameters must be set in a connection, you must use the following statements before releasing the connection to the connection pool:

```
SET SESSION AUTHORIZATION DEFAULT;  
RESET ALL;
```

Clear the connection status.

4.3.2 Object Access Programming Specifications

- [Rule] When operating an object, a user should have the operation permission on the object.

For details about the permission description, see [Users and Permissions](#) and comply with [Permission Design Specifications](#).

- [Rule] Use *schemaname.tablename* to access objects (such as tables and functions).

If the schema name prefix is not added, the system searches all tablespaces in the current **search_path** until a matched table is found, which causes unnecessary performance overhead.

4.3.3 WHERE

- [Rule] Do not compare the same table columns in WHERE conditions.

For example, the following statement does not meet specifications:

```
SELECT * FROM t1 WHERE col1 = col1;
```

The following modification should be considered:

```
SELECT * FROM t1 WHERE col1 IS NOT NULL;
```

- [Rule] Do not include implicit data type conversion in WHERE conditions.

After implicit conversion is performed in the database, the created index may fail to be used.

You are advised to enable the GUC parameter **check_implicit_conversions** and disable **enable_fast_query_shipping** during development to check whether query statements contain implicit data types that may affect performance.

```
SET enable_fast_query_shipping = off;  
SET check_implicit_conversions = true;
```

Implicit data type conversion check requires extra overhead. After the query statement is developed, disable the **check_implicit_conversions** parameter and reset **enable_fast_query_shipping**.

Example:

- The following lines of code do not meet specifications:

```
-- The phonenumber column in the t_tablename table is of the VARCHAR type instead of the numeric type.
```

```
SELECT column1  
  INTO i_l_variable1  
  FROM t_tablename  
 WHERE phonenumber = 13512345678;
```

- The following modification should be considered:

```
-- The phonenumber column in the t_tablename table is of the VARCHAR type instead of the numeric type.
```

```
SELECT column1  
  INTO i_l_variable1  
  FROM t_tablename  
 WHERE phonenumber = '13512345678';
```

- [Rule] Do not use expressions or functions in WHERE condition columns.

When expressions or functions are used in condition columns, indexes become invalid, and each row of data is calculated, causing unnecessary performance consumption.

Example:

- The following lines of code do not meet specifications:

```
SELECT income FROM table WHERE abs(income) > ?;
SELECT income FROM table WHERE income * 10 > ?;
SELECT create_time
FROM table
WHERE date_format(create_time, '%Y%m%d %H:%i:%s') = '20090101 00:00:0';
```

- The following modification should be considered:

```
SELECT income FROM table WHERE income > ? OR income < (-1) * ?;
SELECT income FROM table WHERE income > ?/10;
SELECT create_time
FROM table
WHERE create_time = str_to_date('20090101 00:00:0', '%Y%m%d %H:%i:%s');
```

- [Rule] Do not use comparison operator (!=) when comparing with NULL in query conditions. Instead, use IS NULL or IS NOT NULL.
- [Rule] Do not use comparison operator (!=) in the index columns of the query condition to avoid invalid indexes.
- [Rule] Do not compare the index columns of the query condition with NULL (IS NULL and IS NOT NULL instead).
- [Rule] Do not use NOT in the index columns of the query condition.
- [Rule] Do not use NOT IN in the index columns of the query condition.
- [Rule] Do not place % at the beginning of the LIKE statement for fuzzy query. If % is placed at the beginning of the LIKE statement, the index cannot be used and the entire table will be scanned.
- [Recommendation] The number of IN clauses in the WHERE statement shall be less than or equal to 500.
 - During the query, the values of all the IN clauses will be compared to check whether they are equal, which increases the overhead.
 - If the included values are relatively fixed, consider creating a temporary table, writing the clause data into the table, and then using INNER JOIN to implement the inclusion query.
- [Recommendation] If the IN clause in the WHERE condition is not a constant but a column in the table, you are advised to change it to a subquery. In this case, it is actually an unequal JOIN, which is executed through the nestloop plan. If the table is too large, the execution efficiency is low. You are advised to change the query to a subquery of the JOIN type.

Example:

- The following code does not meet specifications:

```
SELECT col1, COALESCE(max(col2 - 1), 0)
FROM t1, t2
WHERE t1.col1 = ANY(VALUES(id1), (id2))
GROUP BY col1;
```

- The following modification should be considered:

```
SELECT col1, COALESCE(max(tmp), 0) FROM
(
  (
    SELECT col1, (col2-1) AS tmp
    FROM t1, t2
    WHERE t1.col1 = t2.id1 AND t1.col1 != t2.id2
  ) UNION ALL (
    SELECT col1, (col2-1) AS tmp
    FROM t1, t2
    WHERE t1.col1 = t2.id2
  )
) GROUP BY col1;
```


- [Recommendation] Use equal operators, instead of not equal operators.
If a not equal operator (IN, BETWEEN, <, <=, >, or >=) is used in the WHERE condition, no index can be used in the following conditions because two range conditions cannot be used at the same time.

4.3.4 SELECT

- [Rule] Do not use the wildcard character (*) in the SELECT statement.

If the table structure is changed due to service or database upgrade when a wildcard column is used to query a table, the table structure may be incompatible with service statements. Therefore, the service must specify the name of the table column to be queried and avoid using the wildcard character.

- [Rule] Do not perform operations that may cause sorting, such as ORDER BY, DISTINCT, GROUP BY, and UNION, on large columns (such as VARCHAR(2000)).

These operations consume a large number of CPU and memory resources, resulting in low execution efficiency.

- [Rule] Do not use the LOCK TABLE statement to lock a table. Instead, use the SELECT .. FOR UPDATE statement.

LOCK TABLE provides multiple levels of locks. However, if you do not fully understand the database principles and services, misuse of table locks may trigger deadlocks, causing the database to be unavailable.

- [Recommendation] Use UNION ALL instead of UNION. Consider data deduplication if necessary.

UNION ALL does not deduplicate data and does not require sorting operations. Therefore, UNION ALL is faster than UNION. If deduplication is not required, UNION ALL is preferred.

- [Recommendation] Do not frequently use count() to obtain the number of rows in a large table. This operation consumes a large number of resources and affects the execution efficiency of parallel jobs.

If you do not need the real-time row statistics, run the following statement to obtain the number of rows in the table:

```
SELECT reltuples FROM pg_class WHERE relname = 'tablename';
```

The number of rows recorded in pg_class is updated only after ANALYZE is executed on the table.

Currently, ANALYZE is triggered in either of the following conditions:

- The service sends an ANALYZE statement. For example:

```
-- Analyze all tables in the connection library.
ANALYZE;
-- Analyze the specified table.
ANALYZE tablename;
```
- This event is triggered when the number of rows added or deleted at a specified interval or in a table reaches a specified value by using the AUTO VACUUM mechanism. The interval and addition/deletion ratio can be set through the GUC parameters.

4.3.5 INSERT

- [Specification] INSERT ON DUPLICATE KEY UPDATE does not support UPDATE on columns with PRIMARY KEY or UNIQUE constraints.

The semantics of **INSERT ON DUPLICATE KEY UPDATE** is to update the rows that have **UNIQUE** constraint conflicts. In this process, the value of the constraint should not be updated.

- [Rule] Do not run **INSERT ON DUPLICATE KEY UPDATE** on a table that has multiple **UNIQUE** constraints.

A table may have multiple unique indexes, or both primary keys and unique indexes. When multiple **UNIQUE** constraints exist, the system checks all the **UNIQUE** constraints by default. If any constraint conflicts, the system updates the conflicting rows. That is, multiple records may be updated, which does not meet the service expectation. More specific conditions for inserting and updating data shall be added.

- [Recommendation] In the case of batch insertion, you are advised to use `executeBatch` to execute **INSERT INTO VALUES (?)**. The execution efficiency is higher than that of executing multiple **INSERT INTO VALUES()** or **INSERT INTO VALUES(?), ..., (?)** statements.

4.3.6 UPDATE

- [Specification] **LIMIT** cannot be directly used in the **UPDATE** statement. The **WHERE** condition must be used to specify the target row to be updated.

- [Specification] Multi-table update is not supported.

Multi-table update is to update multiple tables in a single SQL statement.

- [Rule] The **UPDATE** statement must contain the **WHERE** clause to avoid full table scan.

- [Rule] Do not use the updated columns as the update source when the **UPDATE** clause updates multiple columns simultaneously.

Even if multiple columns are updated simultaneously from the same source, the behavior varies depending on the database. To avoid compatibility issues, avoid the preceding operations at the service layer.

Example:

```
UPDATE table SET col1 = col2, col3 = col1 WHERE col1 = 1;
```

In , the value of **col3** is the original value of **col1**. In database B, the value of **col3** is the value of **col2** (because the value of **col2** is assigned to **col1**).

- [Rule] Do not use **ORDER BY** or **GROUP BY** in the **UPDATE** statement to avoid unnecessary sorting.
- [Recommendation] If a table has a primary key or index, the **WHERE** condition must be used together with the primary key or index during update.

4.3.7 DELETE

- [Specification] **LIMIT** cannot be used in the **DELETE** statement. The **WHERE** condition should be used to specify the target row to be deleted.

- [Specification] Multi-table deletion is not supported.

Multi-table deletion indicates that multiple tables are deleted in a single SQL statement.

- [Rule] The **DELETE** statement must contain the **WHERE** clause to avoid full table scan.

- [Rule] Do not use **ORDER BY** or **GROUP BY** in the **DELETE** statement to avoid unnecessary sorting.

- [Rule] Use TRUNCATE instead of DELETE to clear a table.
TRUNCATE creates a new physical file and physically deletes the original file when the transaction ends to clear the disk space. However, the DELETE statement marks data in the table and does not clear the disk space until the VACUUM FULL phase.
- [Recommendation] If **DELETE** is executed on a table that has a primary key or index, the WHERE condition must be used together with the primary key or index to improve execution efficiency.

4.3.8 Join Query

- [Rule] The nesting depth of multi-table join must be less than 8.
If the join nesting is too deep, slow SQL statements may be generated. You need to optimize the join nesting at the service layer.
- [Rule] Specify the join condition (ON) of each table to avoid Cartesian product.
For example, in B database, JOIN is equivalent to CROSS JOIN and INNER JOIN. However, in the SQL standards, JOIN is equivalent to INNER JOIN only and must be used together with the ON condition.
- [Rule] Specify the JOIN mode based on the SQL standards during join. Do not use the JOIN keyword directly. Instead, use CROSS JOIN, INNER JOIN, LEFT JOIN, or RIGHT JOIN.
- [Rule] When multiple tables are joined for query, aliases must be added to the tables to ensure that the statement logic is clear and easy to maintain.
- [Recommendation] Different columns have different comparison overheads. Use a column type with high efficiency for joined columns.
 - The comparison efficiency of the numeric type is much higher than that of the string type.
 - The comparison efficiency of integers is much higher than that of numeric and floating-point types.
- [Recommendation] The joined columns must be of the same data type to avoid the impact of implicit type conversion on the execution efficiency.
- [Recommendation] Do not use nested subqueries. Use table join if possible because subqueries will generate temporary tables, which greatly affects SQL performance.
- [Recommendation] If a large number of NULL values exist in the joined columns, you are advised to add the IS NOT NULL condition to the WHERE condition to improve the execution efficiency.

4.3.9 Subquery

- [Rule] Do not use repeated subquery statements in an SQL statement.
- [Recommendation] Do not use scalar subqueries.

A scalar subquery is a subquery whose result is one value and whose condition expression uses an equal operator.

Example:

The following statement does not meet specifications:

```
SELECT * FROM t1 WHERE id = (SELECT id FROM t2 LIMIT 1);
```

You are advised to split the preceding statement into two SQL statements and execute the subquery first.

- [Recommendation] Do not use subqueries in the SELECT target columns. Otherwise, the plan cannot be pushed down, affecting the execution performance.
- [Recommendation] The nesting depth of a subquery cannot exceed 2. Subqueries cause temporary table overhead. Therefore, complex queries must be optimized based on service logic.

4.3.10 Transaction

- [Specification] Large object operations do not support transactions. Large object operations include creating and deleting DATABASE, ANALYZE, and VACUUM jobs.
- [Rule] When accessing the database through the JDBC client, disable the **autocommit** parameter and explicitly execute the transaction COMMIT.
 - On the one hand, enabling the **autocommit** parameter will cause some parameters (such as **fetchsize**) to become invalid.
 - On the other hand, the service should clarify the service logic and reduce the dependence on the database.
- [Rule] Do not combine multiple SQL statements into one statement when accessing the database through JDBC.

When multiple statements are combined into one statement that contains object operations, if the intermediate object operation fails, a new transaction is started to execute subsequent statements.

Example:

- The following statement does not meet specifications:

```
Connection conn = ....
try {
    Statement stmt = null;
    try {
        stmt = conn.createStatement();
        stmt.executeUpdate("CREATE TABLE t1 (a int); DROP TABLE t1");
    } finally {
        stmt.close();
    }
    conn.commit();
} catch (Exception e) {
    conn.rollback();
} finally {
    conn.close();
}
```

In the preceding statement, if

```
CREATE TABLE t1;
```

The creation fails, and a new transaction will be started.

```
DROP TABLE t1;
```

The execution fails.

- You are advised to split the statement into two statements and send them separately:

```
Connection conn = ....
try {
    Statement stmt = null;
    try {
```

```
stmt = conn.createStatement();
stmt.executeUpdate("CREATE TABLE t1 (a int)");
stmt.executeUpdate("DROP TABLE t1");
} finally {
    stmt.close();
}
conn.commit();
} catch(Exception e) {
    conn.rollback();
} finally {
    conn.close();
}
```

4.3.11 SQL Compilation

DDL

- [Recommendation] In GaussDB, you are advised to perform DDL operations (such as CREATE TABLE and COMMENT) in a unified manner. Avoid DDL operations in batch processing. to avoid performance deterioration caused by many concurrent transactions.
- [Recommendation] Perform the TRUNCATE operation immediately after an unlogged table is used, because GaussDB cannot ensure the security of unlogged tables in abnormal scenarios.
- [Recommendation] Suggestions on the storage model of temporary and unlogged tables are the same as those on base tables. If the base table is a row-store table, you are advised to create temporary tables and unlogged tables as row-store tables.
- [Recommendation] The total length of an index column cannot exceed 50 bytes. Otherwise, the index size will increase greatly, resulting in large storage cost and low index performance.
- [Recommendation] Do not delete objects using DROP...CASCADE, unless the dependency between objects is specified. Otherwise, the objects may be deleted by mistake.

Data Loading and Unloading

- [Recommendation] Explicitly provide the inserted column list in the INSERT statement. Example:

```
INSERT INTO task(name,id,comment) VALUES ('task1','100','100th task');
```
- [Recommendation] After data is imported to the database in batches or the data increment reaches the threshold, you are advised to analyze tables to prevent the execution plan from being degraded due to inaccurate statistics.
- [Recommendation] To clear all data in a table, you are advised to use TRUNCATE TABLE instead of DELETE TABLE. DELETE TABLE is not efficient and cannot release disk space occupied by the deleted data.

Type Conversion

- [Recommendation] Convert data types explicitly. If you perform implicit conversion, the result may differ from expected.
- [Recommendation] During data query, explicitly specify the data type for constants, and do not attempt to perform any implicit data type conversion.

- [Description] If **sql_compatibility** is set to **A**, null strings will be automatically converted to NULL during data import. If null strings need to be retained, set **sql_compatibility** to **C**.

Query Operation

- [Recommendation] Do not return a large number of result sets to a client except the ETL program. If a large result set is returned, consider modifying your service design.
- [Recommendation] Perform DDL and DML operations encapsulated in transactions. For example, operations such as TRUNCATE TABLE, UPDATE TABLE, DELETE TABLE, and DROP TABLE cannot be revoked once they are committed. You are advised to encapsulate such operations in transactions so that you can roll back the operations if necessary.
- [Recommendation] During query compilation, you are advised to list all columns to be queried and avoid using **SELECT ***. Doing so reduces output lines, improves query performance, and avoids the impact of adding or deleting columns on front-end service compatibility.
- [Recommendation] During table object access, add the schema prefix to the table object to avoid accessing an unexpected table due to schema switchover.
- [Recommendation] The cost of joining more than three tables or views, especially FULL JOIN, is difficult to be estimated. You are advised to use the WITH TABLE AS statement to create interim tables to improve the readability of SQL statements.
- [Recommendation] Avoid using a Cartesian product or FULL JOIN. Cartesian products and full joins will result in a sharp expansion of result sets and poor performance.
- [Description] Only IS NULL and IS NOT NULL can be used to determine NULL value comparison results. If any other method is used, NULL is returned. For example, NULL instead of expected Boolean values is returned for NULL<>NULL, NULL=NULL, and NULL<>1.
- [Description] Do not use count(col) instead of count(*) to count the total number of records in a table. count(*) counts the NULL value (actual rows) while count(col) does not.
- [Description] While executing count(col), the number of NULL record rows is counted as 0. While executing sum(col), NULL is returned if all records are NULL. If not all the records are NULL, the number of NULL record rows is counted as 0.
- [Description] To count multiple columns using count(), column names must be enclosed in parentheses. For example, count ((col1, col2, col3)). Note: When multiple columns are used to count the number of NULL record rows, a row is counted even if all the selected columns are NULL. The result is the same as that when count(*) is executed.
- [Description] Null records are not counted when count(distinct col) is used to calculate the number of non-NULL columns that are not repeated.
- [Description] If all statistical columns are NULL when count(distinct (col1,col2,...)) is used to count the number of unique values in multiple columns, null records are also counted, and the records are considered the same.

- [Recommendation] Use the connection operator || to replace the **concat** function for string connection, because the concat function needs to query type tables and function tables, which slows down the basic performance. In addition, because the concat output is related to the data type, the execution plan generated by the concat function cannot calculate results in advance. As a result, the query performance severely deteriorates.
- [Recommendation] Use the following time-related macros to replace the now function and obtain the current time, because the execution plan generated by the now function cannot be pushed down to disks. As a result, the query performance deteriorates severely.

Table 4-9 Time-related macros

Macro Name	Description	Example
CURRENT_DATE	Obtains the current date, excluding the hour, minute, and second details.	openGauss=# SELECT CURRENT_DATE; date ----- 2018-02-02 (1 row)
CURRENT_TIME	Obtains the current time, excluding the year, month, and day.	openGauss=# SELECT CURRENT_TIME; timetz ----- 00:39:34.633938+08 (1 row)
CURRENT_TIMESTAMP(n)	Obtains the current date and time, including year, month, day, hour, minute, and second. NOTE n indicates the number of digits after the decimal point in the time string.	openGauss=# SELECT CURRENT_TIMESTAMP(6); timestampz ----- 2018-02-02 00:39:55.231689+08 (1 row)

- [Recommendation] Do not use scalar subquery statements. A scalar subquery appears in the output list of a SELECT statement. In the following example, the underlined part is a scalar subquery statement:

```
SELECT id, (SELECT COUNT(*) FROM films f WHERE f.did = s.id) FROM staffs_p1 s;
```

Scalar subqueries often result in query performance deterioration. During application development, scalar subqueries need to be converted into equivalent table joins based on the service logic.

- [Recommendation] In WHERE clauses, the filter conditions should be collated. The condition that few records are selected for reading (the number of filtered records is small) is listed at the beginning.
- [Recommendation] Filter conditions in WHERE clauses should comply with unilateral rules, that is, to place the column name on one side of a comparison operator. In this way, the optimizer automatically performs pruning optimization in some scenarios. The format is *col op expression*, where *col* indicates a table column, *op* indicates a comparison operator, such

as = and >, and *expression* indicates an expression that does not contain a column name. Example:

```
SELECT id, from_image_id, from_person_id, from_video_id FROM face_data WHERE  
current_timestamp(6) - time < '1 days'::interval;
```

The modification is as follows:

```
SELECT id, from_image_id, from_person_id, from_video_id FROM face_data where time >  
current_timestamp(6) - '1 days'::interval;
```

- [Recommendation] Do not perform unnecessary collation operations. Collation requires a large amount of memory and CPU. If service logic permits, ORDER BY and LIMIT can be combined to reduce resource overhead. By default, GaussDB performs collation by ASC & NULL LAST.
- [Recommendation] When the ORDER BY clause is used for collation, specify collation modes (ASC or DESC), and use NULL FIRST or NULL LAST for NULL record sorting.
- [Recommendation] Do not rely on only the LIMIT clause to return the result set displayed in a specific sequence. Combine ORDER BY and LIMIT clauses for some specific result sets and use OFFSET to skip specific results if necessary.
- [Recommendation] If the service logic is accurate, you are advised to use UNION ALL instead of UNION.
- [Recommendation] If a filter condition contains only an OR expression, convert the OR expression to UNION ALL to improve performance. SQL statements that use **OR** expressions cannot be optimized, resulting in slow execution. Example:

```
SELECT * FROM scdc.pub_menu  
WHERE (cdp= 300 AND inline=301) OR (cdp= 301 AND inline=302) OR (cdp= 302 AND inline=301);
```

Convert the statement to the following:

```
SELECT * FROM scdc.pub_menu  
WHERE (cdp= 300 AND inline=301)  
union all  
SELECT * FROM scdc.pub_menu  
WHERE (cdp= 301 AND inline=302)  
union all  
SELECT * FROM tablename  
WHERE (cdp= 302 AND inline=301)
```

- [Recommendation] If an IN(val1, val2, val3...) expression contains a large number of columns, you are advised to replace it with the IN (VALUES (val1), (val2), (val3)...) statement. The optimizer will automatically convert the IN constraint into a non-correlated subquery to improve the query performance.
- [Recommendation] Replace (NOT) IN with (NOT) EXIST when joined columns do not contain **NULL** values. For example, in the following query statement, if the **T1.C1** column does not contain any **NULL** value, add the NOT NULL constraint to the **T1.C1** column, and then rewrite the statements.

```
SELECT * FROM T1 WHERE T1.C1 NOT IN (SELECT T2.C2 FROM T2);
```

Rewrite the statement as follows:

```
SELECT * FROM T1 WHERE NOT EXISTS (SELECT * FROM T2 WHERE T1.C1=T2.C2);
```

NOTE

- If the value of the T1.C1 column is not **NOT NULL**, the preceding rewriting cannot be performed.
- If the **T1.C1** column is the output of a subquery, check whether the output is **NOT NULL** based on the service logic.
- [Recommendation] Use cursors instead of the LIMIT OFFSET syntax to perform pagination queries to avoid resource overheads caused by multiple

executions. A cursor must be used in a transaction, and you must disable the cursor and commit the transaction once the query is finished.

4.4 Client Programming Specifications

4.4.1 JDBC

- [Specification] A database must be specified for each JDBC instance. Once the instance is created, it cannot switch to another database.
- [Specification] The length of a single SQL statement cannot exceed 2 GB bytes. Considering the communications cost, it is recommended that the length of a single SQL statement be less than or equal to 5 KB.
- [Specification] Currently, only the DEFAULT value of the columns in CREATE/ALTER TABLE can be set through parameters. Other DDL statements do not support parameter settings using the Prepare Execute method.
- [Rule] Each PreparedStatement of JDBC can contain a maximum of 32767 parameters.
- [Specification] The connection parameter **fetchsize** must be used when **autocommit** is disabled. Otherwise, the **fetchsize** configuration is invalid.
- [Rule] Use the default GUC parameters and avoid sending SET requests through JDBC to modify GUC parameters. For details, see [GUC Parameter Programming Specifications](#).
- [Recommendation] Use Prepare Execute method to execute query statements to improve execution efficiency.
- [Rule] Execute SQL statements one by one in the same transaction. Do not combine multiple SQL statements and send them as one statement.

For details, see [\[Rule\] Do not combine multiple SQL statements into one statement when accessing the database through JDBC](#).

- [Rule] The time zone of the host where the JDBC client is located, the time zone of the host where the database is located, and the time zone during database configuration must be the same.
- [Rule] If a temporary table is created in a connection, delete the temporary table before releasing the connection to the connection pool to avoid service errors.
- [Description] When a third-party tool connects to GaussDB through JDBC, JDBC sends a connection request to GaussDB. By default, the following parameters are added. For details, see the implementation of the ConnectionFactoryImpl JDBC code.

```
params = {  
  { "user", user },  
  { "database", database },  
  { "client_encoding", "UTF8" },  
  { "DateStyle", "ISO" },  
  { "extra_float_digits", "3" },  
  { "TimeZone", createPostgresTimeZone() },  
};
```

These parameters may cause the JDBC and gsql clients to display inconsistent data, for example, date data display mode, floating point precision representation, and timezone.

If the result is not as expected, you are advised to explicitly set these parameters in the Java connection setting.

When the database is connected through JDBC, **extra_float_digits** is set to **3**. When the database is connected using gsql, **extra_float_digits** is set to **0**. As a result, the precision of the same data displayed in JDBC clients may be different from that displayed in gsql clients.

In precision-sensitive scenarios, the numeric type is recommended.

- [Recommendation] When connecting to the database through JDBC, ensure that the following three time zones are the same:
 - Time zone of the host where the JDBC client is located
 - Time zone of the host where the GaussDB database instance is located.
 - Time zone used during GaussDB database instance configuration.

NOTE

For details about how to set the time zone, contact the administrator.

- [Recommendation] Enable **autocommit** in the code for connecting to GaussDB by the JDBC. If **autocommit** needs to be disabled to improve performance or for other purposes, applications need to ensure that transactions are committed. For example, explicitly commit transactions after specifying service SQL statements. Particularly, ensure that all transactions are committed before the client exits.
- [Recommendation] You are advised to use connection pools to limit the number of connections from applications. You are advised not to connect to a database each time an SQL statement is executed.
- [Recommendation] After an application completes its jobs, disconnect it from GaussDB to release occupied resources. You are advised to set the session timeout interval in the jobs.
- [Recommendation] Reset the session environment before releasing connections to the JDBC connection tool. Otherwise, historical session information may cause object conflicts.
 - If GUC parameters are set in the connection, run **SET SESSION AUTHORIZATION DEFAULT;RESET ALL;** to clear the connection status before you return the connection to the connection pool.
 - If a temporary table is used, delete the temporary table before you return the connection to the connection pool.
- [Recommendation] In the scenario where the ETL tool is not used and real-time data import is required, it is recommended that you use the CopyManager API driven by the GaussDB JDBC to import data in micro-batches during application development.
- [Recommendation] Set **prepareThreshold** to a proper value. If the query statement is fixed, set it to **1**.
- [Recommendation] You are advised to set **batchMode** to **on** to use the batch connection mode to improve the execution performance.
- [Recommendation] Set a proper JDBC connection timeout interval based on the upper-layer request timeout of the service to prevent the job from always occupying database resources.

Timeout parameters include **loginTimeout**, **connectTimeout**, and **socketTimeout**.

- **loginTimeout**: integer type. Specifies the waiting time for establishing the database connection, in seconds. The default value is **0**, indicating that the timeout mechanism is disabled.
- **connectTimeout**: integer type. This parameter specifies the timeout interval for connecting to a server. If the time taken to connect to a server exceeds the value specified, the connection is interrupted. The unit of the timeout interval is second. The default value **0** indicates that the timeout mechanism is disabled.
- **socketTimeout**: integer type. This parameter indicates the timeout interval for a socket read operation. If the time taken to read data from a server exceeds the value specified, the connection is closed. The unit of the timeout interval is second. The default value **0** indicates that the timeout mechanism is disabled.
- **cancelSignalTimeout**: integer type. A cancel message may cause a block. This attribute controls "connect timeout" and "socket timeout" in a cancel command. The default value is **10**.
- **tcpKeepAlive**: Boolean type. This parameter is used to enable or disable TCP keepalive detection. The default value is **false**.

5 Application Development Guide

5.1 GaussDB Application Development Guide

You can connect to and perform operations on a database through APIs such as JDBC, ODBC, libpq, Psycopg, Go, and ecpg.

JDBC

Java Database Connectivity (JDBC) is a Java API for running SQL statements. It provides unified access APIs for different relational databases. It allows Java applications to use SQL statements to perform database operations, such as querying, inserting, updating, and deleting data.

Some key features and usage of JDBC are as follows:

1. Database connection management: JDBC allows applications to establish connections to databases and manages the lifecycle of these connections.
2. SQL execution: JDBC can be used to query, update, and delete SQL statements. This is achieved by constructing SQL statements in Java code and sending them to the database.
3. Transaction management: JDBC APIs can be used to start, commit, or roll back transactions to ensure the consistency and integrity of database operations.
4. Exception handling: JDBC defines a group of exception classes to handle various exceptions that may occur during database operations. Developers can use the try-catch block to capture and handle these exceptions.
5. Batch processing: JDBC provides the batch processing function, which allows multiple SQL statements to be executed concurrently, improving the efficiency of database operations.
6. Metadata access: JDBC can be used to obtain the metadata information of the database, such as the table structure, column name, and data type. Developers then can dynamically construct SQL query statements or perform operations based on the database structure.

In conclusion, JDBC provides a flexible and powerful bridge that enables Java applications to interact with various databases and to implement data persistence

and management. GaussDB supports JDBC 4.0 and requires JDK 1.8 for code compiling. It does not support JDBC-ODBC bridging.

For details about JDBC-based development, see [Development Based on JDBC](#).

ODBC


Open Database Connectivity (ODBC) is a Microsoft API for accessing databases in C/C++ based on the X/OPEN CLI. It provides a unified method for applications to access various database management systems (DBMSs) without considering specific database types or OS platforms. ODBC allows applications to use SQL to query, insert, update, and delete data in a database. Applications interact with the database through the APIs provided by ODBC, which enhances their portability, scalability, and maintainability.

The ODBC architecture consists of three main components: application, ODBC driver manager, and ODBC driver. Applications use ODBC APIs to communicate with the ODBC driver manager which loads and manages ODBC drivers. The ODBC drivers are responsible for communicating with a specific database, executing SQL queries, and returning results. [Figure 5-1](#) shows the system structure of ODBC.

In conclusion, ODBC provides a flexible and cross-platform method that allows users to easily connect applications to various databases without worrying about the details of a specific database system.

For details about ODBC-based development, see [Development Based on ODBC](#).

Figure 5-1 ODBC system structure



GaussDB supports ODBC in the following environments.

Table 5-1 Platforms supported by ODBC

OS	Platform
CentOS 6.4/6.5/6.6/6.7/6.8/6.9/7.0/7.1/7.2/7.3/7.4	x86_64
CentOS 7.6	Arm64
EulerOS 2.0 SP2/SP3	x86_64
EulerOS 2.0 SP8	Arm64
Kylin V10	x86_64
Kylin V10	Arm64

The ODBC Driver Manager running on Unix or Linux can be unixODBC or iODBC. unixODBC-2.3.7 is used as the component for connecting to the database.

Windows has a native ODBC Driver Manager. You can locate **Data Sources (ODBC)** by choosing **Control Panel > Administrative Tools**.

NOTE

The current database ODBC driver is based on an open-source version and may be incompatible with data types tinyint, smalldatetime, nvarchar, and nvarchar2.

ODBC limitations:

- ODBC does not support read on standby.
- ODBC does not support user-defined types and does not support user-defined parameters in stored procedures.
- ODBC does not support DR failover.
- When the **proc_outparam_override** parameter is enabled for the database, ODBC cannot properly call the stored procedure that contains the **out** parameter.

libpq

libpq is a GaussDB C API. libpq contains a set of library functions that allow client programs to send query requests to the GaussDB servers and obtain query results. It is also the underlying engine of other GaussDB APIs, such as ODBC.

For details about libpq-based development, see [Development Based on libpq](#).

5.2 Development Specifications

If the connection pool mechanism is used during application development, comply with the following specifications:

- If GUC parameters are set in the connection, run **SET SESSION AUTHORIZATION DEFAULT;RESET ALL;** to clear the connection status before you return the connection to the connection pool.
- If a temporary table is used, delete the temporary table before you return the connection to the connection pool.

If you do not do so, the connection in the connection pool will be stateful, which affects subsequent operations on the connection pool.

[Table 5-2](#) describes the compatibility of application development drivers.

Table 5-2 Compatibility

Driver	Compatibility
JDBC, Go, ODBC, libpq, Psycopg, and ecpg	The drivers are compatible with the earlier database version. However, to use the new features added to the drivers and database, you must upgrade the database.

NOTICE

Setting **behavior_compat_options** to '**proc_outparam_override**' is applicable only in A-compatible mode.

If the driver is used in a multi-thread environment:


The JDBC driver is not thread-safe and does not guarantee that the connection methods are synchronized. The caller synchronizes the calls to the driver.

5.3 Development Based on JDBC

5.3.1 Development Process

Application development based on JDBC covers obtaining the driver JAR package and configuring JDK 1.8, connecting to a database, running SQL statements, processing result sets, and closing the connection.

Figure 5-2 Application development process based on JDBC



5.3.2 Development Process

5.3.2.1 Obtaining the Driver JAR Package and Configuring JDK 1.8

Environment preparation includes obtaining the driver JAR package and configuring JDK 1.8.

Obtaining the Driver JAR Package

Decompress the obtained driver package, find the JDBC driver package **GaussDB-Kernel_Database version number_OS_64bit_Jdbc.tar.gz**, and decompress it. After the decompression, you will obtain the following driver packages in JAR format:

- **gsjdbc4.jar**: The main class name is **org.postgresql.Driver**, and the URL prefix of the database connection is **jdbc:postgresql**. This driver package applies to the scenario where services are migrated from PG. The driver class and loading path are the same as those before the migration, but the supported APIs are different. The APIs that are not supported need to be adjusted on the service side.
- **opengaussjdbc.jar**: The main class name is **com.huawei.opengauss.jdbc.Driver**. The URL prefix of the database connection is **jdbc:opengauss**. This driver package is recommended. This driver package is used when both PG and GaussDB are accessed in a JVM process.

NOTICE

- The loading paths and URL prefixes of driver classes vary in different driver packages, but their functions are the same.
 - **gsjdbc200.jar**: This driver package applies to the scenario where services are migrated from Gauss200. The driver class and loading path are the same as those before the migration, but the supported APIs are different. The APIs that are not supported need to be adjusted on the service side.
 - The **gsjdbc4.jar** driver package cannot be used to operate the PG database. Although the connection can be successfully established in some versions, some API behaviors are different from those of PG JDBC, which may cause unknown errors.
 - The PG driver package cannot be used to operate the GaussDB database. Although the connection can be successfully established in some versions, some API behaviors are different from those of GaussDB JDBC, which may cause unknown errors.
-

Configuring JDK 1.8

JDK 1.8 must be configured on the client. JDK supports multiple platforms such as Windows and Linux. The following uses Windows as an example to describe the configuration method:

- Step 1** Enter the following command in the MS-DOS window (command prompt in Windows) to check the JDK version:

```
java -version
```

Ensure that the JDK version is JDK 1.8. If JDK is not installed, download the installation package from the [official website](#) and install it.

Step 2 Configure system environment variables.

1. Right-click **My computer** and choose **Properties**.
2. In the **System** window, click **Advanced system settings** in the navigation pane.
3. In the **System Properties** dialog box, click **Environment Variables** on the **Advanced** tab page.
4. In the **System variables** area of the **Environment Variables** dialog box, click **New** or **Edit** to configure system variables. For details, see [Table 5-3](#).

Table 5-3 Variables

Variable	Operation	Variable Value
JAVA_HOME	– If the variable exists, click Edit. – If the variable does not exist, click New.	Specifies the Java installation directory. Example: C:\Program Files\Java\jdk1.8.0_131
Path	Click Edit .	– If <i>JAVA_HOME</i> is configured, add %JAVA_HOME%\bin before the variable value. – If <i>JAVA_HOME</i> is not configured, add the full Java installation path before the variable value: C:\Program Files\Java\jdk1.8.0_131\bin
CLASSPATH	Click New .	.;%JAVA_HOME%\lib;%JAVA_HOME%\lib\tools.jar

----End

5.3.2.2 Connecting to a Database

After you have obtained the driver JAR package and configured JDK 1.8, GaussDB allows you to connect to a database in three methods. This section describes how to connect to a database in the three methods.

5.3.2.2.1 Connection Methods

The database connection method is usually selected based on factors such as security, performance, scalability, and implementation complexity. The following lists advantages, disadvantages, and selection principles of three database connection methods.

Table 5-4 Advantages, disadvantages, and selection principles of different methods

Connection Method	Advantage	Disadvantage	Selection Principle
Non-encrypted	- This method is easy to implement and does not require additional encryption and decryption operations. - It features high performance because no encryption or decryption operation is required, which reduces the overhead of data transmission.	- It provides low security and data is easily intercepted or tampered with during transmission. - It is not suitable for sensitive data transmission and may violate privacy protection laws and regulations.	It applies only to internal networks or environments that have low security requirements.
SSL	- It provides encrypted transmission. Data is protected during transmission, ensuring high security. - It supports client/server identity authentication, which enhances security.	- It is complex in configuration and management, and certificates need to be issued and updated. - Some performance overhead may be added, especially when it is used for encrypting and decrypting large amounts of data.	It applies to environments that require high data transmission security, such as finance and healthcare industries.

Conn ectio n Met hod	Advantage	Disadvantage	Selection Principle
UDS	- It applies to local communication without passing through the network, featuring a high transmission speed. - No extra network overhead is required, and it is not vulnerable to network attacks.	It is used only for the communication between processes on the same host but cannot be used for remote connections.	It applies to local communication scenarios, for example, data exchange between local services and applications.

5.3.2.2.2 Connection Parameter Reference

All connection properties of the **info** parameter are case-sensitive. [Table 5-5](#) describes the common properties.

Table 5-5 Connection properties of the info parameter

Property	Description	Value
PGDBNAME	Specifies the database name. You do not need to set this parameter in the URL because the database name is automatically parsed from the .properties file.	Property type: string
PGHOST	Specifies the host IP address.	Property type: string
PGPORT	Specifies the host port number.	Property type: integer
user	Specifies the database user who creates the connection.	Property type: string
password	Specifies the password of the database user.	Property type: string

Property	Description	Value
driverInfo Mode	Controls the output mode of the driver description information.	Property type: string Value range: postgresql and gaussdb. • postgresql: The driver description related to PG is displayed. • gaussdb: The driver description related to GaussDB is displayed. Default value: postgresql
loggerLevel	Specifies the logging level.	Property type: string Value range: The following logging levels are supported: OFF, INFO, DEBUG, and TRACE. • OFF: The logging function is disabled. • INFO, DEBUG, and TRACE logs record information of different levels. Default value: No default value is provided. If this property is not set, the value INFO is used.
loggerFile	Specifies the log output path (directory and file name). This parameter has been deprecated and does not take effect. To specify the log output path, you can configure the java.util.logging property file or system property.	Property type: string
logger	Specifies the log output framework used by the JDBC driver. The JDBC driver supports the log output framework used for interconnecting with user applications.	Property type: string Value range: Slf4JLogger • If it is left empty, JDK LOGGER is used. • Currently, only the Slf4j-API-based third-party log output framework is supported. For details, see Log Management.

Property	Description	Value
allowEncodingChanges	If this parameter is set to true , the character set type can be changed. This parameter is used together with characterEncoding to set the character set. The two parameters are separated by ampersands (&). The value of characterEncoding can be UTF8 , GBK , LATIN1 , or GB18030 . Example: allowEncodingChanges=true&characterEncoding=UTF8 .	Property type: Boolean Value range: true: The character set type can be changed. false: The character set type cannot be changed. Default value: false
currentSchema	Specifies the schema of the current connection. You need to specify a schema in search-path . If the schema name contains special characters except letters, digits, and underscores (_), you are advised to enclose the schema name in quotation marks. Note that the schema name is case-sensitive after quotation marks are added. If multiple schemas need to be configured, separate them with commas (,). Schemas containing special characters also need to be enclosed in quotation marks. Example: currentSchema=schema_a,"schema-b","schema/c"	Property type: string Default value: If this parameter is not set, the default schema is the username used for the connection.
hostRecheckSeconds	After JDBC attempts to connect to a host, the host status is saved: connection success or connection failure. This status is trusted within the period specified by hostRecheckSeconds . After the period expires, the status becomes invalid.	Property type: integer Unit: s Value range: 0 to 2147483647. Default value: 10

Property	Description	Value
ssl	Specifies that the database is connected in SSL mode. When ssl is set to true, the NonValidatingFactory channel and certificate mode are supported. 1. For the NonValidatingFactory channel, configure the username and password and set SSL to true. 2. In certification mode, configure the client certificate, key, and root certificate, and set SSL to true.	Property type: Boolean Value range: • true: The database is connected in SSL mode. • false: The database is not connected in SSL mode. Default value: No default value is provided. If this property is not set, the value false is used.

Property	Description	Value
sslmode	This parameter specifies the SSL authentication mode.	Property type: string Value range: disable, allow, prefer, require, verify-ca, and verify-full. • disable: SSL connection is disabled. • allow: If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified. • prefer: If the database supports SSL connection, SSL connection is preferred. However, authenticity of the database server will not be verified. • require: The system attempts to set up an SSL connection. If there is a CA file, the system performs verification as if the parameter was set to verify-ca. • verify-ca: The system attempts to set up an SSL connection and checks whether the server certificate is issued by a trusted CA. • verify-full: The system attempts to set up an SSL connection, checks whether the server certificate is issued by a trusted CA, and checks whether the host name of the server is the same as that in the certificate. Default value: No default value is provided. If this property is not set, the value require is used.

Property	Description	Value
sslcert	Specifies the complete path of the certificate file. The certificate type is End Entity .	Property type: string Default value: No default value is provided. If this property is not set, the root directory of the user is read.
sslkey	Specifies the complete path of the key file. You need to convert the key to the DER format before using it. <code>openssl pkcs8 -topk8 -outform DER -in client.key -out client.key.pk8 -nocrypt</code>	Property type: string Default value: No default value is provided. If this property is not set, the root directory of the user is read.
sslrootcert	Specifies the name of the SSL root certificate. The root certificate type is CA .	Property type: string Default value: No default value is provided. If this property is not set, the root directory of the user is read.
sslpassword	Specifies the SSL password, which is provided for ConsoleCallbackHandler.	Property type: string
sslpasswordcallback	Specifies the class name of the SSL password provider.	Property type: string Default value: No default value is provided.
sslfactory	Specifies the class name used by SSLSocketFactory to establish an SSL connection.	Property type: string Default value: No default value is provided.
sslprivatekeyfactory	This parameter specifies the fully qualified name of the implementation class of the org.postgresql.ssl.PrivateKeyFactory API that implements the private key decryption method. If this parameter is not specified, try the default JDK private key decryption algorithm. If the decryption fails, use org.postgresql.ssl.BouncyCastlePrivateKeyFactory. You need to provide the bcpkix-jdk15on.jar package. The recommended version is 1.65 or later.	Property type: string Default value: No default value is provided.
sslfactoryarg	Specifies the optional parameter of the constructor function of the sslfactory class. (This parameter is not recommended.)	Property type: string Default value: No default value is provided.

Property	Description	Value
sslhostnamerverifier	Specifies the class name of the host name verifier. The API must implement javax.net.ssl.HostnameVerifier. The default value is com.huawei.opengauss.jdbc.ssl.PGjdbcHostnameVerifier .	Property type: string Default value: No default value is provided. If this property is not set, com.huawei.opengauss.jdbc.ssl.PGjdbcHostnameVerifier is used.

Property	Description	Value
loginTimeout	Specifies the waiting time for establishing the database connection. When multiple IP addresses are configured in the URL, if the time for obtaining the connection exceeds the value of this parameter, the connection fails and the subsequent IP addresses are not tried. The value of loginTimeout is related to connectTimeout and socketTimeoutInConnecting. The calculation formula is as follows: loginTimeout = (connectTimeout + Connection authentication time + Initialization statement execution time) x Number of nodes.	Property type: integer Unit: s Value range: 0 to 2147483647. The value 0 indicates that the timeout mechanism does not take effect. Default value: 0 Suggestion: After this parameter is set, an asynchronous thread is started each time a connection is established. If there are a large number of connections, the pressure on the client may increase. If this parameter needs to be set in a centralized environment, you are advised to set it as follows: $\max(\text{connectTimeout}, \text{socketTimeoutInConnecting}) \times \text{Number of nodes}$. This prevents connection failures when the network is abnormal and the nth IP address is the IP address of the primary node.

Property	Description	Value
	NOTICE • This parameter sets the time for attempting to connect to all IP addresses in a list. If this parameter is set to a small value, the subsequent IP addresses in the list may fail to be connected. For example, if three IP addresses are set, loginTimeout is set to 5s, and it takes 5s to connect to the first two IP addresses, the third IP address cannot be connected. In a centralized environment, the last IP address is the IP address of the primary node. As a result, the automatic search for the primary node may fail. • When any of the CPU, memory, and I/O load approaches 100%, the connection is slow, which may cause connection timeout. You can locate the fault as follows: 1. Log in to a physical machine with slow connections or use a management tool to query the resource load. You can run the top command to check the CPU usage, run the free command to check the memory usage, and run the iostat command to check the I/O load. In addition, you can check the monitoring logs in the CM Agent and the monitoring records on the database O&M platform. 2. For peak load scenarios caused by a large number of slow queries in a short period of time, you can use the port specified by <i>[Port number of the database server + 1]</i> to query the pg_stat_activity view. For slow queries, you can use the system function pg_terminate_backend(pid int) to kill sessions. 3. If service overloading exists for a long time (that is, there is no obvious slow query, or new queries still become slow after slow queries are killed), reduce the service load and increase database resources.	

Property	Description	Value
connectTimeout	Specifies the timeout interval for connecting to a server OS. If the time taken to connect to a server OS exceeds the value specified, the connection is interrupted. When multiple IP addresses are configured in the URL, this parameter indicates the timeout interval for connecting to a single IP address.	Property type: integer Unit: s Value range: 0 to 2147483647. The value 0 indicates that the timeout mechanism does not take effect. Default value: 0
socketTimeout	Specifies the timeout period for a socket read operation. If the time taken to read data streams from a server exceeds the value specified, the connection is closed. If this parameter is not set, the client waits for a long time when the database process is abnormal. You are advised to set this parameter based on the SQL execution time acceptable to services. NOTE When the timeout is triggered on the JDBC side and the connection is closed, the running services delivered by the JDBC to the database are forcibly terminated. The capability of forcibly terminating a service is controlled by the GUC parameter check_disconnect_query . If this parameter is set to on , the capability is supported. If this parameter is set to off , the capability is not supported.	Property type: integer Unit: s Value range: 0 to 2147483647. The value 0 indicates that the timeout mechanism does not take effect. Default value: 0
socketTimeoutInConnecting	Specifies the timeout interval for a socket read operation during the connection establishment. If the time of reading data streams from the server exceeds the threshold, it attempts to search for the next node for connection.	Property type: integer Unit: s Value range: 0 to 2147483647. Default value: 5

Property	Description	Value
cancelSignalTimeout	Controls "connect timeout" and "socket timeout" of a cancel command, which may cause blocking. If the cancel command does not respond within the specified time, the connection is interrupted to reduce the occupation of client resources.	Property type: integer Unit: s Value range: 0 to 2147483647. Default value: 10
tcpKeepAlive	Specifies whether to enable TCP keepalive detection.	Property type: Boolean Value range: • true: TCP keepalive detection is enabled. • false: TCP keepalive detection is disabled. Default value: false
logUnclosedConnections	The client may leak a connection object because it does not call the close() method to close the connection object. Such object will be recycled and finalized using the finalize() method.	Property type: Boolean Value range: • true: If the close() method is not called, the finalize() method will be used to close the corresponding connection object. • false: If the close() method is not called, the corresponding connection object will not be closed. Default value: false
assumeMinServerVersion (Deprecated)	Specifies the version of the server to be connected. If the value of assumeMinServerVersion is greater than or equal to 9.0 , the number of packets to be sent are reduced during connection establishment. The client does not send a request to set the floating point to 3 (that is, retain the original floating point 2).	Property type: string Default value: No default value is provided.

Property	Description	Value
ApplicationName	Specifies the name of the JDBC driver that is being connected. You can query the pg_stat_activity table on the primary database node to view information about the client that is being connected. The JDBC driver name is displayed in the application_name column.	Property type: string Default value: PostgreSQL JDBC Driver
connectionExtraInfo	Specifies whether the driver reports the driver deployment path, process owner, and URL connection configuration information to the database.	Property type: Boolean Value range: • true: The JDBC driver reports the driver deployment path, process owner, and URL connection configuration information to the database and displays the information in the connection_info parameter. In this case, you can query the information from PG_STAT_ACTIVITY. • false: The driver does not report the driver deployment path, process owner, and URL connection configuration information to the database. Default value: false
autosave	Specifies the action that the driver should perform upon a query failure.	Property type: string Value range: always , never , and conservative . • always: The JDBC driver sets a savepoint before each query and rolls back to the savepoint if the query fails. • never: There is no savepoint. • conservative: A savepoint is set for each query. However, the system rolls back and retries only when there is an invalid statement. Default value: never

Property	Description	Value
protocolVersion	Specifies the connection protocol version.	Property type: integer Value range: 1 and 3. - If it is set to 1, only the V1 server is connected. - If it is set to 3, only the V5 server is connected. Default value: No default value is provided. If this property is not set, the value 3 is used.
prepareThreshold	Specifies the number of times that the PreparedStatement object is executed before the parsed statement on the server is used. NOTE You are advised not to use JDBC and prepareStatement to execute password-related statements (for example, CREATE USER user_name WITH PASSWORD '*****'). This is because when the number of execution times reaches the value specified by prepareThreshold, the database caches SQL statements but does not cache the passwords for security purposes. When the prepareStatement is executed again, error message "Password must contain at least 8 characters." will be displayed.	Property type: integer Value range: 0 to 2147483647. Default value: 5. The default value indicates that when the same PreparedStatement object is executed for five or more times, the parse message is not sent to the server to parse the statement. Instead, the statement that has been parsed on the server is used.
preparedStatementCacheQueries	Specifies the maximum number of queries generated by caching preparedStatement objects in each connection.	Property type: integer Value range: 0 to 2147483647. The value 0 indicates that the cache function is disabled. Default value: 256. If more than 256 different queries are used in the prepareStatement() call, the least recently used query cache will be discarded. Suggestion: Set the value based on service requirements. This parameter is used together with prepareThreshold.

Property	Description	Value
preparedStatementCacheSizeMiB	Specifies the maximum size of queries generated by caching preparedStatement objects in each connection.	Property type: integer Unit: MB Value range: 0 to 2147483647. The value 0 indicates that the cache function is disabled. Default value: 5 . If the size of the cached queries exceeds 5 MB, the least recently used query cache will be discarded.
databaseMetadataCacheFields	Specifies the maximum number of columns that can be cached in each connection.	Property type: integer Value range: 0 to 2147483647. The value 0 indicates that the cache function is disabled. Default value: 65536
databaseMetadataCacheFieldsMiB	Specifies the maximum size of columns that can be cached in each connection.	Property type: integer Unit: MB Value range: 0 to 2147483647. The value 0 indicates that the cache function is disabled. Default value: 5
stringtype	Specifies the type of the parameter passed to the database when the java.sql.PreparedStatement#setString method is called.	Property type: string Value range: unspecified and varchar . • varchar: Parameters are sent to the server as varchar parameters. • unspecified: Parameters are sent to the server without their types specified, and the server will attempt to infer the appropriate types. Default value: varchar

Property	Description	Value
batchMode	Specifies whether to use the batch mode. NOTE If batchMode is set to on, the data type of each column is the same as that specified by the first data record. If the data types are mixed, an error may be reported or the inserted data may be abnormal.	Property type: string Value range: • on: The batch mode is enabled to improve the batch update performance. If batchMode is set to on, the returned result is [count, 0, 0...0]. The first element in the array is the total number of records affected in batches. • off: The batch mode is disabled. If batchMode is set to off, the returned result is [1, 1, 1...1]. Each element in the array corresponds to the number of affected records in a single modification. Default value: on
fetchsize	Specifies the default fetchsize for statements in the created connection. Specifies the number of rows fetched from ResultSet each time. It has the same function as defaultRowFetchSize. If fetchsize and defaultRowFetchSize are set at the same time, fetchsize prevails.	Property type: integer Value range: 0 to 2147483647. Default value: 0 It indicates that all results are obtained from the database at a time. Suggestion: You are advised to set this parameter based on the amount of data queried by services and the memory of the client. When setting fetchsize, disable automatic commit (set autocommit to false). Otherwise, the setting of fetchsize does not take effect.
defaultRowFetchSize	Specifies the number of rows fetched from ResultSet each time. Limiting the number of rows read each time in a database access request can avoid unnecessary memory consumption, thereby avoiding the out of memory exception.	Property type: integer Value range: 0 to 2147483647. Default value: 0. It indicates that all results are obtained from the database at a time.

Property	Description	Value
rewriteBatchInserts	Specifies whether to rewrite SQL statements during batch insertion. When this parameter is used, batchMode must be set to off .	Property type: Boolean Value range: • true: <i>N</i> insert statements can be combined into one: insert into TABLE_NAME values(values1, ..., valuesN), ..., (values1, ..., valuesN) • false: SQL statements are not rewritten during batch insertion. Default value: false
unknownLength	Specifies the length of some GaussDB types (such as TEXT) whose length is unknown and not defined when they are returned by functions such as ResultSetMetaData.getColumnDisplaySize and ResultSetMetaData.getPrecision.	Property type: integer Value range: 0 to 2147483647 Default value: Integer.MAX_VALUE
uppercaseAttributeName	Converts the query result of the API for obtaining metadata to uppercase letters. The application scenario is as follows: All metadata stored in the database is in lowercase, but they must be used as the input and output parameters in uppercase. For details about the involved APIs, see java.sql.DatabaseMetaData and java.sql.ResultSetMetaData. NOTE After the uppercaseAttributeName parameter is enabled, if the database contains metadata with a mixture of uppercase and lowercase letters, only the metadata in lowercase letters can be queried and output in uppercase letters. Before using the metadata, ensure that the metadata is stored in lowercase letters to prevent data errors.	Property type: string Value range: • true: Metadata is converted to uppercase letters. • false: GUC parameter configuration is used. Default value: false

Property	Description	Value
binaryTransfer	Specifies whether data is sent and received in binary format.	Property type: Boolean Value range: true: enabled. false: disabled. Default value: false
binaryTransferEnable	Specifies types for which binary transmission is enabled. Types are separated by commas (,), for example, binaryTransferEnable=Integer4_ARRAY,Integer8_ARRAY . You can select either the OID or name. For example, if the name is BLOB and the OID is 88, set binaryTransferEnable to BLOB or 88 .	Property type: string Default value: No default value is provided.
binaryTransferDisable	Specifies types for which binary transmission is disabled. Types are separated by commas (,). You can select either the OID or name. If binaryTransferDisable and binaryTransferEnable have the same value, the disabling function is used.	Property type: string Default value: No default value is provided.
blobMode	Sets the setBinaryStream() method to assign values to different types of data. You are advised to set this parameter to on for systems migrated from database A or B and to off for systems migrated from database PG.	Property type: string Value range: on: Values are assigned to BLOB data. off: Values are assigned to bytea data. Default value: on
socketFactory	Specifies the name of the class used to create a socket connection with the server. This class must implement the javax.net.SocketFactory API and define a constructor with no parameter or a single string parameter.	Property type: string
socketFactoryArg	The value is an optional parameter of the constructor function of the socketFactory class and is not recommended.	Property type: string

Property	Description	Value
receiveBufferSize	Sets SO_RCVBUF on a connection stream.	Property type: integer Unit: byte Value range: -1 to 2147483647 Default value: -1 . It indicates that the buffer size is not set.
sendBufferSize	Sets SO_SNDBUF on a connection stream.	Property type: integer Unit: byte Value range: -1 to 2147483647 Default value: -1 . It indicates that the buffer size is not set.
preferQueryMode	Specifies the query mode.	Property type: string Value range: simple , extended , extendedForPrepared , and extendedCacheEverything . - In simple mode, the query is executed without parsing or binding. - In extended mode, the query is executed and bound. - The extendedForPrepared mode is used for prepared statement extension. - In extendedCacheEverything mode, each statement is cached. Default value: extended

Property	Description	Value
targetServerType	Identifies the primary DN and standby DN by querying whether a DN allows the write operation in the URL connection string.	Property type: string Value range: any, master, slave, preferSlave, and clusterMainNode. • master: The system attempts to connect to the IP addresses configured in the URL connection string in sequence until the primary DN in the database instance is connected. If the primary DN cannot be found, an exception is thrown. • slave: The system attempts to connect to the IP addresses configured in the URL connection string in sequence until a standby DN in the database instance is connected. If no standby DN can be found, an exception is thrown. • preferSlave: The system prefers to connect to a standby DN (if available) in the URL connection string; if no standby DN can be found, it connects to the primary DN. • any: The system attempts to connect to any DN in the URL connection string. • clusterMainNode: The system attempts to connect to the primary DN or main standby DN in the URL string. If the primary DN or main standby DN cannot be found, an exception is thrown. Default value: any Query statement: select local_role, db_state from pg_stat_get_stream_replications(); Suggestion: You are advised to set this parameter to master for services with write

Property	Description	Value
		operations to ensure that the primary DN can be properly connected after a switchover. However, if the standby DN is not completely promoted to primary during the switchover, the connection cannot be established. As a result, service statements cannot be executed.
priorityServers	Specifies the first n nodes configured in the URL as the primary database instance to be connected preferentially. It is used in streaming DR scenarios. Example: jdbc:opengauss://host1:port1,host2:port2,host3:port3,host4:port4/database?priorityServers=2 . That is, host1 and host2 are primary database instance nodes, and host3 and host4 are standby database instance nodes for DR.	Property type: integer Value range: a number greater than 0 and less than the number of DNs configured in the URL. Default value: NULL
forceTargetServerSlave	Specifies whether to enable the function of forcibly connecting to the standby node and forbid the existing connections to the standby node that is promoted to primary after a switchover of the database instance.	Property type: Boolean Value range: • false: The function of forcibly connecting to the standby node is disabled. • true: The function of forcibly connecting to the standby node is enabled. Default value: false
tracerInterfaceClass	Specifies the fully qualified name of the implementation class of the com.huawei.opengauss.jdbc.log.Tracer API that implements the method for obtaining traceld .	Property type: string Default value: NULL
use_boolean	Sets the OID type bound to the setBoolean() method in extended mode.	Property type: Boolean Value range: • false: The int2 type is bound. • true: The Boolean type is bound. Default value: false

Property	Description	Value
allowRead Only	Specifies whether to enable the read-only mode for a connection.	Property type: Boolean Value range: • true: The read-only mode is enabled. • false: The read-only mode is disabled. In this case, calling connection.setReadOnly(true) does not take effect, and data can still be modified. Default value: true
TLS_CIPHERS	Specifies the supported TLS cipher suite.	Property type: string Default value: TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
stripTrailingZeros	Specifies whether trailing 0s of the numeric type are removed. It is valid only for ResultSet.getObject(int columnIndex).	Property type: Boolean Value range: • true: Trailing 0s of the numeric type are removed. • false: Trailing 0s of the numeric type are not removed. Default value: false
enableTimeZone	Specifies whether to enable the time zone setting on the client.	Property type: Boolean Value range: • true: The time zone setting on the client is enabled. The JVM time zone is obtained to specify the database time zone. • false: The time zone setting on the client is disabled. Instead, the database time zone is used. Default value: true

Property	Description	Value
loadBalanceHosts	Specifies whether to enable the load balancing function. In a centralized environment, ensure that no write operation is in services when using this parameter.	Property type: Boolean Value range: • true: The shuffle algorithm is used to randomly select a host from the candidates to establish a connection. • false: Multiple hosts specified in the URL are connected in sequence. Default value: false

5.3.2.2.3 Connecting to a Database in Non-Encrypted Mode

To connect to a database in non-encrypted mode, you need to load a driver and then establish a database connection. This section describes methods of loading drivers, APIs for creating database connections, and APIs for non-encrypted connections.

Methods of Loading Drivers

You can load drivers in the following ways:

- Implicitly load a driver at any position in code before creating a connection.
`Class.forName("com.huawei.opengauss.jdbc.Driver");`
- Pass parameters when a JVM is started. **jdbctest** is the name of the test case program.
`java -Djdbc.drivers=com.huawei.opengauss.jdbc.Driver jdbctest`

NOTE

1. GaussDB is compatible with PG in the use of JDBC. Therefore, when two JDBC drivers are used in the same process, class names may conflict.
2. Compared with the PG driver, the GaussDB JDBC driver has the following enhanced features:
 1. SHA-256 is supported for encrypted login.
 2. The third-party log framework that implements the sf4j API can be connected.
 3. DR failover is supported.

APIs for Creating Database Connections

JDBC provides three APIs for creating database connections. For details about the **url**, **info**, **user**, and **password** parameters, see [Table 5-6](#).

API 1: `DriverManager.getConnection(String url)`. You need to write a database username and password in a URL, which is insecure and not recommended.

API 2: `DriverManager.getConnection(String url, String user, String password)`. For details, see [Connecting to a Database Using API 2](#).

API 3: `DriverManager.getConnection(String url, Properties info)`. For details, see [Connecting to a Database Using API 3](#).

Table 5-6 Database connection parameters

Parameter	Description
url	Database connection descriptor when opengaussjdbc.jar is used. The format is as follows: • jdbc:opengauss:database • jdbc:opengauss://host/database • jdbc:opengauss://host:port/database • jdbc:opengauss://host:port/database?param1=value1&param2=value2 • jdbc:opengauss://host1:port1,host2:port2/database?param1=value1&param2=value2 NOTE • database indicates the name of the database to connect. • host indicates the name or IP address of the database server. For security purposes, the primary database node forbids access from other nodes in the database without authentication. To access the primary database node from inside the database, deploy the JDBC program on the host where the primary database node is located and set host to 127.0.0.1. Otherwise, the error message "FATAL: Forbid remote connection with trust method!" may be displayed. It is recommended that the service system be deployed outside the database. If it is deployed inside, database performance may be affected. By default, the localhost is used to connect to the server. • port indicates the port number of the database server. By default, the database on port 5431 of the localhost is connected. • param indicates a database connection attribute. The parameter can be configured in the URL. The URL starts with a question mark (?), uses an equal sign (=) to assign a value to the parameter, and uses an ampersand (&) to separate parameters. • value indicates the database connection attribute values. • The connectTimeout and socketTimeout parameters must be set for connection. If they are not set, the default value 0 is used, indicating that the connection will not time out. When the network between a DN and the client is faulty, the client cannot receive the ACK packet from the DN and retries transmission repeatedly. A timeout error is reported only when the number of retransmission times reaches 15 (the default value). As a result, the RTO is high. • You are advised to ensure the validity of the URL when using the standard JDBC API to establish a connection. An invalid URL may cause an exception, and the exception contains the original URL character string, which may cause sensitive information leakage.
info	Database connection property. For details about all parameters, see Connection Parameter Reference .
user	Database user.
password	Password of the database user.

Connecting to a Database Using API 2

Use the `DriverManager.getConnection(String url, String user, String password)` API to establish a database connection. The commands are as follows:

Step 1 Import `java.sql.Connection` and `java.sql.DriverManager`.

`java.sql.Connection` is a database connection API. The `getConnection()` method of `java.sql.DriverManager` is used to connect applications to a database. In addition, you need to import other APIs and classes based on the actual application scenario. For details, see [JDBC API Reference](#).

```
import java.sql.Connection;
import java.sql.DriverManager;
```

Step 2 Specify the database **sourceURL** (change *\$ip*, *\$port*, and *database* as required), username, and password.

Writing the username and password to code has great security risks. You are advised to store the username and password in environment variables.

```
String sourceURL = "jdbc:opengauss://$ip:$port/database";
String userName = System.getenv("EXAMPLE_USERNAME_ENV");
String password = System.getenv("EXAMPLE_PASSWORD_ENV");
```

Step 3 Load the driver.

1. Add the **opengaussjdbc.jar** package to the code running tool (such as IDE).
2. Load the database driver **com.huawei.opengauss.jdbc.Driver** as follows:

```
String driver = "com.huawei.opengauss.jdbc.Driver";
Class.forName(driver);
```

Step 4 Establish a database connection.

Call `DriverManager.getConnection(String url, String user, String password)` to connect to the database.

```
Connection conn = DriverManager.getConnection(sourceURL, userName, password);
```

----End

Connecting to a Database Using API 3

Use the `DriverManager.getConnection(String url, Properties info)` API to establish a database connection. The commands are as follows:

Step 1 Import `java.sql.Connection`, `java.sql.DriverManager`, and `java.util.Properties`.

The `setProperty()` method of `java.util.Properties` is used to set property values of a `Properties` object. In addition, you need to import other APIs and classes based on the actual application scenario. For details, see [JDBC API Reference](#).

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.util.Properties;
```

Step 2 Specify the database **sourceURL** (change *\$ip*, *\$port*, and *database* as required), username, and password.

Writing the username and password to code has great security risks. You are advised to store the username and password in environment variables.

```
String sourceURL = "jdbc:opengauss://$ip:$port/database";
String userName = System.getenv("EXAMPLE_USERNAME_ENV");
String password = System.getenv("EXAMPLE_PASSWORD_ENV");
```

Step 3 Create a Properties object and set **userName** and **password** as the property values of the object.

```
Properties info = new Properties();  
info.setProperty("user", userName);  
info.setProperty("password", password);
```

Step 4 Load the driver.

1. Add the **opengaussjdbc.jar** package to the code running tool (such as IDE).
2. Load the database driver **com.huawei.opengauss.jdbc.Driver** as follows:

```
String driver = "com.huawei.opengauss.jdbc.Driver";  
Class.forName(driver);
```

Step 5 Establish a database connection.

Call `DriverManager.getConnection(String url, Properties info)` to connect to the database.

```
Connection conn = DriverManager.getConnection(sourceURL, info);
```

----End

5.3.2.2.4 Connecting to a Database in SSL Mode

When establishing connections to the GaussDB server using JDBC, you can enable SSL connections to encrypt client and server communications for security of sensitive data transmission on the Internet. You can connect to a database in SSL mode using the **NonValidatingFactory** channel or certificate-based authentication. In certificate-based authentication mode, a client and a server authenticate each other. In this section, the `DriverManager.getConnection(String url, Properties info)` API is used to connect to a database.

Method 1: NonValidatingFactory Channel

Prerequisites: You have obtained the certificates and private key file required by the server and configured the server. For details about how to generate and obtain the certificates and configure the server, contact an administrator or refer to related OpenSSL documents and commands.

Connect to a database through the **NonValidatingFactory** channel as follows:

Step 1 Import `java.sql.Connection`, `java.sql.DriverManager`, and `java.util.Properties`.

In addition, you need to import other APIs and classes based on the actual application scenario. For details, see [JDBC API Reference](#).

```
import java.sql.Connection;  
import java.sql.DriverManager;  
import java.util.Properties;
```

Step 2 Specify the database **sourceURL** (change *\$ip*, *\$port*, and *database* as required), username, and password.

Writing the username and password to code has great security risks. You are advised to store the username and password in environment variables.

```
String sourceURL = "jdbc:opengauss://$ip:$port/database";  
Properties urlProps = new Properties();  
urlProps.setProperty("user", System.getenv("EXAMPLE_USERNAME_ENV"));  
urlProps.setProperty("password", System.getenv("EXAMPLE_PASSWORD_ENV"));
```

Step 3 Set the SSL property to **true** to use the **NonValidatingFactory** channel.

```
urlProps.setProperty("ssl", "true");  
urlProps.setProperty("sslfactory", "com.huawei.opengauss.jdbc.ssl.NonValidatingFactory");
```

Step 4 Load the driver.

1. Add the **opengaussjdbc.jar** package to the code running tool (such as IDE).
2. Load the database driver **com.huawei.opengauss.jdbc.Driver** as follows:

```
Class.forName("com.huawei.opengauss.jdbc.Driver");
```

Step 5 Establish a database connection.

Call `DriverManager.getConnection(String url, Properties info)` to connect to the database.

```
Connection conn = DriverManager.getConnection(sourceURL, urlProps);
```

----End

Method 2: Certificate-based Authentication

Prerequisites: You have obtained the certificates and private key file required by the server and configured the server. You have obtained the **client.crt** client certificate, **cacert.pem** root certificate, and **client.key.pk8** client private key file required by the client. [Step 3](#) describes how to configure the certificates and private key file on the client. For details about how to generate and obtain the certificates and configure the server, contact an administrator or refer to related OpenSSL documents and commands.

Configure certificates on the client to connect to a database as follows:

Step 1 Import `java.sql.Connection`, `java.sql.DriverManager`, and `java.util.Properties`.

In addition, you need to import other APIs and classes based on the actual application scenario. For details, see [JDBC API Reference](#).

```
import java.sql.Connection;  
import java.sql.DriverManager;  
import java.util.Properties;
```

Step 2 Specify the database **sourceURL** (change *\$ip*, *\$port*, and *database* as required), username, and password.

Writing the username and password to code has great security risks. You are advised to store the username and password in environment variables.

```
String sourceURL = "jdbc:opengauss://$ip.$port/database";  
Properties urlProps = new Properties();  
urlProps.setProperty("user", System.getenv("EXAMPLE_USERNAME_ENV"));  
urlProps.setProperty("password", System.getenv("EXAMPLE_PASSWORD_ENV"));
```

Step 3 Set the SSL property to **true** and configure the **client.crt** client certificate, **cacert.pem** root certificate, and **client.key.pk8** client private key on the client.

```
urlProps.setProperty("ssl", "true");  
urlProps.setProperty("sslcert", "client.crt");  
urlProps.setProperty("sslrootcert", "cacert.pem");  
urlProps.setProperty("sslkey", "client.key.pk8");
```

 NOTE

Before using the client private key file, convert **client.key** to **client.key.pk8**.

```
/**
 * openssl pkcs8 -topk8 -outform DER -in client.key -out client.key.pk8 -nocrypt
 * openssl pkcs8 -topk8 -inform PEM -in client.key -outform DER -out client.key.der -v1 PBE-MD5-DES
 * openssl pkcs8 -topk8 -inform PEM -in client.key -outform DER -out client.key.der -v1 PBE-SHA1-3DES
 * The preceding algorithms are not recommended due to their low security.
 * If the customer needs to use a higher-level private key encryption algorithm, the following private
 * key encryption algorithms can be used after the BouncyCastle or a third-party private key is used to
 * decrypt the password package:
 * openssl pkcs8 -in client.key -topk8 -outform DER -out client.key.der -v2 AES128
 * openssl pkcs8 -in client.key -topk8 -outform DER -out client.key.der -v2 aes-256-cbc -iter 1000000
 * openssl pkcs8 -in client.key -topk8 -out client.key.der -outform Der -v2 aes-256-cbc -v2prf
 * hmacWithSHA512
 * Enable BouncyCastle: Introduce the bcpkix-jdk15on.jar package for projects that use JDBC. The
 * recommended version is 1.65 or later.
 */
```

Step 4 Configure **sslmode**.

Set **sslmode** to **require**, **verify-ca**, or **verify-full**. For details about the parameters, see [sslmode](#). You can select one of them based on the application scenario.

```
/* Set sslmode to require. */
urlProps.setProperty("sslmode", "require");
/* Set sslmode to verify-ca. */
urlProps.setProperty("sslmode", "verify-ca");
/* Set sslmode to verify-full (verification in Linux). */
urlProps.setProperty("sslmode", "verify-full");
```

Step 5 Load the driver.

1. Add the **opengaussjdbc.jar** package to the code running tool (such as IDE).
2. Load the database driver **com.huawei.opengauss.jdbc.Driver** as follows:

```
Class.forName("com.huawei.opengauss.jdbc.Driver");
```

Step 6 Establish a database connection.

Call `DriverManager.getConnection(String url, Properties info)` to connect to the database.

```
Connection conn = DriverManager.getConnection(sourceURL,urlProps);
```

----End

5.3.2.2.5 Connecting to a Database in UDS Mode

The UDS is used for data exchange between different processes on the same host. You can add **unixsocket** to obtain the socket factory.

Prerequisites: You have referenced **unixsocket-core-XXX.jar**, **unixsocket-common-XXX.jar**, and **unixsocket-native-common-XXX.jar**. XXX indicates the version number. The version numbers of the referenced JAR packages must be consistent.

To connect to a database in UDS mode, perform the following steps:

Step 1 Import `java.sql.Connection`, `java.sql.DriverManager`, and `java.util.Properties`.

In addition, you need to import other APIs and classes based on the actual application scenario. For details, see [JDBC API Reference](#).

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.util.Properties;
```

Step 2 Specify the username and password of the database.

Writing the username and password to code has great security risks. You are advised to store the username and password in environment variables. Set the username and password to the property values of the Properties object.

```
String userName = System.getenv("EXAMPLE_USERNAME_ENV");
String password = System.getenv("EXAMPLE_PASSWORD_ENV");
Properties properties = new Properties();
properties.setProperty("user", userName);
properties.setProperty("password", password);
```

Step 3 Load the driver.

1. Add the **opengaussjdbc.jar** package to the code running tool (such as IDE).
2. Load the database driver **com.huawei.opengauss.jdbc.Driver** as follows:

```
String driver = "com.huawei.opengauss.jdbc.Driver";
Class.forName(driver);
```

Step 4 Specify the database *\$ip*, *\$port*, *database*, **socketFactory**, and **socketFactoryArg**.

Change *\$ip*, *\$port* as required. Set *database* to **localhost**, **socketFactory** to **org.newsclub.net.unix.AFUNIXSocketFactory\$FactoryArg**, and **socketFactoryArg** to *[path-to-the-unix-socket]* to connect to the database in UDS mode. For details about the **socketFactory** and **socketFactoryArg** parameters, see [socketFactory](#) and [socketFactoryArg](#).

 NOTE

Set the **socketFactoryArg** parameter based on the actual path. The value must be the same as that of the GUC parameter **unix_socket_directory**.

```
Connection conn = DriverManager.getConnection("jdbc:opengauss://$ip:$port/database?
socketFactory=org.newsclub.net.unix" +
".AFUNIXSocketFactory$FactoryArg&socketFactoryArg=[path-to-the-unix-socket]",properties);
System.out.println("Connection Successful!");
```

----End

5.3.2.3 Running SQL Statements

In this section, you can [execute basic SQL statements](#) to create the **customer_t1** table, [execute a prepared statement](#) to insert data in batches, [execute a prepared statement](#) to update data, and [create and call stored procedures](#).

Running a Common SQL Statement

SQL statements are run on applications to operate a database. Operations such as SELECT, UPDATE, INSERT, and DELETE can be performed on XML data.

The prerequisite is that you have connected to the database using the connection object **conn**. Execute basic SQL statements to create the **customer_t1** table as follows:

Step 1 Create a statement object by calling the `createStatement` method of the Connection API.

```
Statement stmt = conn.createStatement();
```

Step 2 Run the SQL statement by calling the `executeUpdate` method of the Statement API.


```
int rc = stmt.executeUpdate("CREATE TABLE customer_t1(c_customer_sk INTEGER, c_customer_name  
VARCHAR(32));");
```

Step 3 Close the statement object **stmt** by calling the close method of the Statement API.

```
stmt.close();
```

----End

NOTE

- If an execution request (not in a transaction block) received in the database contains multiple statements, the request is packed into a transaction. **VACUUM** is not supported in a transaction block. If one of the statements fails, the entire request will be rolled back.
- Use semicolons (;) to separate statements. Stored procedures, functions, and anonymous blocks do not support multi-statement execution. When **preferQueryMode** is set to **simple**, the statement does not execute the parsing logic, and the semicolons (;) cannot be used to separate statements in this scenario.
- The slash (/) can be used as the terminator for creating a single stored procedure, function, anonymous block, or package body. When **preferQueryMode** is set to **simple**, the statement does not execute the parsing logic, and the slash (/) cannot be used as the terminator in this scenario.
- JDBC caches SQL statements in PreparedStatement, which may cause memory bloat. If the JVM memory is small, you are advised to adjust **preparedStatementCacheSizeMiB** or **preparedStatementCacheQueries**.

Executing a Prepared Statement for Insertion

When a prepared statement processes multiple pieces of similar data, the database creates only one execution plan. This improves compilation and optimization efficiency.

The prerequisite is that the **customer_t1** table has been created by executing the preceding basic SQL statements. Execute the prepared statement to insert data in batches as follows:

Step 1 Create the prepared statement object **pst** by calling the prepareStatement method of the Connection API.

```
PreparedStatement pst = conn.prepareStatement("INSERT INTO customer_t1 VALUES (?,?)");
```

Step 2 Call the corresponding API to set parameters for each piece of data and call addBatch to add SQL statements to the batch processing.

```
for (int i = 0; i < 3; i++) {  
    pst.setInt(1, i);  
    pst.setString(2, "data " + i);  
    pst.addBatch();  
}
```

Step 3 Perform batch processing by calling the executeBatch method of the PreparedStatement API.

```
pst.executeBatch();
```

Step 4 Close the prepared statement object **pst** by calling the close method of the PreparedStatement API.

```
pst.close();
```

----End

 NOTE

Do not terminate a batch processing action when it is ongoing; otherwise, database performance will deteriorate. Therefore, disable automatic commit during batch processing. Manually commit several lines at a time. The statement for disabling automatic commit is as follows:

```
conn.setAutoCommit(false);
```

Running a Prepared SQL Statement

Prepared statements are compiled and optimized once but can be used in different scenarios by assigning different values. Using prepared statements improves execution efficiency. If you want to run a statement for several times, use a prepared statement.

The prerequisite is that the preceding prepared statement has been executed and data has been inserted into the **customer_t1** table in batches. Execute the prepared SQL statement to update data as follows:

- Step 1** Create the prepared statement object **pstmt** by calling the `prepareStatement` method of the Connection API.

```
PreparedStatement pstmt = conn.prepareStatement("UPDATE customer_t1 SET c_customer_name = ?  
WHERE c_customer_sk = 1");
```

- Step 2** Set parameters by calling the `setString` method of the PreparedStatement API.

```
pstmt.setString(1, "new Data");
```

- Step 3** Execute the prepared statement by calling the `executeUpdate` method of the PreparedStatement API.

```
int rowcount = pstmt.executeUpdate();
```

- Step 4** Close the prepared statement object **pstmt** by calling the `close` method of the PreparedStatement API.

```
pstmt.close();
```

----End

NOTICE

After binding parameters are set in `PreparedStatement`, a B packet or U packet is constructed and sent to the server when the SQL statement is executed. However, the maximum length of a B packet or U packet cannot exceed 1023 MB. If the data bound at a time is too large, an exception may occur because the packet is too long. Therefore, when setting binding parameters in `PreparedStatement`, you need to evaluate and control the size of the bound data to avoid exceeding the upper limit of the packet.

Creating and Calling a Stored Procedure

GaussDB can call stored procedures through JDBC. The prerequisite is that the database connection is established using the connection object `conn`.

Create the **testproc** stored procedure as follows:

```
// Create the stored procedure (containing the out parameter) in the database as follows:  
create or replace procedure testproc
```

```
(
    psv_in1 in integer,
    psv_in2 in integer,
    psv_inout inout integer
)
as
begin
    psv_inout := psv_in1 + psv_in2 + psv_inout;
end;
/
```

Call the **testproc** stored procedure as follows:

Step 1 Create a call statement object **cstmt** by calling the `prepareCall` method of `Connection`.

```
CallableStatement cstmt = conn.prepareCall("{? = CALL TESTPROC(?,?,?)}");
```

Step 2 Set parameters by calling the `setInt` method of `CallableStatement`.

```
cstmt.setInt(2, 50);
cstmt.setInt(1, 20);
cstmt.setInt(3, 90);
```

Step 3 Register an output parameter by calling the **registerOutParameter** method in **CallableStatement**.

```
cstmt.registerOutParameter(4, Types.INTEGER); // Register an OUT parameter of the integer type.
```

Step 4 Execute SQL statements by calling the `execute` method of `CallableStatement`.

```
cstmt.execute();
```

Step 5 Obtain the out parameter by calling the `getInt` method of `CallableStatement`.

```
int out = cstmt.getInt(4);
```

Step 6 Close the call statement object **cstmt** by calling the `close` method of `CallableStatement`.

```
cstmt.close();
```

----End

NOTICE

- If JDBC is used to call a stored procedure whose returned value is a cursor, the returned cursor cannot be used.
- A stored procedure and a basic SQL statement must be run separately.
- Output parameters must be registered for parameters of the `inout` type in the stored procedure.

 NOTE

- Many database classes such as Connection, Statement, and ResultSet have a close() method. Close these classes after using their objects. Closing Connection objects is to close all the related Statement objects, and closing a Statement object is to close its ResultSet object.
- Some JDBC drivers support named parameters, which can be used to set parameters by name rather than sequence. If a parameter has the default value, you do not need to specify any parameter value but can use the default value directly. Even though the parameter sequence changes during a stored procedure, the application does not need to be modified. Currently, the GaussDB JDBC driver does not support this method.
- GaussDB does not support functions containing OUT parameters, or stored procedures and function parameters containing default values.
- When you bind parameters in `conn.prepareStatement("{? = CALL TESTPROC(?,?,?)})"` during a stored procedure calling, you can bind parameters and register the first parameter as the output parameter according to the placeholder sequence or the fourth parameter as the output parameter according to the parameter sequence in the stored procedure. The preceding example registers the fourth parameter.

Creating and Calling a Stored Procedure (Input Parameters of the Composite Data Type)

The following shows how to create and call a stored procedure whose input parameters are of the composite data type in A-compatible mode. The prerequisite is that the database connection is established using the connection object `conn`.

Create the **test_proc** stored procedure as follows:

```
// Create a composite data type in the database.
CREATE TYPE compfoo AS (f1 int, f3 text);
// Create the following stored procedure (containing the out parameter) in the database:
create or replace procedure test_proc
(
    psv_in in compfoo,
    psv_out out compfoo
)
as
begin
    psv_out := psv_in;
end;
/
```

Call the **test_proc** stored procedure as follows:

Step 1 After setting **behavior_compat_options** to '**proc_outparam_override**', create the call statement object **cs** by calling the `prepareCall` method of Connection.

```
Statement statement = conn.createStatement();
statement.execute("set behavior_compat_options='proc_outparam_override'");
CallableStatement cs = conn.prepareCall("{ CALL TEST_PROC(?,?) }");
```

Step 2 Set parameters by calling the `set` method of CallableStatement.

```
PObject pObject = new PObject();
pObject.setType("public.compfoo"); // Set the composite type name. The format is "schema.typename".
pObject.setValue("1,demo"); // Bind the value of the composite type. The format is "(value1,value2)".
cs.setObject(1, pObject);
```

Step 3 Register an output parameter by calling the **registerOutParameter** method in CallableStatement.

```
//Register an out parameter of the composite type. The format is "schema.typename".
cs.registerOutParameter(2, Types.STRUCT, "public.compfoo");
```

Step 4 Execute SQL statements by calling the execute method of CallableStatement.

```
cs.execute();
```

Step 5 Obtain the output parameter by calling the **getObject** method in **CallableStatement**.

```
PObject result = (PObject)cs.getObject(2); // Obtain the out parameter.  
result.getValue(); // Obtain the string value of the composite type.  
result.getArrayValue(); // Obtain the array values of the composite type and sort the values according to  
the sequence of columns of the composite type.  
result.getStruct(); // Obtain the subtype names of the composite type and sort them according to the  
creation sequence.
```

Step 6 Close the call statement object **cs** by calling the close method of CallableStatement.

```
cs.close();
```

----End

NOTE

- After the A-compatible mode is enabled, you must use the {call proc_name(?,?) format to call a stored procedure and use the {? = call func_name(?,?) format to call a function. The question mark (?) on the left of the equal mark is the placeholder for the return value of the function and is used to register the return value of the function.
- After setting **behavior_compat_options** to 'proc_outparam_override', re-establish the connection. Otherwise, the stored procedures and functions cannot be correctly called.
- If a function or stored procedure contains a composite type, bind and register parameters in the schema.typename format.

5.3.2.4 Processing Data in a Result Set

After running SQL statements, you need to process the result set. This section describes how to set the result set type, locate the result set, obtain the cursor position from a result set, and obtain data from the result set.

Setting a Result Set Type

Different types of result sets apply to different application scenarios. Applications select proper types of result sets based on requirements. Before running an SQL statement, you must create a statement object. Some methods of creating statement objects can set the type of a result set. [Table 5-7](#) lists result set parameters. The related Connection methods are as follows:

```
// Create a Statement object. This object will generate a ResultSet object with a specified type and  
concurrency.  
createStatement(int resultSetType, int resultSetConcurrency);  
  
// Create a PreparedStatement object. This object will generate a ResultSet object with a specified type and  
concurrency.  
prepareStatement(String sql, int resultSetType, int resultSetConcurrency);  
  
// Create a CallableStatement object. This object will generate a ResultSet object with a specified type and  
concurrency.  
prepareCall(String sql, int resultSetType, int resultSetConcurrency);
```

Table 5-7 Result set types

Parameter	Description
resultSetType	Type of a result set. There are three types of result sets: • ResultSet.TYPE_FORWARD_ONLY: The ResultSet object can only be navigated forward. It is the default value. • ResultSet.TYPE_SCROLL_SENSITIVE: You can view the modified result by scrolling to the modified row. • ResultSet.TYPE_SCROLL_INSENSITIVE: The ResultSet object is insensitive to changes in the underlying data source. NOTE After a result set has obtained data from the database, the result set is insensitive to data changes made by other transactions, even if the result set type is ResultSet.TYPE_SCROLL_SENSITIVE. To obtain up-to-date data of the record pointed by the cursor from the database, call the <code>refreshRow()</code> method in a ResultSet object.
resultSetConcurrency	Concurrency type of a result set. There are two types of concurrency. • ResultSet.CONCUR_READ_ONLY: Data in a result set cannot be updated except that an updated statement has been created in the result set data. • ResultSet.CONCUR_UPDATEABLE: changeable result set. The concurrency type for a result set object can be updated if the result set is scrollable.

Positioning a Cursor in a Result Set

ResultSet objects include a cursor pointing to the current data row. The cursor is initially positioned before the first row. The `next` method moves the cursor to the next row from its current position. When a ResultSet object does not have a next row, a call to the `next` method returns **false**. Therefore, this method is used in the while loop for result set iteration. However, the JDBC driver provides more cursor positioning methods for scrollable result sets, which allows positioning cursor in the specified row. [Table 5-8](#) describes these methods.

Table 5-8 Methods for positioning a cursor in a result set

Method	Description
<code>next()</code>	Moves cursor to the next row from its current position.
<code>previous()</code>	Moves cursor to the previous row from its current position.
<code>beforeFirst()</code>	Places cursor before the first row.

Method	Description
afterLast()	Places cursor after the last row.
first()	Places cursor to the first row.
last()	Places cursor to the last row.
absolute(int row)	Places cursor to a specified row.
relative(int rows)	Moves the result set downwards by the number of rows specified by rows when rows is set to a positive number, or moves the result set upwards by the number of rows specified by rows when rows is set to a negative number.

Obtaining the Cursor Position from a Result Set

This cursor positioning method will be used to change the cursor position for a scrollable result set. The JDBC driver provides a method to obtain the cursor position in a result set. [Table 5-9](#) describes these methods.

Table 5-9 Methods for obtaining a cursor position in a result set

Method	Description
isFirst()	Checks whether it is in the first row.
isLast()	Checks whether it is in the last row.
isBeforeFirst()	Checks whether it is before the first row.
isAfterLast()	Checks whether it is after the last row.
getRow()	Obtains its current row number.

Obtaining Data from a Result Set

ResultSet objects provide a variety of methods to obtain data from a result set. [Table 5-10](#) describes the common methods for obtaining data. If you want to know more about other methods, see JDK official documents.

Table 5-10 Common methods for obtaining data from a result set

Method	Description
<code>getInt(int columnIndex)</code>	Retrieves the value of the column designated by a column index in the current row as an integer.
<code>getInt(String columnLabel)</code>	Retrieves the value of the column designated by a column label in the current row as an integer.
<code>getString(int columnIndex)</code>	Retrieves the value of the column designated by a column index in the current row as a string.
<code>getString(String columnLabel)</code>	Retrieves the value of the column designated by a column label in the current row as a string.
<code>getDate(int columnIndex)</code>	Retrieves the value of the column designated by a column index in the current row as a date.
<code>getDate(String columnLabel)</code>	Retrieves the value of the column designated by a column name in the current row as a date.

5.3.2.5 Closing a Database Connection

After you complete required data operations in the database, close the database connection.

You can call the close method to close the connection.

```
conn.close();
```

5.3.3 Typical Application Development Examples

5.3.3.1 Configuring Connection Parameters in Different Scenarios

NOTE

In the following example, **host:port** represents a node, where **host** indicates the name or IP address of the server where the database resides, and **port** indicates the port number of the server where the database resides.

DR

A customer has two database instances. Database instance A is the production database instance, and database instance B is the DR database instance. When the customer performs a DR failover, database instance A is demoted to the DR database instance, and database instance B is promoted the production database instance. In this case, to avoid application restart or re-release caused by

modifications on the configuration file, the customer can write database instances A and B to the connection string when initializing the configuration file. If the primary database instance cannot be connected, the driver attempts to connect to the DR database instance. For example, database instance A consists of *node1*, *node2*, and *node3*, and database instance B consists of *node4*, *node5*, and *node6*.

The URL can be configured as follows:

```
jdbc:opengauss://node1,node2,node3,node4,node5,node6/database?priorityServers=3
```

If you want to connect to both the primary database instance and hosts in the primary database instance, set **targetServerType** to **master**. The URL can be configured as follows:

```
jdbc:opengauss://node1,node2,node3,node4,node5,node6/database?  
priorityServers=3&targetServerType=master
```

Load Balancing

A customer has a centralized database instance that consists of one primary node and two standby nodes {*node1*, *node2*, *node3*}. **node1** is the primary node, and **node2** and **node3** are standby nodes.

If the customer wants to evenly distribute the connections established on the same application to three nodes, the URL can be configured as follows:

```
jdbc:opengauss://node1,node2,node3/database?loadBalanceHosts=true
```

CAUTION

When **loadBalanceHosts** is used, if the connection is established on the standby DN, write operations cannot be performed. If read and write operations are required, do not set this parameter.

Automatic Selection of the Primary Node

A customer has a centralized database instance that consists of one primary node and two standby nodes {*node1*, *node2*, *node3*}. **node1** is the primary node, and **node2** and **node3** are standby nodes.

If the customer requires that the application connection be established on the primary DN and a new primary node be automatically selected to establish the connection during the primary/standby switchover, configure the URL as follows:

```
jdbc:opengauss://node1,node2,node3/database?targetServerType=master
```

Log Diagnosis Scenario

If a customer encounters slow data import or some errors that are difficult to analyze, the trace log function can be enabled for diagnosis. The URL can be configured as follows:

```
jdbc:opengauss://node1/database?loggerLevel=trace
```

High Performance

A customer may execute the same SQL statement for multiple times with different input parameters. To improve the execution efficiency, the **prepareThreshold**

parameter can be enabled to avoid repeatedly generating execution plans. The URL can be configured as follows:

```
jdbc:opengauss://node1/database?prepareThreshold=5
```

A customer queries 10 million data records at a time. To prevent memory overflow caused by simultaneous return of the data records, the **defaultRowFetchSize** parameter can be used. The URL can be configured as follows:

```
jdbc:opengauss://node1/database?defaultRowFetchSize=50000
```

A customer needs to insert 10 million data records in batches. To improve efficiency, the **batchMode** parameter can be used. The URL can be configured as follows:

```
jdbc:opengauss://node1/database?batchMode=on
```

Case Conversion

In GaussDB, metadata is stored in lowercase letters by default. If metadata is migrated from a database where metadata is stored in uppercase letters by default to GaussDB, the metadata in uppercase letters changes to lowercase letters. If the original service involves the processing of uppercase metadata, you can enable **uppercaseAttributeName**. However, you are advised to modify the service code instead of using this method to solve the problem. If you have to use this function, ensure that the metadata in the current database is in lowercase to avoid problems.

```
jdbc:opengauss://node1/database?uppercaseAttributeName=true
```

The APIs involved in DatabaseMetaData can be directly called based on input parameters. The methods of using the APIs involved in ResultSetMetaData are as follows:

```
Statement stmt = conn.createStatement();
ResultSet rs = stmt.executeQuery("select * from test_supplier");
ResultSetMetaData rsmd = rs.getMetaData();
for (int i = 1; i <= rsmd.getColumnCount(); i++) {
    System.out.println(rsmd.getColumnLabel(i) + " " + rsmd.getColumnName(i));
}
```

5.3.3.2 Creating and Calling a Stored Procedure

The following illustrates how to develop applications based on GaussDB JDBC APIs. It also demonstrates how to connect to a database and create and call stored procedures.

Prerequisites for code running: Add the **opengaussjdbc.jar** package as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.

```
// There will be security risks if the username and password used for authentication are directly written into code. It is recommended that the username and password be stored in the configuration file or environment variables (the password must be stored in ciphertext and decrypted when being used) to ensure security.
```

```
// In this example, the username and password are stored in environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV in the local environment (set the environment variable names based on the actual situation).
```

```
// You need to change the values of $ip, $port, and database.
```

```
import java.sql.Connection;
import java.sql.DriverManager;
```

```
import java.sql.PreparedStatement;
import java.sql.SQLException;
import java.sql.Statement;
import java.sql.CallableStatement;
import java.sql.Types;

public class DBTest {

    // Establish a database connection in non-encrypted mode.
    public static Connection GetConnection(String username, String passwd) {
        String driver = "com.huawei.opengauss.jdbc.Driver";
        String sourceURL = "jdbc:opengauss://$p:$port/database";
        Connection conn = null;
        try {
            // Load the database driver.
            Class.forName(driver).newInstance();
        } catch (Exception e) {
            e.printStackTrace();
            return null;
        }

        try {
            // Create a database connection.
            conn = DriverManager.getConnection(sourceURL, username, passwd);
            System.out.println("Connection succeed!");
        } catch (Exception e) {
            e.printStackTrace();
            return null;
        }

        return conn;
    };

    // Create a stored procedure.
    public static void CreateCallable(Connection conn) {
        Statement stmt = null;
        try {
            stmt = conn.createStatement();
            // Create a stored procedure and return the sum of the three input values.
            stmt.execute("create or replace procedure testproc \n" +
                "\n" +
                "    psv_in1 in integer,\n" +
                "    psv_in2 in integer,\n" +
                "    psv_inout inout integer\n" +
                ")\n" +
                "as\n" +
                "begin\n" +
                "    psv_inout := psv_in1 + psv_in2 + psv_inout;\n" +
                "end;\n" +
                "/");
        } catch (SQLException e) {
            throw new RuntimeException(e);
        } finally {
            if (stmt != null) {
                try {
                    stmt.close();
                } catch (SQLException e) {
                    throw new RuntimeException(e);
                }
            }
        }
    }

    -- Call the stored procedure.
    public static void ExecCallableSQL(Connection conn) {
        CallableStatement cstmt = null;
        try {
            cstmt=conn.prepareCall("{? = CALL TESTPROC(?,?,?)");
        }
```

```
cstmt.setInt(2, 50);
cstmt.setInt(1, 20);
cstmt.setInt(3, 90);
cstmt.registerOutParameter(4, Types.INTEGER); // Register an OUT parameter of the integer type.
cstmt.execute();
int out = cstmt.getInt(4); // Obtain the OUT parameter.
System.out.println("The CallableStatement TESTPROC returns:"+out);
cstmt.close();
} catch (SQLException e) {
    if (cstmt != null) {
        try {
            cstmt.close();
        } catch (SQLException e1) {
            e1.printStackTrace();
        }
    }
    e.printStackTrace();
}
}
```

```
/**
 * Main process. Call static methods one by one.
 * @param args
 */
public static void main(String[] args) {
    // Create a database connection.
    String userName = System.getenv("EXAMPLE_USERNAME_ENV");
    String password = System.getenv("EXAMPLE_PASSWORD_ENV");
    Connection conn = GetConnection(userName, password);

    // Create a stored procedure.
    CreateCallable(conn);

    // Run the stored procedure.
    ExecCallableSQL(conn);

    // Close the connection to the database.
    try {
        conn.close();
    } catch (SQLException e) {
        e.printStackTrace();
    }
}
```

The execution result of the preceding example is as follows:

```
Connection succeed!
The CallableStatement TESTPROC returns:160
```

5.3.3.3 Obtaining the Return Value of a Function

JDBC obtains the return value when calling a function. The following example shows the return values of the bit and float8 types. For other data types, refer to this example.

Prerequisites for code running: Add the **opengaussjdbc.jar** package as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.

```
// There will be security risks if the username and password used for authentication are directly written into
code. It is recommended that the username and password be stored in the configuration file or
environment variables (the password must be stored in ciphertext and decrypted when being used) to
ensure security.
```

```
// In this example, the username and password are stored in environment variables. Before running this
// example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV in the local
// environment (set the environment variable names based on the actual situation).
// You need to change the values of $ip, $port, and database.

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.Statement;
import java.sql.CallableStatement;
import java.sql.SQLException;
import java.sql.PreparedStatement;
import java.sql.Types;

public class Type {
    public static void main(String[] args) throws SQLException {
        String driver = "com.huawei.opengauss.jdbc.Driver";
        String username = System.getenv("EXAMPLE_USERNAME_ENV");
        String passwd = System.getenv("EXAMPLE_PASSWORD_ENV");
        String sourceURL = "jdbc:opengauss://$ip:$port/database";
        Connection conn = null;

        try {
            // Load the database driver.
            Class.forName(driver).newInstance();
        } catch (Exception e) {
            e.printStackTrace();
        }

        try {
            // Establish a database connection in non-encrypted mode.
            conn = DriverManager.getConnection(sourceURL, username, passwd);
            System.out.println("Connection succeed!");
        } catch (Exception e) {
            e.printStackTrace();
        }

        // Create a table.
        String createsql = "create table if not exists t_bit(col_bit bit)";
        Statement stmt = conn.createStatement();
        stmt.execute(createsql);
        stmt.close();
        // Example of using the bit type. Note that the value range of the bit type is [0,1].
        Statement st = conn.createStatement();
        String sqlstr = "create or replace function fun_1()\n" + "returns bit AS $$\n"
            + "select col_bit from t_bit limit 1;\n" + "$$\n" + "LANGUAGE SQL;";
        st.execute(sqlstr);
        CallableStatement c = conn.prepareCall("{ ? = call fun_1() }");
        // Register the output type, which is a bit string.
        c.registerOutParameter(1, Types.BIT);
        c.execute();
        // Use the Boolean type to obtain the result.
        System.out.println(c.getBoolean(1));

        // Example of using the float8 type.
        st.execute("create table if not exists t_float(col1 float8)");
        PreparedStatement pstmt = conn.prepareStatement("insert into t_float values(?)");
        pstmt.setDouble(1, 123456.123);
        pstmt.execute();
        pstmt.close();

        // Example of using the function whose return value is of the float8 type.
        st.execute(
            "create or replace function func_float() " + "return float8 " + "as declare " + "var1 float8; " + "begin
"
            + " select col1 into var1 from t_float limit 1; " + " return var1; " + "end;");
        CallableStatement cs = conn.prepareCall("{ ? = call func_float() }");
        cs.registerOutParameter(1, Types.DOUBLE);
        cs.execute();
        System.out.println(cs.getDouble(1));
    }
}
```

```
st.close();

// Close the connection to the database.
try {
    conn.close();
} catch (SQLException e) {
    e.printStackTrace();
}
}
```

The execution result of the preceding example is as follows:

```
Connection succeed!
false
123456.123
```

 CAUTION

Currently, JDBC cannot call stored procedures and functions whose return data type is money.

5.3.3.4 Batch Query

In this example, `setFetchSize` adjusts the memory usage of the client by using the database cursor to obtain server data in batches. It may increase network interaction and affect performance to some extent. The cursor is valid within a transaction. Therefore, disable automatic transaction commit and then manually commit the transaction.

Prerequisites for code running: Add the **opengaussjdbc.jar** package as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.

```
// There will be security risks if the username and password used for authentication are directly written into
// code. It is recommended that the username and password be stored in the configuration file or
// environment variables (the password must be stored in ciphertext and decrypted when being used) to
// ensure security.
// In this example, the username and password are stored in environment variables. Before running this
// example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV in the local
// environment (set the environment variable names based on the actual situation).
// You need to change the values of $ip, $port, and database.
```

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.Statement;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.PreparedStatement;

public class Batch {
    public static void main(String[] args) throws SQLException {
        String driver = "com.huawei.opengauss.jdbc.Driver";
        String username = System.getenv("EXAMPLE_USERNAME_ENV");
        String passwd = System.getenv("EXAMPLE_PASSWORD_ENV");
        String sourceURL = "jdbc:opengauss://$ip:$port/database";
        Connection conn = null;

        try {
            // Load the database driver.
            Class.forName(driver).newInstance();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

```
}

try {
    // Establish a database connection in non-encrypted mode.
    conn = DriverManager.getConnection(sourceURL, username, passwd);
    System.out.println("Connection succeed!");
} catch (Exception e) {
    e.printStackTrace();
}

// Disable automatic commit.
conn.setAutoCommit(false);

// Create a table.
Statement st = conn.createStatement();
st.execute("create table mytable (cal1 int);");

// Insert 200 rows of data into the table.
PreparedStatement pstmt = conn.prepareStatement("insert into mytable values (?)");
for (int i = 0; i < 200; i++) {
    pstmt.setInt(1, i + 1);
    pstmt.addBatch();
}
pstmt.executeBatch();
conn.commit();
pstmt.close();

// Open the cursor and obtain 50 rows of data each time.
st.setFetchSize(50);
int fetchCount = 0;
ResultSet rs = st.executeQuery("SELECT * FROM mytable");
while (rs.next()) {
    fetchCount++;
}
System.out.println(fetchCount == 200);
conn.commit();
rs.close();

// Disable the server cursor.
st.setFetchSize(0);
fetchCount = 0;
rs = st.executeQuery("SELECT * FROM mytable");
while (rs.next()) {
    fetchCount++;
}
System.out.println(fetchCount == 200);
conn.commit();
rs.close();

// Close the statement object and database connection.
st.close();
conn.close();

}
}
```

The execution result of the preceding example is as follows:

```
Connection succeed!
true
true
```

Enable automatic commit.

```
conn.setAutoCommit(true);
```

5.3.3.5 SQL Retry at an Application Layer

If the primary database node is faulty and cannot be restored within 10s, GaussDB automatically switches the standby database node to the active state to ensure the normal running of the database. Jobs running during the failover will fail and those started after the failover will not be affected. To prevent upper-layer services from being affected by the failover, refer to the following example to construct an SQL retry mechanism at the service layer.

Prerequisites for code running: Add the **opengaussjdbc.jar** package as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.

```
// There will be security risks if the username and password used for authentication are directly written into
// code. It is recommended that the username and password be stored in the configuration file or
// environment variables (the password must be stored in ciphertext and decrypted when being used) to
// ensure security.
// In this example, the username and password are stored in environment variables. Before running this
// example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV in the local
// environment (set the environment variable names based on the actual situation).
// You need to change the values of $ip, $port, and database.
```

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.PreparedStatement;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;

class ExitHandler extends Thread {
    private Statement cancel_stmt = null;

    public ExitHandler(Statement stmt) {
        super("Exit Handler");
        this.cancel_stmt = stmt;
    }

    public void run() {
        System.out.println("exit handle");
        try {
            this.cancel_stmt.cancel();
        } catch (SQLException e) {
            System.out.println("cancel query failed.");
            e.printStackTrace();
        }
    }
}

public class SQLRetry {
    // Establish a database connection in non-encrypted mode.
    public static Connection GetConnection(String username, String passwd) {
        String driver = "com.huawei.opengauss.jdbc.Driver";
        String sourceURL = "jdbc:opengauss://$ip:$port/database";
        Connection conn = null;
        try {
            // Load the database driver.
            Class.forName(driver).newInstance();
        } catch (Exception e) {
            e.printStackTrace();
            return null;
        }

        try {
            // Create a database connection.
            conn = DriverManager.getConnection(sourceURL, username, passwd);
            System.out.println("Connection succeed!");
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

```
        return null;
    }

    return conn;
}

// Run a common SQL statement to create the jdbc_test1 table.
public static void CreateTable(Connection conn) {
    Statement stmt = null;
    try {
        stmt = conn.createStatement();

        Runtime.getRuntime().addShutdownHook(new ExitHandler(stmt));

        // Run a common SQL statement.
        int rc2 = stmt
            .executeUpdate("DROP TABLE if exists jdbc_test1;");

        int rc1 = stmt
            .executeUpdate("CREATE TABLE jdbc_test1(col1 INTEGER, col2 VARCHAR(10));");

        stmt.close();
    } catch (SQLException e) {
        if (stmt != null) {
            try {
                stmt.close();
            } catch (SQLException e1) {
                e1.printStackTrace();
            }
        }
        e.printStackTrace();
    }
}

// Run a prepared statement to insert data in batches.
public static void BatchInsertData(Connection conn) {
    PreparedStatement pst = null;

    try {
        // Generate a prepared statement.
        pst = conn.prepareStatement("INSERT INTO jdbc_test1 VALUES (?,?)");
        for (int i = 0; i < 100; i++) {
            // Add parameters.
            pst.setInt(1, i);
            pst.setString(2, "data " + i);
            pst.addBatch();
        }
        // Perform batch processing.
        pst.executeBatch();
        pst.close();
    } catch (SQLException e) {
        if (pst != null) {
            try {
                pst.close();
            } catch (SQLException e1) {
                e1.printStackTrace();
            }
        }
        e.printStackTrace();
    }
}

// Run a prepared statement to update data.
private static boolean QueryRedo(Connection conn){
    PreparedStatement pstmt = null;
    boolean retValue = false;
    try {
        pstmt = conn
```



```
.prepareStatement("SELECT col1 FROM jdbc_test1 WHERE col2 = ?");

pstmt.setString(1, "data 10");
ResultSet rs = pstmt.executeQuery();

while (rs.next()) {
    System.out.println("col1 = " + rs.getString("col1"));
}
rs.close();

pstmt.close();
retValue = true;
} catch (SQLException e) {
    System.out.println("catch..... retValue " + retValue);
    if (pstmt != null) {
        try {
            pstmt.close();
        } catch (SQLException e1) {
            e1.printStackTrace();
        }
    }
    e.printStackTrace();
}

System.out.println("finesh.....");
return retValue;
}

// Configure the number of retry attempts for the retry of a query statement upon a failure.
public static void ExecPreparedSQL(Connection conn) throws InterruptedException {
    int maxRetryTime = 10;
    int time = 0;
    String result = null;
    do {
        time++;
        try {
            System.out.println("time:" + time);
            boolean ret = QueryRedo(conn);
            if (ret == false) {
                System.out.println("retry, time:" + time);
                Thread.sleep(10000);
                QueryRedo(conn);
            }
        } catch (Exception e) {
            e.printStackTrace();
        }
    } while (null == result && time < maxRetryTime);
}

/**
 * Main process. Call static methods one by one.
 * @param args
 * @throws InterruptedException
 */
public static void main(String[] args) throws InterruptedException {
    // Create a database connection.
    String userName = System.getenv("EXAMPLE_USERNAME_ENV");
    String password = System.getenv("EXAMPLE_PASSWORD_ENV");
    Connection conn = GetConnection(userName, password);

    // Create a table.
    CreateTable(conn);

    // Insert data in batches.
    BatchInsertData(conn);

    // Run a prepared statement to update data.
    ExecPreparedSQL(conn);
}
```

```
// Close the connection to the database.  
try {  
    conn.close();  
} catch (SQLException e) {  
    e.printStackTrace();  
}  
  
}  
  
}
```

The execution result of the preceding example is as follows:

```
Connection succeed!  
time:1  
col1 = 10  
finesh.....  
time:2  
col1 = 10  
finesh.....  
time:3  
col1 = 10  
finesh.....  
time:4  
col1 = 10  
finesh.....  
time:5  
col1 = 10  
finesh.....  
time:6  
col1 = 10  
finesh.....  
time:7  
col1 = 10  
finesh.....  
time:8  
col1 = 10  
finesh.....  
time:9  
col1 = 10  
finesh.....  
time:10  
col1 = 10  
finesh.....  
exit handle
```

5.3.3.6 Importing and Exporting Data Through Local Files

When Java is used for secondary development based on GaussDB, you can use the CopyManager API to export data from the database to a local file or import a local file to the database by streaming. The file can be in CSV or TEXT format.

Prerequisites for code running:

1. The **opengaussjdbc.jar** package is added as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.
2. Tables **migration_table** and **migration_table_1** have been created in the database, and data has been inserted into the **migration_table** table.

```
// There will be security risks if the username and password used for authentication are directly written into  
// code. It is recommended that the username and password be stored in the configuration file or  
// environment variables (the password must be stored in ciphertext and decrypted when being used) to  
// ensure security.  
// In this example, the username and password are stored in environment variables. Before running this
```

example, set environment variables *EXAMPLE_USERNAME_ENV* and *EXAMPLE_PASSWORD_ENV* in the local environment (set the environment variable names based on the actual situation).
// You need to change the values of *\$ip*, *\$port*, and *database*.

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.io.IOException;
import java.io.FileInputStream;
import java.io.FileOutputStream;
import java.sql.SQLException;
import com.huawei.opengauss.jdbc.copy.CopyManager;
import com.huawei.opengauss.jdbc.core.BaseConnection;

public class Copy{

    public static void main(String[] args)
    {
        String urls = new String("jdbc:opengauss://$ip:$port/database"); // Database URL
        String username = System.getenv("EXAMPLE_USERNAME_ENV"); // Username
        String password = System.getenv("EXAMPLE_PASSWORD_ENV"); // Password
        String tablename = new String("migration_table"); // Table information
        String tablename1 = new String("migration_table_1"); // Table information
        String driver = "com.huawei.opengauss.jdbc.Driver";
        Connection conn = null;

        try {
            Class.forName(driver);
            conn = DriverManager.getConnection(urls, username, password);
        } catch (ClassNotFoundException e) {
            e.printStackTrace(System.out);
        } catch (SQLException e) {
            e.printStackTrace(System.out);
        }

        // Export the query result of SELECT * FROM migration_table to the local file d:/data.txt.
        try {
            copyToFile(conn, "d:/data.txt", "(SELECT * FROM migration_table)");
        } catch (SQLException e) {

        }

        e.printStackTrace();
    } catch (IOException e) {

    }

    e.printStackTrace();
}

// Import data from the d:/data.txt file to the migration_table_1 table.
try {
    copyFromFile(conn, "d:/data.txt", tablename1);
} catch (SQLException e) {
    e.printStackTrace();
} catch (IOException e) {

}

e.printStackTrace();
}

// Export the data from the migration_table_1 table to the d:/data1.txt file.
try {
    copyToFile(conn, "d:/data1.txt", tablename1);
} catch (SQLException e) {

}

e.printStackTrace();
} catch (IOException e) {

}

e.printStackTrace();
}

// Use copyIn to import data from a file to the database.
public static void copyFromFile(Connection connection, String filePath, String tableName)
    throws SQLException, IOException {
```

```
FileInputStream fileInputStream = null;

try {
    CopyManager copyManager = new CopyManager((BaseConnection)connection);
    fileInputStream = new FileInputStream(filePath);
    copyManager.copyIn("COPY " + tableName + " FROM STDIN", fileInputStream);
} finally {
    if (fileInputStream != null) {
        try {
            fileInputStream.close();
        } catch (IOException e) {
            e.printStackTrace();
        }
    }
}

// Use copyOut to export data from the database to a file.
public static void copyToFile(Connection connection, String filePath, String tableOrQuery)
    throws SQLException, IOException {

    FileOutputStream fileOutputStream = null;

    try {
        CopyManager copyManager = new CopyManager((BaseConnection)connection);
        fileOutputStream = new FileOutputStream(filePath);
        copyManager.copyOut("COPY " + tableOrQuery + " TO STDOUT", fileOutputStream);
    } finally {
        if (fileOutputStream != null) {
            try {
                fileOutputStream.close();
            } catch (IOException e) {
                e.printStackTrace();
            }
        }
    }
}
```

Execution result of the preceding example: The data in the **data.txt** and **data1.txt** files in local drive D, and in the database tables **migration_table_1** and **migration_table** is the same.

5.3.3.7 Migrating Data from a MySQL Database

The following example shows how to use CopyManager to migrate data from MY to GaussDB.

Prerequisites for code running:

1. The **opengaussjdbc.jar** package is added as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.
2. The **migration_table** table has been created in MySQL and GaussDB databases. Data has been inserted into the **migration_table** table in the MySQL database in advance.

// There will be security risks if the username and password used for authentication are directly written into code. It is recommended that the username and password be stored in the configuration file or environment variables (the password must be stored in ciphertext and decrypted when being used) to ensure security.

// In this example, the username and password are stored in environment variables. Before running this example, set environment variables *EXAMPLE_USERNAME_ENV* and *EXAMPLE_PASSWORD_ENV* in the local environment (set the environment variable names based on the actual situation).

// You need to change the values of *\$ip*, *\$port*, and *database*.

```
import java.io.StringReader;
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;

import com.huawei.opengauss.jdbc.copy.CopyManager;
import com.huawei.opengauss.jdbc.core.BaseConnection;

public class Migration{

    public static void main(String[] args) {
        String url = new String("jdbc.opengauss://$ip:$port/database"); // Database URL
        String user = System.getenv("EXAMPLE_USERNAME_ENV");           // GaussDB username
        String pass = System.getenv("EXAMPLE_PASSWORD_ENV");           // GaussDB password
        String tablename = new String("migration_table"); // Table information
        String delimiter = new String("|"); // Delimiter
        String encoding = new String("UTF8"); // Character set
        String driver = "com.huawei.opengauss.jdbc.Driver";
        StringBuffer buffer = new StringBuffer(); // Buffer to store formatted data

        try {
            // Obtain the query result set of the source database.
            ResultSet rs = getDataSet();

            // Traverse the result set and obtain records row by row.
            // The values of columns in each record are separated by the specified delimiter and end with a
            //linefeed, forming strings.
            // Add the strings to the buffer.
            while (rs.next()) {
                buffer.append(rs.getString(1) + delimiter
                    + rs.getString(2) + delimiter
                    + rs.getString(3) + delimiter
                    + rs.getString(4)
                    + "\n");
            }
            rs.close();

            try {
                // Connect to the database in non-encrypted mode.
                Class.forName(driver);
                Connection conn = DriverManager.getConnection(url, user, pass);
                BaseConnection baseConn = (BaseConnection) conn;
                baseConn.setAutoCommit(false);

                // Initialize the table.
                String sql = "Copy " + tablename + " from STDIN DELIMITER " + "'" + delimiter + "'" + "
ENCODING " + "'" + encoding + "'";

                // Commit data in the buffer.
                CopyManager cp = new CopyManager(baseConn);
                StringReader reader = new StringReader(buffer.toString());
                cp.copyIn(sql, reader);
                baseConn.commit();
                reader.close();
                baseConn.close();
            } catch (ClassNotFoundException e) {
                e.printStackTrace(System.out);
            } catch (SQLException e) {
                e.printStackTrace(System.out);
            }

            } catch (Exception e) {
                e.printStackTrace();
            }
        }
    }
}
```

```

//*****
// Return the query result set from the source database.
//*****
private static ResultSet getDataSet() {
    ResultSet rs = null;
    try {
        Class.forName("com.mysql.jdbc.Driver").newInstance();
        String userName = System.getenv("EXAMPLE_USERNAME_ENV");
        String password = System.getenv("EXAMPLE_PASSWORD_ENV");
        Connection conn = DriverManager.getConnection("jdbc:mysql://$ip:$port/database?
useSSL=false&allowPublicKeyRetrieval=true", userName, password);
        Statement stmt = conn.createStatement();
        rs = stmt.executeQuery("select * from migration_table");
    } catch (SQLException e) {
        e.printStackTrace();
    } catch (Exception e) {
        e.printStackTrace();
    }
    return rs;
}
}

```

Execution result of the preceding example: After the migration, data in the **migration_table** table of GaussDB is consistent with that of MySQL.

5.3.3.8 Logical Replication

The following example demonstrates how to use the logical replication function through the JDBC APIs.

For logical replication, in addition to the configuration items described in [Logical Decoding](#), the following configuration items are added for JDBC streaming decoding:

1. Decoding thread concurrency

Set **parallel-decode-num** to specify the number of decoder threads for parallel decoding. The value is an integer ranging from 1 to 20. The value **1** indicates that decoding is performed based on the original serial logic. Other values indicate that parallel decoding is enabled. The default value is **1**. When this parameter is set to **1**, do not configure the following options: **decode-style**, **sending-batch**, and **parallel-queue-size**.

2. Decoding format

Configure **decode-style** to specify the decoding format. The value can be 'j', 't' or 'b' of the char type, indicating the JSON, text, or binary format, respectively. This option is set only when parallel decoding is allowed and binary decoding is supported only in the parallel decoding scenario. For the JSON and TEXT formats, in the decoding result sent in batches, the uint32 consisting of the first four bytes of each decoding statement indicates the total number of bytes of the statement (the four bytes occupied by the uint32 are excluded, and **0** indicates that the decoding of this batch ends). The 8-byte uint64 indicates the corresponding LSN (**begin** corresponds to **first_lsn**, **commit** corresponds to **end_lsn**, and other values correspond to the LSN of the statement).

 NOTE

The binary encoding rules are as follows:

1. The first four bytes represent the total number of bytes of the decoding result of statements following the statement-level delimiter letter P (excluded) or the batch end character F (excluded). If the value is **0**, the decoding of this batch ends.
2. The next eight bytes (uint64) indicate the corresponding LSN (**begin** corresponds to **first_lsn**, **commit** corresponds to **end_lsn**, and other values correspond to the LSN of the statement).
3. The following 1-byte letter can be **B**, **C**, **I**, **U**, or **D**, representing BEGIN, COMMIT, INSERT, UPDATE, or DELETE.
4. When the one-byte letter in step 3 is **B**:
 1. The following eight bytes (uint64) indicate the CSN.
 2. The next eight bytes (uint64) indicate **first_lsn**.
 3. (Optional) If the next 1-byte letter is **T**, the following four bytes (uint32) indicate the timestamp length for committing the transaction. The following characters with the same length are the timestamp character string.
 4. (Optional) If the next one-byte letter is **N**, the following four bytes (uint32) indicate the length of the transaction username. The following characters with the same length are the transaction username.
 5. Because there may still be a decoding statement subsequently, a 1-byte letter **P** or **F** is used as a separator between statements. **P** indicates that there are still decoded statements in this batch, and **F** indicates that this batch is completed.
5. When the one-byte letter in step 3 is **C**:
 1. (Optional) If the next 1-byte letter is **X**, the following eight bytes (uint64) indicate XID.
 2. (Optional) If the next 1-byte letter is **T**, the following four bytes (uint32) indicate the timestamp length. The following characters with the same length are the timestamp character string.
 3. When logs are sent in batches, decoding results of other transactions may still exist after a COMMIT log is decoded. If the next 1-byte letter is **P**, the batch still needs to be decoded. If the letter is **F**, the batch decoding ends.
6. When the one-byte letter in step 3 is **I**, **U**, or **D**:
 1. The following two bytes (uint16) indicate the length of the schema name.
 2. The schema name is read based on the preceding length.
 3. The following two bytes (uint16) indicate the length of the table name.
 4. The table name is read based on the preceding length.
 5. (Optional) If the next 1-byte letter is **N**, it indicates a new tuple. If the letter is **O**, it indicates an old tuple. In this case, the new tuple is sent first.
 1. The following two bytes (uint16) indicate the number of columns to be decoded for the tuple, which is recorded as **attrnum**.
 2. The following procedure is repeated for *attrnum* times.
 1. The next two bytes (uint16) indicate the length of the column name.
 2. The column name is read based on the preceding length.
 3. The following four bytes (uint32) indicate the OID of the current column type.
 4. The next four bytes (uint32) indicate the length of the value (stored in the character string format) in the current column. If the value is **0xFFFFFFFF**, it indicates null. If the value is **0**, it indicates a character string whose length is 0.
 5. The column value is read based on the preceding length.

6. Because there may still be a decoding statement after, if the next one-byte letter is **P**, it indicates that the batch still needs to be decoded, and if the next one-byte letter is **F**, it indicates that decoding of the batch ends.
3. Decoding only on the standby node
Configure the **standby-connection** option to specify whether to perform decoding only on the standby node. The value is of the Boolean type (**0** or **1**). The value **true** (or **1**) indicates that only the standby node can be connected for decoding. When the primary node is connected for decoding, an error is reported and the system exits. The value **false** (or **0**) indicates that there is no restriction. The default value is **false (0)**.
4. Batch sending
Configure **sending-batch** to specify whether to send results in batches. The value is an integer ranging from 0 to 1. The value **0** indicates that decoding results are sent one by one. The value **1** indicates that decoding results are sent in batches when the accumulated size of decoding results reaches 1 MB. The default value is **0**. This parameter can be set only during parallel decoding. In the scenario where batch sending is enabled, if the decoding format is 'j' or 't', before each original decoding statement, a uint32 type is added indicating the length of the decoding result (excluding the current uint32 type), and a uint64 type is added, indicating the LSN corresponding to the current decoding result.
5. Length of the parallel decoding queue
Configure **parallel-queue-size** to specify the length of the queue for interaction among parallel logical decoding threads. The value ranges from 2 to 1024 and must be a power of 2. The default value is **128**. The queue length is positively correlated with the memory usage during decoding.
6. Memory threshold for logical decoding
The **max-txn-in-memory** configuration item specifies the memory threshold for caching the intermediate decoding result of a single transaction, in MB. The value ranges from 0 to 100. The default value is **0**, indicating that the memory usage is not controlled. The **max-reorderbuffer-in-memory** configuration item specifies the memory threshold for caching intermediate decoding results of all transactions, in GB. The value ranges from 0 to 100. The default value is **0**, indicating that the memory usage is not controlled. When the memory usage exceeds the threshold, intermediate decoding results are written into a temporary file during decoding, affecting the logical decoding performance.
7. Logical decoding sending timeout threshold
The **sender-timeout** configuration item specifies the heartbeat timeout threshold between the kernel and client. If no message is received from the client within the period, the logical decoding stops and disconnects from the client. The unit is ms, and the value range is [0,2147483647]. The default value depends on the value of **logical_sender_timeout**.
8. User blacklist options for logical decoding
Use the user blacklist for logical decoding. The transaction operations of blacklisted users are filtered from the logical decoding result. The options are as follows:
 - a. **exclude-userids**: specifies the OIDs of blacklisted users. Multiple OIDs are separated by commas (,). The system does not check whether the user OIDs exist.

- b. **exclude-users**: specifies blacklisted usernames. Multiple usernames are separated by commas (,). **dynamic-resolution** specifies whether to dynamically parse and identify usernames. If the decoding is interrupted because the user does not exist and the corresponding blacklisted user does not exist at the time when logs are generated, you can set **dynamic-resolution** to **true** or delete the username from the blacklist to start decoding and continue to obtain logical logs.
 - c. **dynamic-resolution**: specifies whether to dynamically parse blacklisted usernames. The default value is **true**. If the parameter is set to **false**, an error is reported and the logical decoding exits when the decoding detects that the user does not exist in blacklist **exclude-users**. If the parameter is set to **true**, decoding continues when it detects that the user does not exist in blacklist **exclude-users**.
9. Output options for transaction logic logs
- a. **include-xids**: specifies whether the BEGIN logical log of a transaction outputs the transaction ID. The default value is **true**.
 - b. **include-timestamp**: specifies whether the BEGIN logical log of a transaction outputs the time when the transaction is committed. The default value is **false**.
 - c. **include-user**: specifies whether the BEGIN logical log of a transaction outputs the username of the transaction. The default value is **false**. The username of a transaction refers to the authorized user, that is, the login user who executes the session corresponding to the transaction. The username does not change during the execution of the transaction.
10. By default, **socketTimeout** of the logical decoding connection is set to **10s**. When the primary node is overloaded during decoding on the standby node, the connection may be closed due to timeout. You can set **withStatusInterval(10000,TimeUnit.MILLISECONDS)** to adjust the timeout interval.

The decoding performance (Xlog consumption) is greater than or equal to 100 Mbps in the following standard parallel decoding scenario: 16-core CPU, 128 GB memory, network bandwidth > 200 Mbps, 10 to 100 columns in a table, 0.1 KB to 1 KB data in a single row, INSERT as main DML operations, less than 4096 statements in a single transaction, **parallel-decode-num** set to **8**, decoding format as **'b'**, and batch sending function enabled. To ensure that the decoding performance meets the requirements and minimize the impact on services, you are advised to set up only one parallel decoding connection on a standby node to ensure that the CPU, memory, and bandwidth resources are sufficient.

 CAUTION

The logical replication class PGReplicationStream is a non-thread-safe class. Concurrent calls may cause data exceptions.

Prerequisites for code running:

1. The **opengaussjdbc.jar** package is added as required. For example, if you use an IDE to run code, you need to add the **opengaussjdbc.jar** package to the local IDE.

2. Add the IP address (for example, 10.11.12.34) of the JDBC user to the whitelist of the data replication permission. The command is as follows:
`gs_guc reload -Z datanode -N all -I all -h 'host replication all 10.11.12.34/32 sha256'`
3. Set **wal_level** to **logical**. For details, contact the administrator.
4. Create tables **t1** and **t2** and perform DDL or DML operations on the tables.

// There will be security risks if the username and password used for authentication are directly written into code. It is recommended that the username and password be stored in the configuration file or environment variables (the password must be stored in ciphertext and decrypted when being used) to ensure security.

// In this example, the username and password are stored in environment variables. Before running this example, set environment variables *EXAMPLE_USERNAME_ENV* and *EXAMPLE_PASSWORD_ENV* in the local environment (set the environment variable names based on the actual situation).

// You need to change the values of *\$ip*, *\$port*, and *database*.

```
import com.huawei.opengauss.jdbc.PGProperty;
import com.huawei.opengauss.jdbc.jdbc.PgConnection;
import com.huawei.opengauss.jdbc.replication.LogSequenceNumber;
import com.huawei.opengauss.jdbc.replication.PGReplicationStream;

import java.nio.ByteBuffer;
import java.sql.DriverManager;
import java.util.Properties;
import java.util.concurrent.TimeUnit;

public class LogicalReplicationDemo {
    private static PgConnection conn = null;
    public static void main(String[] args) {
        String driver = "com.huawei.opengauss.jdbc.Driver";
        // Configure the IP address and haPort number of the database. By default, the port number is the port
        // number of the connected DN plus 1.
        String sourceURL = "jdbc:opengauss://$ip:$port/database";
        // The default name of the logical replication slot is replication_slot.
        // Test mode: Create a logical replication slot.
        int TEST_MODE_CREATE_SLOT = 1;
        // Test mode: Enable logical replication (the prerequisite is that the logical replication slot already exists).
        int TEST_MODE_START_REPL = 2;
        // Test mode: Delete a logical replication slot.
        int TEST_MODE_DROP_SLOT = 3;
        // Enable different test modes. In practice, set testMode to TEST_MODE_CREATE_SLOT to create a
        // replication slot.
        // Set testMode to TEST_MODE_START_REPL to start decoding. After decoding is complete, set
        // testMode to TEST_MODE_DROP_SLOT to delete the current replication slot.
        int testMode = TEST_MODE_START_REPL;

        try {
            Class.forName(driver);
        } catch (Exception e) {
            e.printStackTrace();
            return;
        }

        try {
            Properties properties = new Properties();
            PGProperty.USER.set(properties, System.getenv("EXAMPLE_USERNAME_ENV")); // Specify the
            decoding username.
            PGProperty.PASSWORD.set(properties, System.getenv("EXAMPLE_PASSWORD_ENV")); // Specify the
            corresponding password.
            // For logical replication, the following three attributes are required. Configure them as follows:
            PGProperty.ASSUME_MIN_SERVER_VERSION.set(properties, "9.4"); // Specify the database version.
            PGProperty.REPLICATION.set(properties, "database"); // Specify the connection used for logical
            replication.
            PGProperty.PREFER_QUERY_MODE.set(properties, "simple"); // Specify the protocol mode.
            conn = (PgConnection) DriverManager.getConnection(sourceURL, properties);
            System.out.println("connection success!");

            if(testMode == TEST_MODE_CREATE_SLOT){
                conn.getReplicationAPI()
```

```
.createReplicationSlot()
.logical()
.withSlotName("replication_slot") // Name of the replication slot to be created. If the
character string contains uppercase letters, the uppercase letters are automatically converted to lowercase
letters.
.withOutputPlugin("mppdb_decoding") // Use the mppdb_decoding plug-in.
.make();
} else if (testMode == TEST_MODE_START_REPL) {
    // Create a replication slot before enabling this mode.
    // LogSequenceNumber waitLSN = LogSequenceNumber.valueOf("0/2808340"); // Filter point for
decoding. Transactions committed before this LSN will not be output.
    PGReplicationStream stream = conn
        .getReplicationAPI()
        .replicationStream()
        .logical()
        .withSlotName("replication_slot") // Specify the name of the replication slot for decoding.
        .withSlotOption("include-xids", true) // Determine whether the decoding result contains
the transaction ID.
        .withSlotOption("skip-empty-xacts", true) // Determine whether to skip empty transactions.
        // .withStartPosition(waitLSN) // Set the filter point for decoding. Transactions committed
before this LSN will not be output.
        .withSlotOption("parallel-decode-num", 10) // Decoding thread concurrency
        // .withSlotOption("white-table-list", "public.t1,public.t2") // Whitelist
        // .withSlotOption("standby-connection", true) // Forcible decoding on standby
        .withSlotOption("decode-style", "t") // Decoding format
        // .withSlotOption("sending-batch", 1) // Send decoding results in batches.
        .withSlotOption("max-txn-in-memory", 100) // The memory threshold for flushing a single
decoding transaction to disks is 100 MB.
        .withSlotOption("max-reorderbuffer-in-memory", 50) // The total memory threshold for
flushing decoding transactions that are being handled to disks is 50 GB.
        .withSlotOption("exclude-users", "userA") // The logical log of the transaction executed by
user A is not returned.
        .withSlotOption("include-user", true) // The BEGIN logical log of the transaction contains
the username.
    .start();
    while (true) {
        ByteBuffer byteBuffer = stream.readPending();

        if (byteBuffer == null) {
            TimeUnit.MILLISECONDS.sleep(10L);
            continue;
        }

        int offset = byteBuffer.arrayOffset();
        byte[] source = byteBuffer.array();
        int length = source.length - offset;
        System.out.println(new String(source, offset, length));

        // If the LSN needs to be flushed, call the following APIs based on the service requirements:
        // LogSequenceNumber lastRecv = stream.getLastReceiveLSN();
        // stream.setFlushedLSN(lastRecv);
        // stream.forceUpdateStatus();

    }
} else if (testMode == TEST_MODE_DROP_SLOT) {
    conn.getReplicationAPI()
        .dropReplicationSlot("replication_slot");
}
} catch (Exception e) {
    e.printStackTrace();
    return;
}
}
```

The following is an example of the decoding result in the text format (that is, the 't' format):

```
BEGIN CSN: 2014 first_lsn: 0/2816A28 username: xxx
table public t1 INSERT: a[integer]:1 b[integer]:2 c[text]:'hello'
```

```
commit xid: 15504
BEGIN CSN: 2015 first_lsn: 0/2816C20 username: xxx
table public t1 UPDATE: old-key: a[integer]:1 new-tuple: a[integer]:1 b[integer]:5 c[text]:'hello'
commit xid: 15505
BEGIN CSN: 2016 first_lsn: 0/2816D60 username: xxx
table public t1 DELETE: a[integer]:1
commit xid: 15506
```

5.3.4 JDBC API Reference

This section describes common JDBC APIs. For more APIs, check JDK 1.8 (software package) and JDBC 4.0.

5.3.4.1 java.sql.Connection

java.sql.Connection is an API for connecting to a database.

Table 5-11 Support status for java.sql.Connection

Method	Return Type	JDBC 4 Supported or Not
abort(Executor executor)	void	Yes
clearWarnings()	void	Yes
close()	void	Yes
commit()	void	Yes
createArrayOf(String typeName, Object[] elements)	Array	Yes
createBlob()	Blob	Yes
createClob()	Clob	Yes
createSQLXML()	SQLXML	Yes
createStatement()	Statement	Yes
createStatement(int resultSetType, int resultSetConcurrency)	Statement	Yes
createStatement(int resultSetType, int resultSetConcurrency, int resultSetHoldability)	Statement	Yes
getAutoCommit()	boolean	Yes
getCatalog()	String	Yes
getClientInfo()	Properties	Yes
getClientInfo(String name)	String	Yes

Method	Return Type	JDBC 4 Supported or Not
getHoldability()	int	Yes
getMetaData()	DatabaseMetaData	Yes
getNetworkTimeout()	int	Yes
getSchema()	String	Yes
getTransactionIsolation()	int	Yes
getTypeMap()	Map<String,Class<?>>	Yes
getWarnings()	SQLWarning	Yes
isClosed()	boolean	Yes
isReadOnly()	boolean	Yes
isValid(int timeout)	Boolean	Yes
nativeSQL(String sql)	String	Yes
prepareCall(String sql)	CallableStatement	Yes
prepareCall(String sql, int resultSetType, int resultSetConcurrency)	CallableStatement	Yes
prepareCall(String sql, int resultSetType, int resultSetConcurrency, int resultSetHoldability)	CallableStatement	Yes
prepareStatement(String sql)	PreparedStatement	Yes
prepareStatement(String sql, int autoGeneratedKeys)	PreparedStatement	Yes
prepareStatement(String sql, int[] columnIndexes)	PreparedStatement	Yes
prepareStatement(String sql, int resultSetType, int resultSetConcurrency)	PreparedStatement	Yes
prepareStatement(String sql, int resultSetType, int resultSetConcurrency, int resultSetHoldability)	PreparedStatement	Yes
prepareStatement(String sql, String[] columnNames)	PreparedStatement	Yes
releaseSavepoint(Savepoint savepoint)	void	Yes

Method	Return Type	JDBC 4 Supported or Not
rollback()	void	Yes
rollback(Savepoint savepoint)	void	Yes
setAutoCommit(boolean autoCommit)	void	Yes
setClientInfo(Properties properties)	void	Yes
setClientInfo(String name,String value)	void	Yes
setHoldability(int holdability)	void	Yes
setNetworkTimeout(Executor executor, int milliseconds)	void	Yes
setReadOnly(boolean readOnly)	void	Yes
setSavepoint()	Savepoint	Yes
setSavepoint(String name)	Savepoint	Yes
setSchema(String schema)	void	Yes
setTransactionIsolation(int level)	void	Yes
setTypeMap(Map<String,Class<?>> map)	void	Yes

NOTICE

The AutoCommit mode is used by default within the API. If you disable it by running **setAutoCommit(false)**, all the statements executed later will be packaged in explicit transactions, and you cannot execute statements that cannot be executed within transactions.

5.3.4.2 java.sql.CallableStatement

This section describes java.sql.CallableStatement, the API for executing the stored procedure.

Table 5-12 Support status for java.sql.CallableStatement

Method Name	Return Type	JDBC 4 Is Supported Or Not
getArray(int parameterIndex)	Array	Yes
getBigDecimal(int parameterIndex)	BigDecimal	Yes
getBlob(int parameterIndex)	Blob	Yes
getBoolean(int parameterIndex)	Boolean	Yes
getByte(int parameterIndex)	byte	Yes
getBytes(int parameterIndex)	byte[]	Yes
getClob(int parameterIndex)	Clob	Yes
getDate(int parameterIndex)	Date	Yes
getDate(int parameterIndex, Calendar cal)	Date	Yes
getDouble(int parameterIndex)	double	Yes
getFloat(int parameterIndex)	float	Yes
getInt(int parameterIndex)	int	Yes
getLong(int parameterIndex)	long	Yes
getObject(int parameterIndex)	Object	Yes
getObject(int parameterIndex, Class<T> type)	Object	Yes
getShort(int parameterIndex)	short	Yes
getSQLXML(int parameterIndex)	SQLXML	Yes
getString(int parameterIndex)	String	Yes
getNString(int parameterIndex)	String	Yes
getTime(int parameterIndex)	Time	Yes
getTime(int parameterIndex, Calendar cal)	Time	Yes
getTimestamp(int parameterIndex)	Timestamp	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getTimestamp(int parameterIndex, Calendar cal)	Timestamp	Yes
registerOutParameter(int parameterIndex, int type)	void	Yes
registerOutParameter(int parameterIndex, int sqlType, int type)	void	Yes
wasNull()	boolean	Yes

 NOTE

- The batch operation of statements containing OUT parameter is not allowed.
- The following methods are inherited from **java.sql.Statement**: **close**, **execute**, **executeQuery**, **executeUpdate**, **getConnection**, **getResultSet**, **getUpdateCount**, **isClosed**, **setMaxRows**, and **setFetchSize**.
- The following methods are inherited from **java.sql.PreparedStatement**: **addBatch**, **clearParameters**, **execute**, **executeQuery**, **executeUpdate**, **getMetaData**, **setBigDecimal**, **setBoolean**, **setByte**, **setBytes**, **setDate**, **setDouble**, **setFloat**, **setInt**, **setLong**, **setNull**, **setObject**, **setString**, **setTime**, and **setTimestamp**.
- The **registerOutParameter(int parameterIndex, int sqlType, int type)** method is used only to register the composite data type.

5.3.4.3 java.sql.DatabaseMetaData

This section describes **java.sql.DatabaseMetaData**, the API for defining database objects.

Table 5-13 Support status for **java.sql.DatabaseMetaData**

Method Name	Return Type	JDBC 4 Is Supported Or Not
allProceduresAreCallable()	Boolean	Yes
allTablesAreSelectable()	Boolean	Yes
autoCommitFailureClosesAllResultSets()	Boolean	Yes
dataDefinitionCausesTransactionCommit()	Boolean	Yes
dataDefinitionIgnoredInTransactions()	Boolean	Yes
deletesAreDetected(int type)	Boolean	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
doesMaxRowSizeIncludeBlobs()	Boolean	Yes
generatedKeyAlwaysReturned()	Boolean	Yes
getBestRowIdentifier(String catalog, String schema, String table, int scope, boolean nullable)	ResultSet	Yes
getCatalogs()	ResultSet	Yes
getCatalogSeparator()	String	Yes
getCatalogTerm()	String	Yes
getClientInfoProperties()	ResultSet	Yes
getColumnPrivileges(String catalog, String schema, String table, String columnNamePattern)	ResultSet	Yes
getConnection()	Connection	Yes
getCrossReference(String parentCatalog, String parentSchema, String parentTable, String foreignCatalog, String foreignSchema, String foreignTable)	ResultSet	Yes
getDefaultTransactionIsolation()	int	Yes
getExportedKeys(String catalog, String schema, String table)	ResultSet	Yes
getExtraNameCharacters()	String	Yes
getFunctionColumns(String catalog, String schemaPattern, String functionNamePattern, String columnNamePattern)	ResultSet	Yes
getFunctions(String catalog, String schemaPattern, String functionNamePattern)	ResultSet	Yes
getIdentifierQuoteString()	String	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getImportedKeys(String catalog, String schema, String table)	ResultSet	Yes
getIndexInfo(String catalog, String schema, String table, boolean unique, boolean approximate)	ResultSet	Yes
getMaxBinaryLiteralLength()	int	Yes
getMaxCatalogNameLength()	int	Yes
getMaxCharLiteralLength()	int	Yes
getMaxColumnNameLength()	int	Yes
getMaxColumnsInGroupBy()	int	Yes
getMaxColumnsInIndex()	int	Yes
getMaxColumnsInOrderBy()	int	Yes
getMaxColumnsInSelect()	int	Yes
getMaxColumnsInTable()	int	Yes
getMaxConnections()	int	Yes
getMaxCursorNameLength()	int	Yes
getMaxIndexLength()	int	Yes
getMaxLogicalLobSize()	default long	Yes
getMaxProcedureNameLength()	int	Yes
getMaxRowSize()	int	Yes
getMaxSchemaNameLength()	int	Yes
getMaxStatementLength()	int	Yes
getMaxStatements()	int	Yes
getMaxTableNameLength()	int	Yes
getMaxTablesInSelect()	int	Yes
getMaxUserNameLength()	int	Yes
getNumericFunctions()	String	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getPrimaryKeys(String catalog, String schema, String table)	ResultSet	Yes
getPartitionTablePrimaryKeys(String catalog, String schema, String table)	ResultSet	Yes
getProcedureColumns(String catalog, String schemaPattern, String procedureNamePattern, String columnNamePattern)	ResultSet	Yes
getProcedures(String catalog, String schemaPattern, String procedureNamePattern)	ResultSet	Yes
getProcedureTerm()	String	Yes
getSchemas()	ResultSet	Yes
getSchemas(String catalog, String schemaPattern)	ResultSet	Yes
getSchemaTerm()	String	Yes
getSearchStringEscape()	String	Yes
getSQLKeywords()	String	Yes
getSQLStateType()	int	Yes
getStringFunctions()	String	Yes
getSystemFunctions()	String	Yes
getTablePrivileges(String catalog, String schemaPattern, String tableNamePattern)	ResultSet	Yes
getTimeDateFunctions()	String	Yes
getTypeInfo()	ResultSet	Yes
getUDTs(String catalog, String schemaPattern, String typeNamePattern, int[] types)	ResultSet	Yes
getURL()	String	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getVersionColumns(String catalog, String schema, String table)	ResultSet	Yes
insertsAreDetected(int type)	Boolean	Yes
locatorsUpdateCopy()	Boolean	Yes
othersDeletesAreVisible(int type)	Boolean	Yes
othersInsertsAreVisible(int type)	Boolean	Yes
othersUpdatesAreVisible(int type)	Boolean	Yes
ownDeletesAreVisible(int type)	Boolean	Yes
ownInsertsAreVisible(int type)	Boolean	Yes
ownUpdatesAreVisible(int type)	Boolean	Yes
storesLowerCaseIdentifiers()	Boolean	Yes
storesMixedCaseIdentifiers()	Boolean	Yes
storesUpperCaseIdentifiers()	Boolean	Yes
supportsBatchUpdates()	Boolean	Yes
supportsCatalogsInDataManipulation()	Boolean	Yes
supportsCatalogsInIndexDefinitions()	Boolean	Yes
supportsCatalogsInPrivilegeDefinitions()	Boolean	Yes
supportsCatalogsInProcedureCalls()	Boolean	Yes
supportsCatalogsInTableDefinitions()	Boolean	Yes
supportsCorrelatedSubqueries()	Boolean	Yes
supportsDataDefinitionAndDataManipulationTransactions()	Boolean	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
supportsDataManipulation-TransactionsOnly()	Boolean	Yes
supportsGetGeneratedKeys()	Boolean	Yes
supportsMixedCaseIdentifiers()	Boolean	Yes
supportsMultipleOpenResults()	Boolean	Yes
supportsNamedParameters()	Boolean	Yes
supportsOpenCursorsAcrossCommit()	Boolean	Yes
supportsOpenCursorsAcrossRollback()	Boolean	Yes
supportsOpenStatementsAcrossCommit()	Boolean	Yes
supportsOpenStatementsAcrossRollback()	Boolean	Yes
supportsPositionedDelete()	Boolean	Yes
supportsPositionedUpdate()	Boolean	Yes
supportsRefCursors()	Boolean	Yes
supportsResultSetConcurrency(int type, int concurrency)	Boolean	Yes
supportsResultSetType(int type)	Boolean	Yes
supportsSchemasInIndexDefinitions()	Boolean	Yes
supportsSchemasInPrivilegeDefinitions()	Boolean	Yes
supportsSchemasInProcedureCalls()	Boolean	Yes
supportsSchemasInTableDefinitions()	Boolean	Yes
supportsSelectForUpdate()	Boolean	Yes
supportsStatementPooling()	Boolean	Yes
supportsStoredFunctionsUsingCallSyntax()	Boolean	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
supportsStoredProcedures()	Boolean	Yes
supportsSubqueriesInComparisons()	Boolean	Yes
supportsSubqueriesInExists()	Boolean	Yes
supportsSubqueriesInIns()	Boolean	Yes
supportsSubqueriesInQuantifieds()	Boolean	Yes
supportsTransactionIsolationLevel(int level)	Boolean	Yes
supportsTransactions()	Boolean	Yes
supportsUnion()	Boolean	Yes
supportsUnionAll()	Boolean	Yes
updatesAreDetected(int type)	Boolean	Yes
getTables(String catalog, String schemaPattern, String tableNamePattern, String[] types)	ResultSet	Yes
getColumns(String catalog, String schemaPattern, String tableNamePattern, String columnNamePattern)	ResultSet	Yes
getTableTypes()	ResultSet	Yes
getUserName()	String	Yes
isReadOnly()	Boolean	Yes
nullsAreSortedHigh()	Boolean	Yes
nullsAreSortedLow()	Boolean	Yes
nullsAreSortedAtStart()	Boolean	Yes
nullsAreSortedAtEnd()	Boolean	Yes
getDatabaseProductName()	String	Yes
getDatabaseProductVersion()	String	Yes
getDriverName()	String	Yes
getDriverVersion()	String	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getDriverMajorVersion()	int	Yes
getDriverMinorVersion()	int	Yes
usesLocalFiles()	Boolean	Yes
usesLocalFilePerTable()	Boolean	Yes
supportsMixedCaseIdentifiers()	Boolean	Yes
storesUpperCaseIdentifiers()	Boolean	Yes
storesLowerCaseIdentifiers()	Boolean	Yes
supportsMixedCaseQuotedIdentifiers()	Boolean	Yes
storesUpperCaseQuotedIdentifiers()	Boolean	Yes
storesLowerCaseQuotedIdentifiers()	Boolean	Yes
storesMixedCaseQuotedIdentifiers()	Boolean	Yes
supportsAlterTableWithAddColumn()	Boolean	Yes
supportsAlterTableWithDropColumn()	Boolean	Yes
supportsColumnAliasing()	Boolean	Yes
nullPlusNonNullIsNull()	Boolean	Yes
supportsConvert()	Boolean	Yes
supportsConvert(int fromType, int toType)	Boolean	Yes
supportsTableCorrelationNames()	Boolean	Yes
supportsDifferentTableCorrelationNames()	Boolean	Yes
supportsExpressionsInOrderBy()	Boolean	Yes
supportsOrderByUnrelated()	Boolean	Yes
supportsGroupBy()	Boolean	Yes
supportsGroupByUnrelated()	Boolean	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
supportsGroupByBeyondSelect()	Boolean	Yes
supportsLikeEscapeClause()	Boolean	Yes
supportsMultipleResultSets()	Boolean	Yes
supportsMultipleTransactions()	Boolean	Yes
supportsNonNullableColumns()	Boolean	Yes
supportsMinimumSQLGrammar()	Boolean	Yes
supportsCoreSQLGrammar()	Boolean	Yes
supportsExtendedSQLGrammar()	Boolean	Yes
supportsANSI92EntryLevelSQL()	Boolean	Yes
supportsANSI92IntermediateSQL()	Boolean	Yes
supportsANSI92FullSQL()	Boolean	Yes
supportsIntegrityEnhancementFacility()	Boolean	Yes
supportsOuterJoins()	Boolean	Yes
supportsFullOuterJoins()	Boolean	Yes
supportsLimitedOuterJoins()	Boolean	Yes
isCatalogAtStart()	Boolean	Yes
supportsSchemasInDataManipulation()	Boolean	Yes
supportsSavepoints()	Boolean	Yes
supportsResultSetHoldability(int holdability)	Boolean	Yes
getResultSetHoldability()	int	Yes
getDatabaseMajorVersion()	int	Yes
getDatabaseMinorVersion()	int	Yes
getJDBCMajorVersion()	int	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getJDBCMajorVersion()	int	Yes

 NOTE

If the value of **uppercaseAttributeName** is **true**, the following APIs convert the query result to uppercase letters. The conversion range is the same as that of the **toUpperCase** method in Java.

- `public ResultSet getProcedures(String catalog, String schemaPattern, String procedureNamePattern)`
- `public ResultSet getProcedureColumns(String catalog, String schemaPattern, String procedureNamePattern, String columnNamePattern)`
- `public ResultSet getTables(String catalog, String schemaPattern, String tableNamePattern, String[] types)`
- `public ResultSet getSchemas(String catalog, String schemaPattern)`
- `public ResultSet getColumns(String catalog, String schemaPattern, String tableNamePattern, String columnNamePattern)`
- `public ResultSet getColumnPrivileges(String catalog, String schema, String table, String columnNamePattern)`
- `public ResultSet getTablePrivileges(String catalog, String schemaPattern, String tableNamePattern)`
- `public ResultSet getBestRowIdentifier(String catalog, String schema, String table, int scope, boolean nullable)`
- `public ResultSet getPrimaryKeys(String catalog, String schema, String table)`
- `protected ResultSet getImportedExportedKeys(String primaryCatalog, String primarySchema, String primaryTable, String foreignCatalog, String foreignSchema, String foreignTable)`
- `public ResultSet getIndexInfo(String catalog, String schema, String tableName, boolean unique, boolean approximate)`
- `public ResultSet getUDTs(String catalog, String schemaPattern, String typeNamePattern, int[] types)`
- `public ResultSet getFunctions(String catalog, String schemaPattern, String functionNamePattern)`

 CAUTION

The `getPartitionTablePrimaryKeys(String catalog, String schema, String table)` API is used to obtain the primary key column of a partitioned table that contains global indexes. An example is as follows:

```
PgDatabaseMetaData dbmd = (PgDatabaseMetaData)conn.getMetaData();  
dbmd.getPartitionTablePrimaryKeys("catalogName", "schemaName", "tableName");
```

5.3.4.4 java.sql.Driver

`java.sql.Driver` is a database driver API.

Table 5-14 Support status for java.sql.Driver

Method	Return Type	JDBC 4 Supported or Not
acceptsURL(String url)	boolean	Yes
connect(String url, Properties info)	Connection	Yes
jdbcCompliant()	boolean	Yes
getMajorVersion()	int	Yes
getMinorVersion()	int	Yes
getParentLogger()	Logger	Yes
getPropertyInfo(String url, Properties info)	DriverPropertyInfo[]	Yes

5.3.4.5 java.sql.PreparedStatement

java.sql.PreparedStatement is a prepared statement API.

Table 5-15 Support status for java.sql.PreparedStatement

Method	Return Type	JDBC 4 Supported or Not
clearParameters()	void	Yes
execute()	boolean	Yes
executeQuery()	ResultSet	Yes
excuteUpdate()	int	Yes
executeLargeUpdate()	long	No
getMetaData()	ResultSetMetaData	Yes
getParameterMetaData()	ParameterMetaData	Yes
setArray(int parameterIndex, Array x)	void	Yes
setAsciiStream(int parameterIndex, InputStream x, int length)	void	Yes
setBinaryStream(int parameterIndex, InputStream x)	void	Yes

Method	Return Type	JDBC 4 Supported or Not
setBinaryStream(int parameterIndex, InputStream x, int length)	void	Yes
setBinaryStream(int parameterIndex, InputStream x, long length)	void	Yes
setBlob(int parameterIndex, InputStream inputStream)	void	Yes
setBlob(int parameterIndex, InputStream inputStream, long length)	void	Yes
setBlob(int parameterIndex, Blob x)	void	Yes
setCharacterStream(int parameterIndex, Reader reader)	void	Yes
setCharacterStream(int parameterIndex, Reader reader, int length)	void	Yes
setClob(int parameterIndex, Reader reader)	void	Yes
setClob(int parameterIndex, Reader reader, long length)	void	Yes
setClob(int parameterIndex, Clob x)	void	Yes
setDate(int parameterIndex, Date x, Calendar cal)	void	Yes
setNull(int parameterIndex, int sqlType)	void	Yes

Method	Return Type	JDBC 4 Supported or Not
setNull(int parameterIndex, int sqlType, String typeName)	void	Yes
setObject(int parameterIndex, Object x)	void	Yes
setObject(int parameterIndex, Object x, int targetSqlType)	void	Yes
setObject(int parameterIndex, Object x, int targetSqlType, int scaleOrLength)	void	Yes
setSQLXML(int parameterIndex, SQLXML xmlObject)	void	Yes
setTime(int parameterIndex, Time x)	void	Yes
setTime(int parameterIndex, Time x, Calendar cal)	void	Yes
setTimestamp(int parameterIndex, Timestamp x)	void	Yes
setTimestamp(int parameterIndex, Timestamp x, Calendar cal)	void	Yes
setUnicodeStream(int parameterIndex, InputStream x, int length)	void	Yes
setURL(int parameterIndex, URL x)	void	Yes
setBoolean(int parameterIndex, boolean x)	void	Yes
setBigDecimal(int parameterIndex, BigDecimal x)	void	Yes

Method	Return Type	JDBC 4 Supported or Not
setByte(int parameterIndex, byte x)	void	Yes
setBytes(int parameterIndex, byte[] x)	void	Yes
setDate(int parameterIndex, Date x)	void	Yes
setDouble(int parameterIndex, double x)	void	Yes
setFloat(int parameterIndex, float x)	void	Yes
setInt(int parameterIndex, int x)	void	Yes
setLong(int parameterIndex, long x)	void	Yes
setShort(int parameterIndex, short x)	void	Yes
setString(int parameterIndex, String x)	void	Yes
setNString(int parameterIndex, String x)	void	Yes
addBatch()	void	Yes
executeBatch()	int[]	Yes

 NOTE

- Execute **addBatch()** and **execute()** only after running **clearBatch()**.
- Batch is not cleared by calling **executeBatch()**. Clear batch by explicitly calling **clearBatch()**.
- After bound variables of a batch are added, if you want to reuse these values, you do not need to use **set*()** again. Instead, add a batch.
- The following methods are inherited from **java.sql.Statement**: **close**, **execute**, **executeQuery**, **executeUpdate**, **getConnection**, **getResultSet**, **getUpdateCount**, **isClosed**, **setMaxRows**, and **setFetchSize**.
- The **executeLargeUpdate()** method can only be used in JDBC 4.2 or later.

5.3.4.6 java.sql.ResultSet

This section describes java.sql.ResultSet, the API for execution result sets.

Table 5-16 Support status for java.sql.ResultSet

Method Name	Return Type	JDBC 4 Is Supported Or Not
absolute(int row)	boolean	Yes
afterLast()	void	Yes
beforeFirst()	void	Yes
cancelRowUpdates()	void	Yes
clearWarnings()	void	Yes
close()	void	Yes
deleteRow()	void	Yes
findColumn(String columnLabel)	int	Yes
first()	boolean	Yes
getArray(int columnIndex)	Array	Yes
getArray(String columnLabel)	Array	Yes
getAsciiStream(int columnIndex)	InputStream	Yes
getAsciiStream(String columnLabel)	InputStream	Yes
getBigDecimal(int columnIndex)	BigDecimal	Yes
getBigDecimal(String columnLabel)	BigDecimal	Yes
getBinaryStream(int columnIndex)	InputStream	Yes
getBinaryStream(String columnLabel)	InputStream	Yes
getBlob(int columnIndex)	Blob	Yes
getBlob(String columnLabel)	Blob	Yes
getBoolean(int columnIndex)	boolean	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getBoolean(String columnLabel)	boolean	Yes
getByte(int columnIndex)	byte	Yes
getBytes(int columnIndex)	byte[]	Yes
getByte(String columnLabel)	byte	Yes
getBytes(String columnLabel)	byte[]	Yes
getCharacterStream(int columnIndex)	Reader	Yes
getCharacterStream(String columnLabel)	Reader	Yes
getClob(int columnIndex)	Clob	Yes
getClob(String columnLabel)	Clob	Yes
getConcurrency()	int	Yes
getCursorName()	String	Yes
getDate(int columnIndex)	Date	Yes
getDate(int columnIndex, Calendar cal)	Date	Yes
getDate(String columnLabel)	Date	Yes
getDate(String columnLabel, Calendar cal)	Date	Yes
getDouble(int columnIndex)	double	Yes
getDouble(String columnLabel)	double	Yes
getFetchDirection()	int	Yes
getFetchSize()	int	Yes
getFloat(int columnIndex)	float	Yes
getFloat(String columnLabel)	float	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getInt(int columnIndex)	int	Yes
getInt(String columnLabel)	int	Yes
getLong(int columnIndex)	long	Yes
getLong(String columnLabel)	long	Yes
getMetaData()	ResultSetMetaData	Yes
getObject(int columnIndex)	Object	Yes
getObject(int columnIndex, Class<T> type)	<T> T	Yes
getObject(int columnIndex, Map<String,Class<?>> map)	Object	Yes
getObject(String columnLabel)	Object	Yes
getObject(String columnLabel, Class<T> type)	<T> T	Yes
getObject(String columnLabel, Map<String,Class<?>> map)	Object	Yes
getRow()	int	Yes
getShort(int columnIndex)	short	Yes
getShort(String columnLabel)	short	Yes
getSQLXML(int columnIndex)	SQLXML	Yes
getSQLXML(String columnLabel)	SQLXML	Yes
getStatement()	Statement	Yes
getString(int columnIndex)	String	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
getString(String columnLabel)	String	Yes
getString(int columnIndex)	String	Yes
getString(String columnLabel)	String	Yes
getTime(int columnIndex)	Time	Yes
getTime(int columnIndex, Calendar cal)	Time	Yes
getTime(String columnLabel)	Time	Yes
getTime(String columnLabel, Calendar cal)	Time	Yes
getTimestamp(int columnIndex)	Timestamp	Yes
getTimestamp(int columnIndex, Calendar cal)	Timestamp	Yes
getTimestamp(String columnLabel)	Timestamp	Yes
getTimestamp(String columnLabel, Calendar cal)	Timestamp	Yes
getType()	int	Yes
getWarnings()	SQLWarning	Yes
insertRow()	void	Yes
isAfterLast()	boolean	Yes
isBeforeFirst()	boolean	Yes
isClosed()	boolean	Yes
isFirst()	boolean	Yes
isLast()	boolean	Yes
last()	boolean	Yes
moveToCurrentRow()	void	Yes
moveToInsertRow()	void	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
next()	boolean	Yes
previous()	boolean	Yes
refreshRow()	void	Yes
relative(int rows)	boolean	Yes
rowDeleted()	boolean	Yes
rowInserted()	boolean	Yes
rowUpdated()	boolean	Yes
setFetchDirection(int direction)	void	Yes
setFetchSize(int rows)	void	Yes
updateArray(int columnIndex, Array x)	void	Yes
updateArray(String columnLabel, Array x)	void	Yes
updateAsciiStream(int columnIndex, InputStream x, int length)	void	Yes
updateAsciiStream(String columnLabel, InputStream x, int length)	void	Yes
updateBigDecimal(int columnIndex, BigDecimal x)	void	Yes
updateBigDecimal(String columnLabel, BigDecimal x)	void	Yes
updateBinaryStream(int columnIndex, InputStream x, int length)	void	Yes
updateBinaryStream(String columnLabel, InputStream x, int length)	void	Yes
updateBoolean(int columnIndex, boolean x)	void	Yes
updateBoolean(String columnLabel, boolean x)	void	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
updateByte(int columnIndex, byte x)	void	Yes
updateByte(String columnLabel, byte x)	void	Yes
updateBytes(int columnIndex, byte[] x)	void	Yes
updateBytes(String columnLabel, byte[] x)	void	Yes
updateCharacter-Stream(int columnIndex, Reader x, int length)	void	Yes
updateCharacter-Stream(String columnLabel, Reader reader, int length)	void	Yes
updateDate(int columnIndex, Date x)	void	Yes
updateDate(String columnLabel, Date x)	void	Yes
updateDouble(int columnIndex, double x)	void	Yes
updateDouble(String columnLabel, double x)	void	Yes
updateFloat(int columnIndex, float x)	void	Yes
updateFloat(String columnLabel, float x)	void	Yes
updateInt(int columnIndex, int x)	void	Yes
updateInt(String columnLabel, int x)	void	Yes
updateLong(int columnIndex, long x)	void	Yes
updateLong(String columnLabel, long x)	void	Yes
updateNull(int columnIndex)	void	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
updateNull(String columnLabel)	void	Yes
updateObject(int columnIndex, Object x)	void	Yes
updateObject(int columnIndex, Object x, int scaleOrLength)	void	Yes
updateObject(String columnLabel, Object x)	void	Yes
updateObject(String columnLabel, Object x, int scaleOrLength)	void	Yes
updateRow()	void	Yes
updateShort(int columnIndex, short x)	void	Yes
updateShort(String columnLabel, short x)	void	Yes
updateSQLXML(int columnIndex, SQLXML xmlObject)	void	Yes
updateSQLXML(String columnLabel, SQLXML xmlObject)	void	Yes
updateString(int columnIndex, String x)	void	Yes
updateString(String columnLabel, String x)	void	Yes
updateTime(int columnIndex, Time x)	void	Yes
updateTime(String columnLabel, Time x)	void	Yes
updateTimestamp(int columnIndex, Timestamp x)	void	Yes
updateTimestamp(String columnLabel, Timestamp x)	void	Yes
wasNull()	boolean	Yes

 NOTE

- One statement cannot have multiple open ResultSets.
- The cursor that is used for traversing the ResultSet cannot be open after being committed.

5.3.4.7 java.sql.ResultSetMetaData

java.sql.ResultSetMetaData describes the information about ResultSet objects.

Table 5-17 Support status for java.sql.ResultSetMetaData

Method	Return Type	JDBC 4 Is Supported Or Not
getCatalogName(int column)	String	Yes
getColumnClassName(int column)	String	Yes
getColumnCount()	int	Yes
getColumnDisplaySize(int column)	int	Yes
getColumnLabel(int column)	String	Yes
getColumnName(int column)	String	Yes
getColumnType(int column)	int	Yes
getColumnTypeName(int column)	String	Yes
getPrecision(int column)	int	Yes
getScale(int column)	int	Yes
getSchemaName(int column)	String	Yes
getTableName(int column)	String	Yes
isAutoIncrement(int column)	Boolean	Yes
isCaseSensitive(int column)	Boolean	Yes
isCurrency(int column)	Boolean	Yes
isDefinitelyWritable(int column)	Boolean	Yes
isNullable(int column)	int	Yes

Method	Return Type	JDBC 4 Is Supported Or Not
isReadOnly(int column)	Boolean	Yes
isSearchable(int column)	Boolean	Yes
isSigned(int column)	Boolean	Yes
isWritable(int column)	Boolean	Yes

 NOTE

When **uppercaseAttributeName** is set to **true**, the following APIs convert the query result to uppercase letters. The conversion range is 26 English letters.

- public String getColumnName(int column)
- public String getColumnLabel(int column)

5.3.4.8 java.sql.Statement

This section describes java.sql.Statement, the API for executing SQL statements.

Table 5-18 Support status for java.sql.Statement

Method Name	Return Type	JDBC 4 Is Supported Or Not
addBatch(String sql)	void	Yes
clearBatch()	void	Yes
clearWarnings()	void	Yes
close()	void	Yes
closeOnCompletion()	void	Yes
execute(String sql)	boolean	Yes
execute(String sql, int autoGeneratedKeys)	boolean	Yes
execute(String sql, int[] columnIndexes)	boolean	Yes
execute(String sql, String[] columnNames)	boolean	Yes
executeBatch()	boolean	Yes
executeQuery(String sql)	ResultSet	Yes
executeUpdate(String sql)	int	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
executeUpdate(String sql, int autoGeneratedKeys)	int	Yes
executeUpdate(String sql, int[] columnIndexes)	int	Yes
executeUpdate(String sql, String[] columnNames)	int	Yes
getConnection()	Connection	Yes
getFetchDirection()	int	Yes
getFetchSize()	int	Yes
getGeneratedKeys()	ResultSet	Yes
getMaxFieldSize()	int	Yes
getMaxRows()	int	Yes
getMoreResults()	boolean	Yes
getMoreResults(int current)	boolean	Yes
getResultSet()	ResultSet	Yes
getResultSetConcurrency()	int	Yes
getResultSetHoldability()	int	Yes
getResultSetType()	int	Yes
getQueryTimeout()	int	Yes
getUpdateCount()	int	Yes
getWarnings()	SQLWarning	Yes
isClosed()	boolean	Yes
isCloseOnCompletion()	boolean	Yes
isPoolable()	boolean	Yes
setCursorName(String name)	void	Yes
setEscapeProcessing(boolean enable)	void	Yes

Method Name	Return Type	JDBC 4 Is Supported Or Not
setFetchDirection(int direction)	void	Yes
setMaxFieldSize(int max)	void	Yes
setMaxRows(int max)	void	Yes
setPoolable(boolean poolable)	void	Yes
setQueryTimeout(int seconds)	void	Yes
setFetchSize(int rows)	void	Yes
cancel()	void	Yes
executeLargeUpdate(String sql)	long	No
getLargeUpdateCount()	long	No
executeLargeBatch()	long	No
executeLargeUpdate(String sql, int autoGeneratedKeys)	long	No
executeLargeUpdate(String sql, int[] columnIndexes)	long	No
executeLargeUpdate(String sql, String[] columnNames)	long	No

 NOTE

- Using setFetchSize can reduce the memory occupied by result sets on the client. Result sets are packaged into cursors and segmented for processing, which will increase the communication traffic between the database and the client, affecting performance.
- Database cursors are valid only within their transactions. If **setFetchSize** is set, set **setAutoCommit(false)** and commit transactions on the connection to flush service data to a database.
- The LargeUpdate method can only be used in JDBC 4.2 or later.

5.3.4.9 javax.sql.ConnectionPoolDataSource

This section describes **javax.sql.ConnectionPoolDataSource**, the API for data source connection pools.

Table 5-19 Support status for javax.sql.ConnectionPoolDataSource

Method Name	Return Type	JDBC 4 Is Supported Or Not
getPooledConnection()	PooledConnection	Yes
getPooledConnection(String user,String password)	PooledConnection	Yes

5.3.4.10 javax.sql.DataSource

This section describes **javax.sql.DataSource**, the interface for data sources.

Table 5-20 Support status for javax.sql.DataSource

Method Name	Return Type	Support JDBC 4
getConnection()	Connection	Yes
getConnection(String username,String password)	Connection	Yes
getLoginTimeout()	int	Yes
getLogWriter()	PrintWriter	Yes
setLoginTimeout(int seconds)	void	Yes
setLogWriter(PrintWriter out)	void	Yes

5.3.4.11 javax.sql.PooledConnection

This section describes **javax.sql.PooledConnection**, the connection API created by a connection pool.

Table 5-21 Support status for javax.sql.PooledConnection

Method Name	Return Type	JDBC 4 Is Supported Or Not
addConnectionEventListener(ConnectionEventListener listener)	void	Yes
close()	void	Yes
getConnection()	Connection	Yes
removeConnectionEventListener(ConnectionEventListener listener)	void	Yes

5.3.4.12 javax.naming.Context

javax.naming.Context is the context API for connection configuration.

Table 5-22 Support status for javax.naming.Context

Method	Return Type	Support JDBC 4
bind(Name name, Object obj)	void	Yes
bind(String name, Object obj)	void	Yes
lookup(Name name)	Object	Yes
lookup(String name)	Object	Yes
rebind(Name name, Object obj)	void	Yes
rebind(String name, Object obj)	void	Yes
rename(Name oldName, Name newName)	void	Yes
rename(String oldName, String newName)	void	Yes
unbind(Name name)	void	Yes
unbind(String name)	void	Yes

5.3.4.13 javax.naming.spi.InitialContextFactory

This section describes **javax.naming.spi.InitialContextFactory**, the initial context factory interface.

Table 5-23 Support status for javax.naming.spi.InitialContextFactory

Method Name	Return Type	Support JDBC 4
getInitialContext(Hashtable<?,?> environment)	Context	Yes

5.3.4.14 CopyManager

CopyManager is an API class provided by the JDBC driver in GaussDB. It is used to import data to GaussDB in batches.

Inheritance Relationship of CopyManager

The CopyManager class is in the **com.huawei.opengauss.jdbc.copy** package and inherits the java.lang.Object class. The declaration of the class is as follows:

```
public class CopyManager  
extends Object
```

Construction Method

```
public CopyManager(BaseConnection connection)  
  
throws SQLException
```

Common Methods

Table 5-24 Common methods of CopyManager

Method	Return Type	Description	throws
copyIn(String sql)	CopyIn	-	SQLException
copyIn(String sql, InputStream from)	long	Uses COPY FROM STDIN to quickly load data to tables in the database from InputStream.	SQLException,IOException
copyIn(String sql, InputStream from, int bufferSize)	long	Uses COPY FROM STDIN to quickly load data of a specified length to tables in the database from InputStream.	SQLException,IOException
copyIn(String sql, Reader from)	long	Uses COPY FROM STDIN to quickly load data to tables in the database from Reader.	SQLException,IOException
copyIn(String sql, Reader from, int bufferSize)	long	Uses COPY FROM STDIN to quickly load data of a specified length to tables in the database from Reader.	SQLException,IOException
copyOut(String sql)	CopyOut	-	SQLException
copyOut(String sql, OutputStream to)	long	Sends the result set of COPY TO STDOUT from the database to the OutputStream class.	SQLException,IOException

Method	Return Type	Description	throws
copyOut(String sql, Writer to)	long	Sends the result set of COPY TO STDOUT from the database to the Writer class.	SQLException,IOException

5.3.4.15 PGReplicationConnection

PGReplicationConnection is an API class provided by the JDBC driver in GaussDB. It is used to implement functions related to logical replication.

Inheritance Relationship of PGReplicationConnection

PGReplicationConnection is a logical replication API. Its implementation class is PGReplicationConnectionImpl, which is in the **com.huawei.opengauss.jdbc.replication** package. The declaration of the class is as follows:

```
public class PGReplicationConnection implements PGReplicationConnection
```

Constructor

```
public PGReplicationConnection(BaseConnection connection)
```

Common Methods

Table 5-25 Common methods of PGReplicationConnection

Method	Return Type	Description	Throws
createReplicationSlot()	ChainedCreateReplicationSlotBuilder	Creates a logical replication slot.	-
dropReplicationSlot(String slotName)	void	Drops a logical replication slot.	SQLException,IOException
replicationStream()	ChainedStreamBuilder	Enables logical replication.	-

5.3.4.16 PGReplicationStream

PGReplicationStream is an API class provided by the GaussDB JDBC driver. It is used to operate logical replication streams.

Inheritance Relationship of PGReplicationStream

PGReplicationStream is a logical replication API. Its implementation class is V3PGReplicationStream, which is in the **com.huawei.opengauss.jdbc.core.v3.replication** package. The declaration of the class is as follows:

```
public class V3PGReplicationStream implements PGReplicationStream
```

Constructor

```
public V3PGReplicationStream(CopyDual copyDual, LogSequenceNumber  
startLSN, long updateIntervalMs, ReplicationType replicationType)
```

Common Methods

Table 5-26 Common methods of PGReplicationConnection

Method Name	Return Type	Description	throws
close()	void	Ends the logical replication and releases resources.	SQLException
forceUpdateStatus()	void	Forcibly sends the LSN status received, flushed, and applied last time to the backend.	SQLException
getLastAppliedLSN()	LogSequenceNumber	Obtains the LSN when the primary node replays logs last time.	-
getLastFlushedLSN()	LogSequenceNumber	Obtains the LSN flushed by the primary node last time, that is, the LSN pushed by the current logical decoding.	-
getLastReceivedLSN()	LogSequenceNumber	Obtains the LSN received last time.	-
isClosed()	boolean	Determines whether the replication stream is disabled.	-
read()	ByteBuffer	Reads the next WAL record from the backend. If the data cannot be read, this method blocks the I/O read.	SQLException
readPending()	ByteBuffer	Reads the next WAL record from the backend. If the data cannot be read, this method does not block the I/O read.	SQLException

Method Name	Return Type	Description	throws
setAppliedLSN(LogSequenceNumber applied)	void	Sets the applied LSN.	-
setFlushedLSN(LogSequenceNumber flushed)	void	Sets the flushed LSN, which is sent to the backend at the next update to push the LSN on the server.	-

5.3.4.17 ChainedStreamBuilder

ChainedStreamBuilder is an API class provided by the GaussDB JDBC driver. It is used to build replication streams.

Inheritance Relationship of ChainedStreamBuilder

ChainedStreamBuilder is a logical replication API. Its implementation class is ReplicationStreamBuilder, which is in the **com.huawei.opengauss.jdbc.replication.fluent** package. The declaration of the class is as follows:

```
public class ReplicationStreamBuilder implements ChainedStreamBuilder
```

Constructor

```
public ReplicationStreamBuilder(final BaseConnection connection)
```

Common Methods

Table 5-27 Common methods of ReplicationStreamBuilder

Method Name	Return Type	Description	throws
logical()	ChainedLogicalStreamBuilder	Creates a logical replication stream.	-
physical()	ChainedPhysicalStreamBuilder	Creates a physical replication stream.	-

5.3.4.18 ChainedCommonStreamBuilder

ChainedCommonStreamBuilder is an API class provided by the GaussDB JDBC driver. It is used to specify common parameters for logical and physical replication.

Inheritance Relationship of ChainedCommonStreamBuilder

ChainedCommonStreamBuilder is an API for logical replication. The implementation abstract class is AbstractCreateSlotBuilder. The inheritance class is LogicalCreateSlotBuilder which is in the **com.huawei.opengauss.jdbc.replication.fluent.logical** package. The declaration of this class is as follows:

```
public class LogicalCreateSlotBuilder
    extends AbstractCreateSlotBuilder<ChainedLogicalCreateSlotBuilder>
    implements ChainedLogicalCreateSlotBuilder
```

Constructor

```
public LogicalCreateSlotBuilder(BaseConnection connection)
```

Common Methods

Table 5-28 Common methods of LogicalCreateSlotBuilder


Method	Return Type	Description	throws
withSlotName(String slotName)	T	Specifies the name of a replication slot.	-
withOutputPlugin(String outputPlugin)	ChainedLogicalCreateSlotBuilder	Plug-in name. Currently, mppdb_decoding is supported.	-
make()	void	Creates a slot with the specified parameters in the database.	SQLException
self()	ChainedLogicalCreateSlotBuilder	-	-

5.4 Development Based on ODBC

5.4.1 Development Process

Figure 5-3 shows the ODBC-based development process.

Figure 5-3 ODBC development process



5.4.2 Development Procedure

5.4.2.1 Obtaining a Source Code Package, ODBC Packages, and Dependent Libraries

You need to obtain the unixODBC source code package, ODBC packages, and dependency libraries for ODBC-based development.

Obtaining the Source Code Package

- To obtain the unixODBC source code package, visit <https://www.unixodbc.org/unixODBC-2.3.7.tar.gz>.
- To obtain the MD5 file, visit <https://www.unixodbc.org/unixODBC-2.3.7.tar.gz.md5>.

Use the MD5 file to verify the integrity of the unixODBC source code package. (Check whether the MD5 value is the same as that in the source code package.)

Obtaining ODBC Packages and Dependency Libraries

ODBC packages and dependency libraries support both Linux and Windows systems. Obtain the required files based on the OS you use.

- For Linux:

Obtain the package **GaussDB-Kernel-Database version number-OS information-64bit-Odbc.tar.gz** from the driver package. In the Linux OS, unixODBC header files (including **sql.h** and **sqlext.h**) and the library **libodbc.so** are required in application development.

The header files and the library can be obtained from the unixODBC-2.3.7 source code package.

- For Windows:

Obtain the package **GaussDB-Kernel_Database version number_Windows_X86_Odbc.tar.gz** (32-bit) or **GaussDB-Kernel-Database version number-Windows-X64-Odbc.tar.gz** (64-bit) from the driver package.

In the Windows OS, the required header files and library files are system-resident.

5.4.2.2 Connecting to a Database

5.4.2.2.1 Configuring a Data Source in the Linux OS

Before using ODBC to connect to a database, you need to prepare the required resources. To connect to a database, you need to configure the ODBC data source or use the ODBC APIs or the driver to implement the communication and interaction between an application and a database. This section describes how to configure a data source and connect to a database in the Linux OS.

Procedure

- Step 1** Install unixODBC. (By default, the unixODBC source code package has been obtained during environment preparation.) It does not matter if unixODBC of another version has been installed.

For example, to install unixODBC-2.3.7, run the commands below.

```
tar zxvf unixODBC-2.3.7.tar.gz
cd unixODBC-2.3.7
```

```
./configure --enable-gui=no # To perform compilation on an Arm server, add the configure parameter --  
build=aarch64-unknown-linux-gnu.  
make  
# The installation may require root permissions.  
make install
```

NOTE

- Currently, unixODBC-2.2.1 is not supported.
- By default, it is installed in the **/usr/local** directory. The data source file is generated in the **/usr/local/etc** directory, and the library file is generated in the **/usr/local/lib** directory.
- You can compile unixODBC with the **--enable-fastvalidate=yes** option to achieve higher performance. However, this option may cause an application that passes an invalid handle to the ODBC API to fail instead of returning an SQL_INVALID_HANDLE error.

Step 2 Replace the GaussDB client driver.

Decompress **GaussDB-Kernel-Database version number-OS information-64bit-Odbc.tar.gz**. After the decompression, the **lib** and **odbc** folders are generated. The **odbc** folder contains another **lib** folder. Copy all dynamic libraries in the **/lib** and **/odbc/lib** folders to the **/usr/local/lib** directory.

Step 3 Configure a data source.

1. Configure the ODBC driver file.

Add the following content to the **/usr/local/etc/odbcinst.ini** file:

```
[GaussMPP]
Driver64=/usr/local/lib/psqlodbcw.so
setup=/usr/local/lib/psqlodbcw.so
```

For descriptions of the parameters in the **odbcinst.ini** file, see [Table 5-29](#).

Table 5-29 odbcinst.ini configuration parameters

Parameter	Description	Example
[DriverName]	Driver name, corresponding to the driver in DSN.	[GaussMPP]
Driver64	Path of the dynamic driver library.	Driver64=/usr/local/lib/psqlodbcw.so
setup	Driver installation path, which is the same as the dynamic library path in Driver64.	setup=/usr/local/lib/psqlodbcw.so

2. Configure the data source file.

Add the following content to the **/usr/local/etc/odbc.ini** file:

```
[gaussdb]
Driver=GaussMPP
Servername=127.0.0.1 # Database server IP address
Database=postgres # Database name
Username=omm # Database username
Password= # Database user password
Port=8000 # Database listening port
Sslmode = allow
```

For descriptions of the parameters in the **odbc.ini** file, see [Table 5-30](#).

Table 5-30 odbc.ini configuration parameters

Parameter	Description	Example
[DSN]	Data source name.	[gaussdb]
Driver	Driver name, corresponding to DriverName in odbcinst.ini .	Driver=GaussMPP
Servername	Server IP address. Multiple IP addresses can be configured.	Servername=127.0.0.1

Parameter	Description	Example
Database	Name of the database to connect to.	Database=postgres
Username	Database username.	Username=omm
Password	Database user password. NOTE After a user establishes a connection, the ODBC driver automatically clears their password stored in memory. However, if this parameter is configured, unixODBC will cache data source files, which may cause the password to be stored in the memory for a long time. When you connect to an application, you are advised to send your password through an API instead of writing it in a data source configuration file. After the connection has been established, immediately clear the memory segment where your password is stored. NOTICE The password in the configuration file must comply with the following HTTP rules: 1. Characters must comply with the URL encoding specifications. For example, the exclamation mark (!) must be written as %21, and the percent sign (%) must be written as %25. Therefore, pay attention to the handling of the percent sign (%). 2. A plus sign (+) will be replaced by a space.	Password=*****
Port	Port number of the server.	Port=8000
Sslmode	Specifies whether to enable SSL. NOTE For details about the values of the Sslmode option, see Table 5-31.	Sslmode=allow
Debug	Specifies whether to enable the debugging mode. Value range: 0 to <i>INT_MAX</i>. – If this parameter is set to 0, the function is disabled. – If this parameter is set to a value greater than 0, the mylog file of the psqLODBC driver will be printed. The directory generated for storing logs is /tmp/. The default value is 0.	Debug=1

Parameter	Description	Example
UseServerSidePrepare	Specifies whether to enable the extended query protocol for the database. Value range: 0 and 1 – If this parameter is set to 0, the function is disabled. – If this parameter is set to 1, the function is enabled. The default value is 1.	UseServerSidePrepare=1
UseBatchProtocol	Specifies whether to enable the batch query protocol. If it is enabled, DML performance can be improved. Value range: 0 and 1 – If this parameter is set to 0, the batch query protocol is disabled (mainly for communication with earlier database versions). – If this parameter is set to 1 and support_batch_bind is set to on, the batch query protocol is enabled. The default value is 1.	UseBatchProtocol=1
ForExtensionConnector	Specifies whether the savepoint is sent. Value range: 0 and 1 – If this parameter is set to 0, the savepoint is sent. – If this parameter is set to 1, the savepoint is not sent. The default value is 1.	ForExtensionConnector=1

Parameter	Description	Example
ConnectionExtraInfo	Specifies whether to display the driver deployment path and process owner in the connection_info GUC parameter. Value range: 0 and 1 – If this parameter is set to 0, the information is not displayed. – If this parameter is set to 1, the information is displayed. The default value is 0. NOTE The default value is 0. If this parameter is set to 1, the ODBC driver reports the driver deployment path and process owner to the database and displays the information in the connection_info parameter. In this case, you can query the information from PG_STAT_ACTIVITY.	ConnectionExtraInfo=1
BoolsAsChar	Specifies whether to process Boolean values as characters. Value range: 0 and 1 – If this parameter is set to 0, the Boolean values are mapped to the SQL_BIT type. – If this parameter is set to 1, the Boolean values are mapped to the SQL_CHAR type. The default value is 1.	BoolsAsChar = 1
RowVersioning	Specifies whether to allow the application to check whether a row of data is modified by another user when the row of data is updated. Value range: 0 and 1 – If this parameter is set to 0, application detection is not allowed. – If this parameter is set to 1, application detection is allowed. The default value is 0.	RowVersioning=1

Parameter	Description	Example
ShowSystemTables	Specifies whether to regard the system catalog as a common SQL table by default. Value range: 0 and 1 – If this parameter is set to 0, the driver does not regard the system catalog as a common SQL table by default. – If this parameter is set to 1, the driver regards the system catalog as a common SQL table by default. The default value is 0.	ShowSystemTables=1
TcpUserTimeout	Specifies the maximum duration for which transmitted data can remain unacknowledged before the TCP connection is forcibly closed on an OS that supports the TCP_USER_TIMEOUT socket option. Value range: 0 to <i>INT_MAX</i>. The default value is 0. NOTE 0 indicates that the default value is used. Ignore this parameter for UDS connections. The unit is millisecond.	TcpUserTimeout=5000

The valid values of **Sslmode** are described in the following table.

Table 5-31 Sslmode options

Sslmode	Enable SSL Encryption	Description
disable	No	SSL connection is not enabled.
allow	Possible	If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified.
prefer	Possible	If the database supports SSL connection, SSL connection is recommended. However, authenticity of the database server will not be verified.

Sslmode	Enable SSL Encryption	Description
require	Yes	SSL connection is required and data is encrypted. However, authenticity of the database server will not be verified.
verify-ca	Yes	SSL connection is required and whether the database has a trusted certificate will be verified.
verify-full	Yes	SSL connection is required. In addition to the check scope specified by verify-ca , the system checks whether the name of the host where the database resides is the same as that in the certificate. GaussDB does not support this mode.

 NOTE

When establishing connections to the GaussDB server using ODBC, you can enable SSL connections to encrypt client and server communications. To enable SSL connection, you must have the server certificate, client certificate, and private key files. For details on how to obtain these files, see related documents and commands of OpenSSL.

Step 4 Configure the environment variables on the client.

```
vim ~/.bashrc
```

Add the following information to the configuration file:

```
export LD_LIBRARY_PATH=/usr/local/lib/:$LD_LIBRARY_PATH
export ODBC_SYSINI=/usr/local/etc
export ODBCINI=/usr/local/etc/odbc.ini
```

Step 5 Run the following command to validate the addition:

```
source ~/.bashrc
```

Step 6 Test the connection.

After the installation, the generated binary file is stored in the **/usr/bin** directory. You can run the **isql -v gaussdb** command (*gaussdb* is the data source name).

- If the following information is displayed, the configuration is correct and the connection succeeds:

```
+-----+
| Connected!          |
|                    |
| sql-statement       |
| help [tablename]    |
| quit                |
|                    |
+-----+
```

- If error information is displayed, the configuration is incorrect. Check whether the preceding configuration steps are correctly performed.

NOTE

When ODBC is used to connect to a database, the kernel parameters are set as follows:

```
SET extra_float_digits = 2;  
SET DateStyle = 'ISO';
```

These parameters may cause the ODBC and gsql clients to display inconsistent data, for example, date data display mode and floating-point precision representation. If the result is not as expected, you are advised to explicitly set these parameters in the ODBC application code.

----End

Troubleshooting

- [UnixODBC][Driver Manager]Can't open lib 'xxx/xxx/psqlodbcw.so' : file not found.
Possible causes:
 - The path configured in the **odbcinst.ini** file is incorrect.
Run **ls** to check the path in the error information, ensuring that the **psqlodbcw.so** file exists and you have execution permissions on it.
 - The dependent library of **psqlodbcw.so** does not exist or is not in system environment variables.
Run **ldd** to check the path in the error information. If **libodbc.so.1** or other unixODBC libraries do not exist, configure unixODBC again following the procedure provided in this section, and add the **lib** directory under its installation directory to **LD_LIBRARY_PATH**. If other libraries do not exist, add the **lib** directory under the ODBC driver package to **LD_LIBRARY_PATH**. If other standard libraries are missing, install them.
- [UnixODBC]connect to server failed: no such file or directory
Possible causes:
 - An incorrect or unreachable database IP address or port was configured.
Check the **Servename** and **Port** configuration items in data sources.
 - Server monitoring is improper.
If **Servename** and **Port** are correctly configured, ensure the proper NIC and port are monitored by following the database server configurations in the procedure in this section.
 - Firewall and network gatekeeper settings are improper.
Check firewall settings, and ensure that the database communication port is trusted.
Check to ensure network gatekeeper settings are proper (if any).
- [unixODBC]The password-stored method is not supported.
Possible causes:
The **sslmode** configuration item is not configured in the data sources.
Solution:
Set the configuration item to **allow** or a higher level. For details, see [Table 5-31](#).
- Server common name "xxxx" does not match host name "xxxxx"
Possible causes:

When **verify-full** is used for SSL encryption, the driver checks whether the host name in certificates is the same as the actual one.

Solution:

To solve this problem, use **verify-ca** to stop checking host names, or generate a set of CA certificates containing the actual host names.

- Driver's SQLAllocHandle on SQL_HANDLE_DBC failed

Possible causes:

The executable file (such as the **isql** tool of unixODBC) and the database driver (**psqlodbcw.so**) depend on different library versions of ODBC, such as **libodbc.so.1** and **libodbc.so.2**. You can verify this problem by using the following method:

```
ldd `which isql` | grep odbc  
ldd psqlodbcw.so | grep odbc
```

If the suffix digits of the outputs **libodbc.so** are different or indicate different physical disk files, this problem exists. Both **isql** and **psqlodbcw.so** require **libodbc.so** to be loaded. If they load different physical files, two sets of function lists with the same name are generated in a visible domain (the **libodbc.so.*** function export lists of unixODBC are the same). This results in conflicts and the database driver cannot be loaded.

Solution:

Uninstall the unnecessary unixODBC, such as **libodbc.so.2**, and create a soft link with the same name and the **.so.2** suffix for the remaining **libodbc.so.1** library.

- FATAL: Forbid remote connection with trust method!

For security reasons, the primary database node prevents unauthorized access from other nodes within the database.

To access the primary database node from inside the database, deploy the ODBC program on the host where the primary database node is located and set the server address to **127.0.0.1**. It is recommended that the service system be deployed outside the database. Otherwise, the database performance may be affected.

- [unixODBC][Driver Manager]Invalid attribute value

The unixODBC version may not be the recommended one. You are advised to run the **odbcinst --version** command to check the unixODBC version in the environment.

- authentication method 10 not supported.

If this error occurs on an open-source client, the cause may be:

The database stores only the SHA-256 hash of the password, but the open-source client supports only MD5 hashes.

NOTE

- The database stores the hashes of user passwords instead of actual passwords.
- If a password is updated or a user is created, both types of hashes will be stored, compatible with open-source authentication protocols.
- An MD5 hash can only be generated using the original password, but the password cannot be obtained by reversing its SHA-256 hash. Passwords in the source version will only have SHA-256 hashes and not support MD5 authentication.
- MD5 has lower security and poses security risks. Therefore, you are advised to use a more secure cryptographic algorithm.

To solve this problem, you can update the user password (see [ALTER USER](#)) or create a user (see [CREATE USER](#)) having the same permissions as the faulty user.

- unsupported frontend protocol 3.51: server supports 1.0 to 3.0

The database version is too early or the database is an open-source database. Use the driver of the required version to connect to the database.

- FATAL: GSS authentication method is not allowed because XXXX user password is not disabled.

In **pg_hba.conf** of the target primary database node, the authentication mode is set to **gss** for authenticating the IP address of the current client. However, this authentication algorithm cannot authenticate clients. Change the authentication algorithm to **sha256** and try again.

- isql: error while loading shared libraries:xxx

The dynamic library does not exist in the environment. You need to install the corresponding library.

5.4.2.2.2 Configuring a Data Source in the Windows OS

Configure an ODBC data source using the ODBC data source manager preinstalled in the Windows OS.

Procedure

Step 1 Replace the GaussDB client driver.

Decompress **GaussDB-Kernel-Database version number-Windows-Odbc.tar.gz** and click **psqlodbc.exe** (32-bit) to install the driver.

Step 2 Open the driver manager.

Use the ODBC Driver Manager for the 32-bit OS to configure the data source. (Currently, only the ODBC Driver Manager for the 32-bit OS is supported. The following description assumes that the OS is installed on drive C. If the OS is installed on another drive, change the path accordingly.)

- If you want to use 32-bit ODBC driver manager in a 64-bit OS, open **C:\Windows\SysWOW64\odbcad32.exe**. Do not choose **Control Panel > Administrative Tools > Data Sources (ODBC)**.

NOTE

WoW64 is short for Windows 32-bit on Windows 64-bit. **C:\Windows\SysWOW64** stores the 32-bit environment on a 64-bit system. **C:\Windows\System32** stores the environment consistent with the current OS. For technical details, see Windows technical documents.

- For a 32-bit OS, open **C:\Windows\System32\odbcad32.exe** or choose **Computer > Control Panel > Administrative Tools > Data Sources (ODBC)** to open Driver Manager.

Step 3 Configure a data source.

On the **User DSN** tab, click **Add** and choose **PostgreSQL Unicode** for setup.

The screenshot shows the 'PostgreSQL ANSI ODBC Driver (psqlODBC) Setup' window. It contains the following fields and controls:

- Data Source:** Text box with 'gaussdb1'.
- Description:** Text box.
- Database:** Text box with 'postgres'.
- SSL Mode:** Dropdown menu with 'disable' selected.
- Server:** Text box with '10.145.130.27'.
- Port:** Text box with '25308'.
- User Name:** Text box with 'xxxx'.
- Password:** Password field with 12 dots.
- AutoBalance:** Text box with '1'.
- Refresh Time:** Text box with '1'.
- Priority:** Text box.
- Options:** A group box containing three buttons: 'Datasource', 'Global', and 'Manage DSN'.
- Action Buttons:** 'Test', 'Save', and 'Cancel' buttons at the bottom right.

NOTICE

The entered username and password will be recorded in the Windows registry and you do not need to enter them again when connecting to the database next time. For security purposes, you are advised to delete sensitive information before clicking **Save** and enter the required username and password again when using ODBC APIs to connect to the database.

Step 4 Enable the SSL mode.

Change the value of **SSL Mode** in [Step 3](#) to **require**.

Table 5-32 sslmode options

sslmode	Whether SSL Encryption Is Enabled	Description
disable	No	SSL connection is not enabled.
allow	Possible	If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified.
prefer	Possible	If the database supports SSL connection, SSL connection is preferred. However, authenticity of the database server will not be verified.
require	Yes	SSL connection is required and data is encrypted. However, authenticity of the database server will not be verified.

sslmode	Whether SSL Encryption Is Enabled	Description
verify-ca	Yes	SSL connection is required and whether the database has a trusted certificate will be verified. Currently, Windows ODBC does not support the cert authentication.
verify-full	Yes	SSL connection is required. In addition to the check scope specified by verify-ca , the system checks whether the name of the host where the database resides is the same as that in the certificate. Currently, Windows ODBC does not support the cert authentication.


NOTE

When establishing connections to the GaussDB server using ODBC, you can enable SSL connections to encrypt client and server communications. To enable SSL connections, you must have the server certificate, client certificate, and private key files. For details on how to obtain these files, see related documents and commands of OpenSSL.

Step 5 Test the connection.

Click **Test**.

- If the following information is displayed, the configuration is correct and the connection succeeds.



- If error information is displayed, the configuration is incorrect. Check the configuration.

 NOTE

When ODBC is used to connect to a database, the kernel parameters are set as follows:

```
SET extra_float_digits = 2;  
SET DateStyle = 'ISO';
```

These parameters may cause the ODBC and gsql clients to display inconsistent data, for example, date data display mode and floating-point precision representation. If the result is not as expected, you are advised to explicitly set these parameters in the ODBC application code.

----End

FAQ

- connect to server failed: no such file or directory
Possible causes:
 - An incorrect or unreachable database IP address or port was configured.
Check the **Server** and **Port** configuration items in data sources.
 - Server monitoring is improper.
If **Server** and **Port** are correctly configured, ensure the proper NIC and port are monitored by following the database server configurations in the procedure in this section.
 - Firewall and network gatekeeper settings are improper.
Check firewall settings, and ensure that the database communication port is trusted.
Check to ensure network gatekeeper settings are proper (if any).
- The password-stored method is not supported.
Possible causes:
sslmode is not configured for the data source. Set this configuration item to **allow** or a higher level to enable SSL connections. For details on **sslmode**, see [Table 5-32](#).
- authentication method 10 not supported.
If this error occurs on an open-source client, the cause may be:
The database stores only the SHA-256 hash of the password, but the open-source client supports only MD5 hashes.

 NOTE

- The database stores the hashes of user passwords instead of actual passwords.
 - If a password is updated or a user is created, both types of hashes will be stored, compatible with open-source authentication protocols.
 - An MD5 hash can only be generated using the original password, but the password cannot be obtained by reversing its SHA-256 hash. Passwords in the source version will only have SHA-256 hashes and not support MD5 authentication.
 - The MD5 encryption algorithm has lower security and poses security risks.
Therefore, you are advised to use a more secure encryption algorithm.
- To solve this problem, you can update the user password (see [ALTER USER](#)) or create a user (see [CREATE USER](#)) having the same permissions as the faulty user.
- unsupported frontend protocol 3.51: server supports 1.0 to 3.0

The database version is too early or the database is an open-source database. Use the driver of the required version to connect to the database.

- FATAL: GSS authentication method is not allowed because XXXX user password is not disabled.

In **pg_hba.conf** of the target DN, the authentication mode is set to **gss** for authenticating the IP address of the current client. However, this authentication algorithm cannot authenticate clients. Change the authentication algorithm to **sha256** and try again.

5.4.2.2.3 APIs for Connecting to a Database

After the database connection test is successful, the ODBC APIs provide a group of functions to connect to the database, as shown in [Table 5-33](#).

Table 5-33 API description

Function	API
Allocate a handle.	SQLAllocHandle is a generic function for allocating a handle. It can replace the following functions: • SQLAllocEnv: allocates an environment handle. • SQLAllocConnect: allocates a connection handle. • SQLAllocStmt: allocates a statement handle.
Set environment attributes.	SQLSetEnvAttr
Set connection attributes.	SQLSetConnectAttr
Connect to a database.	SQLConnect

The following uses the development source program **DBtest.c** as an example (for details about the complete example, see [Obtaining and Processing Data in a Database](#)).

```
// DBtest.c (compile with: libodbc.so)
// For details about program header files and global variables, see the complete example.

// Allocate an environment handle.
V_OD_erg = SQLAllocHandle(SQL_HANDLE_ENV,SQL_NULL_HANDLE,&V_OD_Env);
if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
{
    printf("Error AllocHandle\n");
    exit(0);
}

// Set the version information (environment attributes).
SQLSetEnvAttr(V_OD_Env, SQL_ATTR_ODBC_VERSION, (void*)SQL_OV_ODBC3, 0);

// Allocate a connection handle.
V_OD_erg = SQLAllocHandle(SQL_HANDLE_DBC, V_OD_Env, &V_OD_hdbc);
if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
{
```

```
SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
exit(0);
}

//Obtain the username and password.
char *userName;
userName = getenv("EXAMPLE_USERNAME_ENV");
char *password;
password = getenv("EXAMPLE_PASSWORD_ENV");

// Set connection attributes.
SQLSetConnectAttr(V_OD_hdbc, SQL_ATTR_AUTOCOMMIT,(SQLPOINTER *)SQL_AUTOCOMMIT_ON, 0);

// Connect to the database. userName and password indicate the username and password for connecting
to the database respectively.
// If the username and password have been set in the odbc.ini file, you can retain "". However, you are
advised not to do so because the username and password will be disclosed if the permission for odbc.ini is
abused.
V_OD_erg = SQLConnect(V_OD_hdbc, (SQLCHAR*) "gaussdb", SQL_NTS,
(SQLCHAR*) userName, SQL_NTS, (SQLCHAR*) password, SQL_NTS);
if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
{
    printf("Error SQLConnect %d\n",V_OD_erg);
    SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
    exit(0);
}
printf("Connected !\n");
```

5.4.2.3 Executing SQL Statements

To help users interact with the database, the ODBC provides APIs for executing SQL statements, as shown in [Table 5-34](#).

Table 5-34 API description

Function	API
Set statement attributes.	SQLSetStmtAttr
Prepare an SQL statement for execution.	SQLPrepare
Run a prepared SQL statement.	SQLExecute
Bind the parameter marker of an SQL statement to a buffer.	SQLBindParameter
Run an SQL statement directly.	SQLExecDirect

NOTE

- ODBC connects applications to the database and delivers the SQL statements sent by an application to the database. It does not parse the SQL syntax. Therefore, when confidential information (such as a plaintext password) is written into the SQL statement sent by an application, the confidential information is exposed in the driver log.
- If an execution request (not in a transaction block) received in the database contains multiple statements, the request is packed into a transaction. If one of the statements fails, the entire request will be rolled back.

The following is an example (for details about the complete example, see [Obtaining and Processing Data in a Database](#)):

```
// Set statement attributes.
SQLSetStmtAttr(V_OD_hstmt, SQL_ATTR_QUERY_TIMEOUT, (SQLPOINTER *)3, 0);

// Allocate a statement handle.
SQLAllocHandle(SQL_HANDLE_STMT, V_OD_hdbc, &V_OD_hstmt);

// Run an SQL statement directly.
SQLExecDirect(V_OD_hstmt, "drop table IF EXISTS customer_t1", SQL_NTS);
SQLExecDirect(V_OD_hstmt, "CREATE TABLE customer_t1(c_customer_sk INTEGER, c_customer_name
VARCHAR(32));", SQL_NTS);
SQLExecDirect(V_OD_hstmt, "insert into customer_t1 values(25, 'li')", SQL_NTS);

// Prepare for execution.
SQLPrepare(V_OD_hstmt, "insert into customer_t1 values(?)", SQL_NTS);

// Bind parameters.
SQLBindParameter(V_OD_hstmt,1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER, 0, 0,
&value, 0, NULL);

// Run the prepared statement.
SQLExecute(V_OD_hstmt);    SQLExecDirect(V_OD_hstmt, "select c_customer_sk from customer_t1",
SQL_NTS);
```

5.4.2.4 Processing Data in a Result Set

The ODBC obtains data from the database and provides the data to the application program for processing. Functions of ODBC processing include but are not limited to: retrieving data, displaying data, processing data, transmitting data, and implementing service logic.

The ODBC provides APIs for processing data in a result set, as shown in [Table 5-35](#).

Table 5-35 API description

Function	API
Bind a buffer to a column in the result set.	SQLBindCol
Fetch the next row (or rows) from the result set.	SQLFetch
Return data in a column of the result set.	SQLGetData
Get the column information from the result set.	SQLColAttribute
Return the error message of the last operation.	SQLGetDiagRec

The following is an example (for details about the complete example, see [Obtaining and Processing Data in a Database](#)):

```
// After the SQL statement is executed, obtain the attributes of a column in the result set.
SQLColAttribute(V_OD_hstmt,1,SQL_DESC_TYPE,typename,100,NULL,NULL);
```



```
printf("SQLColAttribute %s\n",typename);

// Bind the result set.
SQLBindCol(V_OD_hstmt,1,SQL_C_SLONG, (SQLPOINTER*)&V_OD_buffer,150,
           (SQLLEN *)&V_OD_err);

// Use SQLFetch to obtain data from the result set.
V_OD_erg=SQLFetch(V_OD_hstmt);

// Use SQLGetData to obtain and return data.
while(V_OD_erg != SQL_NO_DATA)
{
    SQLGetData(V_OD_hstmt,1,SQL_C_SLONG,(SQLPOINTER*)&V_OD_id,0,NULL);
    printf("SQLGetData ----ID = %d\n",V_OD_id);
    V_OD_erg=SQLFetch(V_OD_hstmt);
};
printf("Done !\n");
```

5.4.2.5 Closing a Connection

The ODBC provides APIs for disconnecting from a database and releasing resources, as shown in [Table 5-36](#).

Table 5-36 API description

Function	API
Disconnect from a data source.	SQLDisconnect
Release a handle.	SQLFreeHandle is a generic function for releasing a handle. It can replace the following functions: • SQLFreeEnv: releases an environment handle. • SQLFreeConnect: releases a connection handle. • SQLFreeStmt: releases a statement handle.

The following is an example (for details about the complete example, see [Obtaining and Processing Data in a Database](#)):

```
// Disconnect from the data source and release handles.
SQLFreeHandle(SQL_HANDLE_STMT,V_OD_hstmt);
SQLDisconnect(V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_DBC,V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
```

5.4.3 Typical Application Development Examples

5.4.3.1 Typical Application Scenarios and Configurations

Log Diagnosis Scenario

ODBC logs are classified into unixODBC driver manager logs and psqLODBC driver logs. The former is used to trace whether the application API is successfully

executed, and the latter is used to locate problems based on DFX logs generated during underlying implementation.

The unixODBC log needs to be configured in the **odbcinst.ini** file:

```
[ODBC]
Trace=Yes
TraceFile=/path/to/odbctrace.log

[GaussMPP]
Driver64=/usr/local/lib/psqlodbcw.so
setup=/usr/local/lib/psqlodbcw.so
```

You only need to add the following information to the **odbc.ini** file:

```
[gaussdb]
Driver=GaussMPP
Servername=10.10.0.13 (database server IP address)
...
Debug=1 (Enable the debug log function of the driver.)
```

NOTE

The unixODBC logs are generated in the path configured by **TraceFile**. The psqlODBC generates the **mylog_XXX.log** file in the **/tmp/** directory.

Automatic Primary/Standby Switchover

Example Scenario

If a database instance is configured with one primary DN and multiple standby DNs, write the IP addresses of all DNs into the configuration file. ODBC automatically searches for the primary DN and establishes a connection with it. When a primary/standby switchover occurs, ODBC can also connect to the new primary DN.

5.4.3.2 Obtaining and Processing Data in a Database

NOTE

- In Windows, ODBC application code can be compiled using the Minimalist GNU for Windows (MinGW) compiler. The compilation command is as follows:

```
gcc odbctest.c -o odbctest -lodb32
```

Run the following command:

```
./odbctest.exe
```

- In Linux, ODBC application code can be compiled using the GNU Compiler Collection (GCC). The compilation command is as follows:

```
gcc odbctest.c -o odbctest -lodb
```

Run the following command:

```
./odbctest
```

If **sql.h** or API cannot be found during compilation, manually connect to the header file and dynamic library of unixODBC.

```
gcc -I /home/omm/unixodbc/include -L /home/omm/unixodbc/lib odbctest.c -o odbctest -lodb
```

This example illustrates how to obtain and process data in GaussDB through ODBC.

Prerequisite: The data source has been configured successfully. For Linux OS, see [Configuring a Data Source in the Linux OS](#). For Windows OS, see [Configuring a Data Source in the Windows OS](#).

```
// DBtest.c (compile with: libodbc.so)
// In this example, the username and password are stored in environment variables. Before running this
// example, set the environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV in the
// local environment.
#ifdef WIN32
#include <windows.h>
#else
#include <stdio.h>
#include <stdlib.h>
#include <sql.h>
#include <sqlext.h>
SQLHENV    V_OD_Env;    // Handle ODBC environment
SQLHSTMT    V_OD_hstmt; // Handle statement
SQLHDBC    V_OD_hdbc;   // Handle connection
char        typename[100];
SQLINTEGER  value = 100;
SQLINTEGER  V_OD_erg,V_OD_buffer,V_OD_err,V_OD_id;
int main(int argc,char *argv[])
{
    // 1. Allocate an environment handle.
    V_OD_erg = SQLAllocHandle(SQL_HANDLE_ENV,SQL_NULL_HANDLE,&V_OD_Env);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        printf("Error AllocHandle\n");
        exit(0);
    }

    // 2. Set the version information (environment attributes).
    SQLSetEnvAttr(V_OD_Env, SQL_ATTR_ODBC_VERSION, (void*)SQL_OV_ODBC3, 0);

    // 3. Allocate a connection handle.
    V_OD_erg = SQLAllocHandle(SQL_HANDLE_DBC, V_OD_Env, &V_OD_hdbc);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
        exit(0);
    }

    // Obtain the username and password.
    char *userName;
    userName = getenv("EXAMPLE_USERNAME_ENV");
    char *password;
    password = getenv("EXAMPLE_PASSWORD_ENV");

    // 4. Set connection attributes.
    SQLSetConnectAttr(V_OD_hdbc, SQL_ATTR_AUTOCOMMIT,(SQLPOINTER *)SQL_AUTOCOMMIT_ON, 0);

    // 5. Connect to the database. userName and password indicate the username and password for
    // connecting to the database, respectively.
    // If the username and password have been set in the odbc.ini file, you can retain "". However, you are
    // advised not to do so because the username and password will be disclosed if the permission for odbc.ini is
    // abused.
    V_OD_erg = SQLConnect(V_OD_hdbc, (SQLCHAR*) "gaussdb", SQL_NTS,
        (SQLCHAR*) userName, SQL_NTS, (SQLCHAR*) password, SQL_NTS);
    if ((V_OD_erg != SQL_SUCCESS) && (V_OD_erg != SQL_SUCCESS_WITH_INFO))
    {
        printf("Error SQLConnect %d\n",V_OD_erg);
        SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
        exit(0);
    }
    printf("Connected !\n");

    // 6. Set statement attributes.
    SQLSetStmtAttr(V_OD_hstmt, SQL_ATTR_QUERY_TIMEOUT, (SQLPOINTER *)3,0);

    // 7. Allocate a statement handle.
    SQLAllocHandle(SQL_HANDLE_STMT, V_OD_hdbc, &V_OD_hstmt);

    // 8. Run SQL statements.
```

```
SQLExecDirect(V_OD_hstmt, "drop table IF EXISTS customer_t1", SQL_NTS);
SQLExecDirect(V_OD_hstmt, "CREATE TABLE customer_t1(c_customer_sk INTEGER, c_customer_name
VARCHAR(32));", SQL_NTS);
SQLExecDirect(V_OD_hstmt, "insert into customer_t1 values(25,'li')", SQL_NTS);

// 9. Prepare for execution.
SQLPrepare(V_OD_hstmt, "insert into customer_t1 values(?)", SQL_NTS);

// 10. Bind parameters.
SQLBindParameter(V_OD_hstmt, 1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER, 0, 0,
&value, 0, NULL);

// 11. Run prepared statements.
SQLExecute(V_OD_hstmt);
SQLExecDirect(V_OD_hstmt, "select c_customer_sk from customer_t1", SQL_NTS);

// 12. Obtain attributes of a specific column in the result set.
SQLColAttribute(V_OD_hstmt, 1, SQL_DESC_TYPE, typename, 100, NULL, NULL);
printf("SQLColAttribute %s\n", typename);

// 13. Bind the result set.
SQLBindCol(V_OD_hstmt, 1, SQL_C_SLONG, (SQLPOINTER)&V_OD_buffer, 150,
(SQLLEN *)&V_OD_err);

// 14. Obtain data in the result set by executing SQLFetch.
V_OD_erg=SQLFetch(V_OD_hstmt);

// 15. Obtain and return data by executing SQLGetData.
while(V_OD_erg != SQL_NO_DATA)
{
    SQLGetData(V_OD_hstmt, 1, SQL_C_SLONG, (SQLPOINTER)&V_OD_id, 0, NULL);
    printf("SQLGetData ----ID = %d\n", V_OD_id);
    V_OD_erg=SQLFetch(V_OD_hstmt);
};
printf("Done !\n");

// 16. Disconnect data source connections and release handles.
SQLFreeHandle(SQL_HANDLE_STMT, V_OD_hstmt);
SQLDisconnect(V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_DBC, V_OD_hdbc);
SQLFreeHandle(SQL_HANDLE_ENV, V_OD_Env);
return(0);
}
```

The running result is as follows:

```
Connected !
SQLColAttribute
SQLGetData ----ID = 25
SQLGetData ----ID = 100
Done!
```

5.4.3.3 Batch Binding

Prerequisite: The data source has been configured successfully. For Linux OS, see [Configuring a Data Source in the Linux OS](#). For Windows OS, see [Configuring a Data Source in the Windows OS](#).

```
/*
*****
* Enable UseBatchProtocol in the data source and set the database parameter support_batch_bind to on.
* The CHECK_ERROR command is used to check and print error information.
* This example is used to interactively obtain the DSN, data volume to be processed, and volume of ignored
data from users, and insert required data into the test_odbc_batch_insert table.
*****
*/
#ifdef WIN32
#include <windows.h>
#endif
#include <stdio.h>
```

```
#include <stdlib.h>
#include <sql.h>
#include <sqlext.h>
#include <string.h>
void Exec(SQLHDBC hdbc, SQLCHAR* sql)
{
    SQLRETURN retcode;          // Return status
    SQLHSTMT hstmt = SQL_NULL_HSTMT; // Statement handle
    SQLCHAR loginfo[2048];

    // Allocate a statement handle.
    retcode = SQLAllocHandle(SQL_HANDLE_STMT, hdbc, &hstmt);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLAllocHandle(SQL_HANDLE_STMT) failed");
        return;
    }

    // Prepare a statement.
    retcode = SQLPrepare(hstmt, (SQLCHAR*) sql, SQL_NTS);
    sprintf((char*)loginfo, "SQLPrepare log: %s", (char*)sql);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLPrepare(hstmt, (SQLCHAR*) sql, SQL_NTS) failed");
        return;
    }

    // Execute the statement.
    retcode = SQLExecute(hstmt);
    sprintf((char*)loginfo, "SQLExecute stmt log: %s", (char*)sql);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLExecute(hstmt) failed");
        return;
    }

    // Release the handle.
    retcode = SQLFreeHandle(SQL_HANDLE_STMT, hstmt);
    sprintf((char*)loginfo, "SQLFreeHandle stmt log: %s", (char*)sql);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLFreeHandle(SQL_HANDLE_STMT, hstmt) failed");
        return;
    }
}

int main ()
{
    SQLHENV henv = SQL_NULL_HENV;
    SQLHDBC hdbc = SQL_NULL_HDBC;
    long int batchCount = 1000; // Amount of data that is bound in batches
    SQLLEN rowsCount = 0;
    int ignoreCount = 0;        // Amount of data that is not imported to the database among the data
    // that is bound in batches
    int i = 0;
    SQLRETURN retcode;
    SQLCHAR dsn[1024] = {'\0'};
    SQLCHAR loginfo[2048];
    do
    {
        if (ignoreCount > batchCount)
        {
            printf("ignoreCount(%d) should be less than batchCount(%d)\n", ignoreCount, batchCount);
        }
    }while(ignoreCount > batchCount);

    // Allocate an environment handle.
    retcode = SQLAllocHandle(SQL_HANDLE_ENV, SQL_NULL_HANDLE, &henv);

    if (!SQL_SUCCEEDED(retcode)) {
```

```
    printf("SQLAllocHandle failed");
    goto exit;
}

// Set the ODBC version.
retcode = SQLSetEnvAttr(henv, SQL_ATTR_ODBC_VERSION,
                        (SQLPOINTER*)SQL_OV_ODBC3, 0);

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLSetEnvAttr failed");
    goto exit;
}

// Allocate connections.
retcode = SQLAllocHandle(SQL_HANDLE_DBC, henv, &hdbc);

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLAllocHandle failed");
    goto exit;
}

// Set login timeout.
retcode = SQLSetConnectAttr(hdbc, SQL_LOGIN_TIMEOUT, (SQLPOINTER)5, 0);

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLSetConnectAttr failed");
    goto exit;
}

// Set the automatic committing.
retcode = SQLSetConnectAttr(hdbc, SQL_ATTR_AUTOCOMMIT,
                            (SQLPOINTER)(1), 0);

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLSetConnectAttr failed");
    goto exit;
}

// Connect to the database.
sprintf(loginfo, "SQLConnect(DSN:%s)", dsn);
retcode = SQLConnect(hdbc, (SQLCHAR*) "gaussdb", SQL_NTS,
                    (SQLCHAR*) NULL, 0, NULL, 0);

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLConnect failed");
    goto exit;
}

// Initialize the table information.
Exec(hdbc, "drop table if exists test_odbc_batch_insert");
Exec(hdbc, "create table test_odbc_batch_insert(id int primary key, col varchar2(50))");
// The code constructs the data to be inserted based on the data volume entered by users.
{
    SQLRETURN retcode;
    SQLHSTMT hstmtinsrt = SQL_NULL_HSTMT;
    SQLCHAR    *sql = NULL;
    SQLINTEGER  *ids = NULL;
    SQLCHAR    *cols = NULL;
    SQLLEN      *bufLenIds = NULL;
    SQLLEN      *bufLenCols = NULL;
    SQLUSMALLINT *operptr = NULL;
    SQLUSMALLINT *statusptr = NULL;
    SQLULEN     process = 0;

    // Data is constructed by column. Each column is stored continuously.
    ids = (SQLINTEGER*)malloc(sizeof(ids[0]) * batchCount);
    cols = (SQLCHAR*)malloc(sizeof(cols[0]) * batchCount * 50);

    // Data size in each row for a column
```

```
    bufLenIds = (SQLLEN*)malloc(sizeof(bufLenIds[0]) * batchCount);
    bufLenCols = (SQLLEN*)malloc(sizeof(bufLenCols[0]) * batchCount);

    // The value SQL_PARAM_IGNORE or SQL_PARAM_PROCEED specifies whether a row needs to be
    processed.
    operptr = (SQLUSMALLINT*)malloc(sizeof(operptr[0]) * batchCount);
    memset(operptr, 0, sizeof(operptr[0]) * batchCount);

    // Processing result of the row
    // Note: In the database, a statement belongs to one transaction. Therefore, data is processed as a
    unit. Either all data is inserted successfully or all data fails to be inserted.
    statusptr = (SQLUSMALLINT*)malloc(sizeof(statusptr[0]) * batchCount);
    memset(statusptr, 88, sizeof(statusptr[0]) * batchCount);
    if (NULL == ids || NULL == cols || NULL == bufLenCols || NULL == bufLenIds)
    {
        fprintf(stderr, "FAILED:\tmalloc data memory failed\n");
        goto exit;
    }
    for (i = 0; i < batchCount; i++)
    {
        ids[i] = i;
        sprintf(cols + 50 * i, "column test value %d", i);
        bufLenIds[i] = sizeof(ids[i]);
        bufLenCols[i] = strlen(cols + 50 * i);
        operptr[i] = (i < ignoreCount) ? SQL_PARAM_IGNORE : SQL_PARAM_PROCEED;
    }

    // Allocate a statement handle.
    retcode = SQLAllocHandle(SQL_HANDLE_STMT, hdbc, &hstmtinesrt);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLAllocHandle failed");
        goto exit;
    }

    // Prepare a statement.
    sql = (SQLCHAR*)"insert into test_odbc_batch_insert values(?, ?)";
    retcode = SQLPrepare(hstmtinesrt, (SQLCHAR*) sql, SQL_NTS);
    sprintf((char*)loginfo, "SQLPrepare log: %s", (char*)sql);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLPrepare failed");
        goto exit;
    }

    retcode = SQLSetStmtAttr(hstmtinesrt, SQL_ATTR_PARAMSET_SIZE, (SQLPOINTER)batchCount,
    sizeof(batchCount));

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLSetStmtAttr failed");
        goto exit;
    }

    retcode = SQLBindParameter(hstmtinesrt, 1, SQL_PARAM_INPUT, SQL_C_SLONG, SQL_INTEGER,
    sizeof(ids[0]), 0,&(ids[0]), 0, bufLenIds);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLBindParameter failed");
        goto exit;
    }

    retcode = SQLBindParameter(hstmtinesrt, 2, SQL_PARAM_INPUT, SQL_C_CHAR, SQL_CHAR, 50, 50,
    cols, 50, bufLenCols);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLBindParameter failed");
        goto exit;
    }

    retcode = SQLSetStmtAttr(hstmtinesrt, SQL_ATTR_PARAMS_PROCESSED_PTR, (SQLPOINTER)&process,
    sizeof(process));

    if (!SQL_SUCCEEDED(retcode)) {
```

```
        printf("SQLSetStmtAttr failed");
        goto exit;
    }
    retcode = SQLSetStmtAttr(hstmtinesrt, SQL_ATTR_PARAM_STATUS_PTR, (SQLPOINTER)statusptr,
sizeof(statusptr[0]) * batchCount);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLSetStmtAttr failed");
        goto exit;
    }
    retcode = SQLSetStmtAttr(hstmtinesrt, SQL_ATTR_PARAM_OPERATION_PTR, (SQLPOINTER)operptr,
sizeof(operptr[0]) * batchCount);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLSetStmtAttr failed");
        goto exit;
    }
    retcode = SQLExecute(hstmtinesrt);
    sprintf((char*)loginfo, "SQLExecute stmt log: %s", (char*)sql);

    if (!SQL_SUCCEEDED(retcode)) {
        printf("SQLExecute(hstmtinesrt) failed");
        goto exit;
        retcode = SQLRowCount(hstmtinesrt, &rowsCount);

        if (!SQL_SUCCEEDED(retcode)) {
            printf("SQLRowCount failed");
            goto exit;
        }
        if (rowsCount != (batchCount - ignoreCount))
        {
            sprintf(loginfo, "(batchCount - ignoreCount)(%d) != rowsCount(%d)", (batchCount -
ignoreCount), rowsCount);

            if (!SQL_SUCCEEDED(retcode)) {
                printf("SQLExecute failed");
                goto exit;
            }
        }
        else
        {
            sprintf(loginfo, "(batchCount - ignoreCount)(%d) == rowsCount(%d)", (batchCount -
ignoreCount), rowsCount);

            if (!SQL_SUCCEEDED(retcode)) {
                printf("SQLExecute failed");
                goto exit;
            }
        }
    }

    // Check the number of returned rows.
    if (rowsCount != process)
    {
        sprintf(loginfo, "process(%d) != rowsCount(%d)", process, rowsCount);

        if (!SQL_SUCCEEDED(retcode)) {
            printf("SQLExecute failed");
            goto exit;
        }
    }
    else
    {
        sprintf(loginfo, "process(%d) == rowsCount(%d)", process, rowsCount);

        if (!SQL_SUCCEEDED(retcode)) {
            printf("SQLExecute failed");
            goto exit;
        }
    }
}
```



```
for (i = 0; i < batchCount; i++)
{
    if (i < ignoreCount)
    {
        if (statusptr[i] != SQL_PARAM_UNUSED)
        {
            sprintf(loginfo, "statusptr[%d](%d) != SQL_PARAM_UNUSED", i, statusptr[i]);

            if (!SQL_SUCCEEDED(retcode)) {
                printf("SQLExecute failed");
                goto exit;
            }
        }
    }
    else if (statusptr[i] != SQL_PARAM_SUCCESS)
    {
        sprintf(loginfo, "statusptr[%d](%d) != SQL_PARAM_SUCCESS", i, statusptr[i]);

        if (!SQL_SUCCEEDED(retcode)) {
            printf("SQLExecute failed");
            goto exit;
        }
    }
}
retcode = SQLFreeHandle(SQL_HANDLE_STMT, hstmtinesrt);
sprintf((char*)loginfo, "SQLFreeHandle hstmtinesrt");

if (!SQL_SUCCEEDED(retcode)) {
    printf("SQLFreeHandle failed");
    goto exit;
}
}
}
exit:
(void) printf ("\nComplete.\n");

// Close the connection.
if (hdbc != SQL_NULL_HDBC) {
    SQLDisconnect(hdbc);
    SQLFreeHandle(SQL_HANDLE_DBC, hdbc);
}
// Release an environment handle.
if (henv != SQL_NULL_HENV)
    SQLFreeHandle(SQL_HANDLE_ENV, henv);
return 0;
}
```

The running result is as follows:

Complete.

Part of database query result is as follows:

id	col
0	column test value 0
1	column test value 1
2	column test value 2
3	column test value 3
4	column test value 4
5	column test value 5
6	column test value 6
7	column test value 7
8	column test value 8
.....	

5.4.3.4 High-Performance Binding

If a large amount of data needs to be inserted, you are advised to perform the following operations:

- You need to set **UseBatchProtocol** to **1** in the **odbc.ini** file and **support_batch_bind** to **on** in the database.
- The ODBC program binding type must be the same as that in the database.
- The character set of the client is the same as that of the database.
- The transaction is committed manually.

odbc.ini configuration file:

```
[gaussdb]
Driver=GaussMPP
Servername=10.10.0.13 # Database server IP address
...
UseBatchProtocol=1 # Enabled by default
ConnSettings=set client_encoding=UTF8 # Set the character code on the client to be the same as that on the server.
```

The binding type example is as follows:

Prerequisite: The data source has been configured successfully. For Linux OS, see [Configuring a Data Source in the Linux OS](#). For Windows OS, see [Configuring a Data Source in the Windows OS](#).

```
#ifdef WIN32
#include <windows.h>
#endif
#include <stdio.h>
#include <stdlib.h>
#include <sql.h>
#include <sqlext.h>
#include <string.h>
#include <sys/time.h>

#define MESSAGE_BUFFER_LEN 128
SQLHANDLE h_env = NULL;
SQLHANDLE h_conn = NULL;
SQLHANDLE h_stmt = NULL;
void print_error()
{
    SQLCHAR Sqlstate[SQL_SQLSTATE_SIZE+1];
    SQLINTEGER NativeError;
    SQLCHAR MessageText[MESSAGE_BUFFER_LEN];
    SQLSMALLINT TextLength;
    SQLRETURN ret = SQL_ERROR;

    ret = SQLGetDiagRec(SQL_HANDLE_STMT, h_stmt, 1, Sqlstate, &NativeError, MessageText,
MESSAGE_BUFFER_LEN, &TextLength);
    if ( SQL_SUCCESS == ret)
    {
        printf("\n STMT ERROR-%05d %s", NativeError, MessageText);
        return;
    }

    ret = SQLGetDiagRec(SQL_HANDLE_DBC, h_conn, 1, Sqlstate, &NativeError, MessageText,
MESSAGE_BUFFER_LEN, &TextLength);
    if ( SQL_SUCCESS == ret)
    {
        printf("\n CONN ERROR-%05d %s", NativeError, MessageText);
        return;
    }

    ret = SQLGetDiagRec(SQL_HANDLE_ENV, h_env, 1, Sqlstate, &NativeError, MessageText,
```

```
MESSAGE_BUFFER_LEN, &TextLength);
    if ( SQL_SUCCESS == ret)
    {
        printf("\n ENV ERROR-%05d %s", NativeError, MessageText);
        return;
    }

    return;
}

/* Expect the function to return SQL_SUCCESS. */
#define RETURN_IF_NOT_SUCCESS(func) \
{\
    SQLRETURN ret_value = (func);\
    if (SQL_SUCCESS != ret_value)\
    {\
        print_error();\
        printf("\n failed line = %u: expect SQL_SUCCESS, but ret = %d", __LINE__, ret_value);\
        return SQL_ERROR; \
    }\
}

/* Expect the function to return SQL_SUCCESS. */
#define RETURN_IF_NOT_SUCCESS_I(i, func) \
{\
    SQLRETURN ret_value = (func);\
    if (SQL_SUCCESS != ret_value)\
    {\
        print_error();\
        printf("\n failed line = %u (i=%d): : expect SQL_SUCCESS, but ret = %d", __LINE__, (i), ret_value);\
        return SQL_ERROR; \
    }\
}

/* Expect the function to return SQL_SUCCESS_WITH_INFO. */
#define RETURN_IF_NOT_SUCCESS_INFO(func) \
{\
    SQLRETURN ret_value = (func);\
    if (SQL_SUCCESS_WITH_INFO != ret_value)\
    {\
        print_error();\
        printf("\n failed line = %u: expect SQL_SUCCESS_WITH_INFO, but ret = %d", __LINE__, ret_value);\
        return SQL_ERROR; \
    }\
}

/* Expect the values are the same. */
#define RETURN_IF_NOT(expect, value) \
if ((expect) != (value))\
{\
    printf("\n failed line = %u: expect = %u, but value = %u", __LINE__, (expect), (value)); \
    return SQL_ERROR;\
}

/* Expect the character strings are the same. */
#define RETURN_IF_NOT_STRCMP_I(i, expect, value) \
if (( NULL == (expect) ) || (NULL == (value)))\
{\
    printf("\n failed line = %u (i=%u): input NULL pointer !", __LINE__, (i)); \
    return SQL_ERROR; \
}\
else if (0 != strcmp((expect), (value)))\
{\
    printf("\n failed line = %u (i=%u): expect = %s, but value = %s", __LINE__, (i), (expect), (value)); \
    return SQL_ERROR;\
}

/* prepare + execute SQL statement */
```

```
int execute_cmd(SQLCHAR *sql)
{
    if ( NULL == sql )
    {
        return SQL_ERROR;
    }

    if ( SQL_SUCCESS != SQLPrepare(h_stmt, sql, SQL_NTS))
    {
        return SQL_ERROR;
    }

    if ( SQL_SUCCESS != SQLExecute(h_stmt))
    {
        return SQL_ERROR;
    }

    return SQL_SUCCESS;
}

/* execute + commit handle */
int commit_exec()
{
    if ( SQL_SUCCESS != SQLExecute(h_stmt))
    {
        return SQL_ERROR;
    }

    /* Manual committing */
    if ( SQL_SUCCESS != SQLEndTran(SQL_HANDLE_DBC, h_conn, SQL_COMMIT))
    {
        return SQL_ERROR;
    }

    return SQL_SUCCESS;
}

int begin_unit_test()
{
    SQLINTEGER    ret;

    /* Allocate an environment handle. */
    ret = SQLAllocHandle(SQL_HANDLE_ENV, SQL_NULL_HANDLE, &h_env);
    if ((SQL_SUCCESS != ret) && (SQL_SUCCESS_WITH_INFO != ret))
    {
        printf("\n begin_unit_test::SQLAllocHandle SQL_HANDLE_ENV failed ! ret = %d", ret);
        return SQL_ERROR;
    }

    /* Set the version number before connection. */
    if (SQL_SUCCESS != SQLSetEnvAttr(h_env, SQL_ATTR_ODBC_VERSION, (SQLPOINTER)SQL_OV_ODBC3,
0))
    {
        print_error();
        printf("\n begin_unit_test::SQLSetEnvAttr SQL_ATTR_ODBC_VERSION failed ! ret = %d", ret);
        SQLFreeHandle(SQL_HANDLE_ENV, h_env);
        return SQL_ERROR;
    }

    /* Allocate a connection handle. */
    ret = SQLAllocHandle(SQL_HANDLE_DBC, h_env, &h_conn);
    if (SQL_SUCCESS != ret)
    {
        print_error();
        printf("\n begin_unit_test::SQLAllocHandle SQL_HANDLE_DBC failed ! ret = %d", ret);
        SQLFreeHandle(SQL_HANDLE_ENV, h_env);
        return SQL_ERROR;
    }

    /* Establish a connection. */
```

```
ret = SQLConnect(h_conn, (SQLCHAR*) "gaussdb", SQL_NTS,
                (SQLCHAR*) NULL, 0, NULL, 0);
if (SQL_SUCCESS != ret)
{
    print_error();
    printf("\n begin_unit_test::SQLConnect failed ! ret = %d", ret);
    SQLFreeHandle(SQL_HANDLE_DBC, h_conn);
    SQLFreeHandle(SQL_HANDLE_ENV, h_env);
    return SQL_ERROR;
}

/* Allocate a statement handle. */
ret = SQLAllocHandle(SQL_HANDLE_STMT, h_conn, &h_stmt);
if (SQL_SUCCESS != ret)
{
    print_error();
    printf("\n begin_unit_test::SQLAllocHandle SQL_HANDLE_STMT failed ! ret = %d", ret);
    SQLFreeHandle(SQL_HANDLE_DBC, h_conn);
    SQLFreeHandle(SQL_HANDLE_ENV, h_env);
    return SQL_ERROR;
}

return SQL_SUCCESS;
}

void end_unit_test()
{
    /* Release a statement handle. */
    if (NULL != h_stmt)
    {
        SQLFreeHandle(SQL_HANDLE_STMT, h_stmt);
    }

    /* Release a connection handle. */
    if (NULL != h_conn)
    {
        SQLDisconnect(h_conn);
        SQLFreeHandle(SQL_HANDLE_DBC, h_conn);
    }

    /* Release an environment handle. */
    if (NULL != h_env)
    {
        SQLFreeHandle(SQL_HANDLE_ENV, h_env);
    }

    return;
}

int main()
{
    /* begin test */
    if (begin_unit_test() != SQL_SUCCESS)
    {
        printf("\n begin_test_unit failed.");
        return SQL_ERROR;
    }

    /* The handle configuration is the same as that in the preceding case. */
    int i = 0;
    SQLCHAR* sql_drop = "drop table if exists test_bindnumber_001";
    SQLCHAR* sql_create = "create table test_bindnumber_001("
        "f4 number, f5 number(10, 2)"
        ")";
    SQLCHAR* sql_insert = "insert into test_bindnumber_001 values(?, ?)";
    SQLCHAR* sql_select = "select * from test_bindnumber_001";
    SQLLEN RowCount;
    SQL_NUMERIC_STRUCT st_number;
    SQLCHAR getValue[2][MESSAGE_BUFFER_LEN];
```

```
/* Step 1. Create a table. */
RETURN_IF_NOT_SUCCESS(execute_cmd(sql_drop));
RETURN_IF_NOT_SUCCESS(execute_cmd(sql_create));

/* Step 2.1 Bind parameters using the SQL_NUMERIC_STRUCT structure. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));

/* First line: 1234.5678 */
memset(st_number.val, 0, SQL_MAX_NUMERIC_LEN);
st_number.precision = 8;
st_number.scale = 4;
st_number.sign = 1;
st_number.val[0] = 0x4E;
st_number.val[1] = 0x61;
st_number.val[2] = 0xBC;

RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 4, &st_number, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 4, &st_number, 0, NULL));

/* Disable the automatic commit function. */
SQLSetConnectAttr(h_conn, SQL_ATTR_AUTOCOMMIT, (SQLPOINTER)SQL_AUTOCOMMIT_OFF, 0);

RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Second line: 12345678 */
memset(st_number.val, 0, SQL_MAX_NUMERIC_LEN);
st_number.precision = 8;
st_number.scale = 0;
st_number.sign = 1;
st_number.val[0] = 0x4E;
st_number.val[1] = 0x61;
st_number.val[2] = 0xBC;

RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 0, &st_number, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 0, &st_number, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Third line: 12345678 */
memset(st_number.val, 0, SQL_MAX_NUMERIC_LEN);
st_number.precision = 0;
st_number.scale = 4;
st_number.sign = 1;
st_number.val[0] = 0x4E;
st_number.val[1] = 0x61;
st_number.val[2] = 0xBC;

RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 4, &st_number, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_NUMERIC,
SQL_NUMERIC, sizeof(SQL_NUMERIC_STRUCT), 4, &st_number, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.2 Bind parameters by using the SQL_C_CHAR character string in the fourth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
SQLCHAR* szNumber = "1234.5678";
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_CHAR,
SQL_NUMERIC, strlen(szNumber), 0, szNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_CHAR,
```

```
SQL_NUMERIC, strlen(szNumber), 0, szNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.3 Bind parameters by using SQL_C_FLOAT in the fifth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
SQLREAL fNumber = 1234.5678;
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_FLOAT,
SQL_NUMERIC, sizeof(fNumber), 4, &fNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_FLOAT,
SQL_NUMERIC, sizeof(fNumber), 4, &fNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.4 Bind parameters by using SQL_C_DOUBLE in the sixth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
SQLDOUBLE dNumber = 1234.5678;
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_DOUBLE,
SQL_NUMERIC, sizeof(dNumber), 4, &dNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_DOUBLE,
SQL_NUMERIC, sizeof(dNumber), 4, &dNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

SQLBIGINT bNumber1 = 0xFFFFFFFFFFFFFFFF;
SQLBIGINT bNumber2 = 12345;

/* Step 2.5 Bind parameters by using SQL_C_SBIGINT in the seventh line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_SBIGINT,
SQL_NUMERIC, sizeof(bNumber1), 4, &bNumber1, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_SBIGINT,
SQL_NUMERIC, sizeof(bNumber2), 4, &bNumber2, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.6 Bind parameters by using SQL_C_UBIGINT in the eighth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_UBIGINT,
SQL_NUMERIC, sizeof(bNumber1), 4, &bNumber1, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_UBIGINT,
SQL_NUMERIC, sizeof(bNumber2), 4, &bNumber2, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

SQLLEN lNumber1 = 0xFFFFFFFFFFFFFFFF;
SQLLEN lNumber2 = 12345;

/* Step 2.7 Bind parameters by using SQL_C_LONG in the ninth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_LONG,
SQL_NUMERIC, sizeof(lNumber1), 0, &lNumber1, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_LONG,
SQL_NUMERIC, sizeof(lNumber2), 0, &lNumber2, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.8 Bind parameters by using SQL_C_ULONG in the tenth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_ULONG,
SQL_NUMERIC, sizeof(lNumber1), 0, &lNumber1, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_ULONG,
SQL_NUMERIC, sizeof(lNumber2), 0, &lNumber2, 0, NULL));
```

```
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

SQLSMALLINT sNumber = 0xFFFF;

/* Step 2.9 Bind parameters by using SQL_C_SHORT in the eleventh line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_SHORT,
SQL_NUMERIC, sizeof(sNumber), 0, &sNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_SHORT,
SQL_NUMERIC, sizeof(sNumber), 0, &sNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.10 Bind parameters by using SQL_C_USHORT in the twelfth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_USHORT,
SQL_NUMERIC, sizeof(sNumber), 0, &sNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_USHORT,
SQL_NUMERIC, sizeof(sNumber), 0, &sNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

SQLCHAR cNumber = 0xFF;

/* Step 2.11 Bind parameters by using SQL_C_TINYINT in the thirteenth line. */
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_TINYINT,
SQL_NUMERIC, sizeof(cNumber), 0, &cNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_TINYINT,
SQL_NUMERIC, sizeof(cNumber), 0, &cNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Step 2.12 Bind parameters by using SQL_C_UTINYINT in the fourteenth line.*/
RETURN_IF_NOT_SUCCESS(SQLPrepare(h_stmt, sql_insert, SQL_NTS));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 1, SQL_PARAM_INPUT, SQL_C_UTINYINT,
SQL_NUMERIC, sizeof(cNumber), 0, &cNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(SQLBindParameter(h_stmt, 2, SQL_PARAM_INPUT, SQL_C_UTINYINT,
SQL_NUMERIC, sizeof(cNumber), 0, &cNumber, 0, NULL));
RETURN_IF_NOT_SUCCESS(commit_exec());
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));
RETURN_IF_NOT(1, RowCount);

/* Use the character string type to unify the expectation. */
SQLCHAR* expectValue[14][2] = {{ "1234.5678", "1234.57"},
    { "12345678", "12345678"},
    { "0", "0"},
    { "1234.5678", "1234.57"},
    { "1234.5677", "1234.57"},
    { "1234.5678", "1234.57"},
    { "-1", "12345"},
    { "18446744073709551615", "12345"},
    { "-1", "12345"},
    { "4294967295", "12345"},
    { "-1", "-1"},
    { "65535", "65535"},
    { "-1", "-1"},
    { "255", "255"},
    };

RETURN_IF_NOT_SUCCESS(execute_cmd(sql_select));
while ( SQL_NO_DATA != SQLFetch(h_stmt))
{
    RETURN_IF_NOT_SUCCESS_I(i, SQLGetData(h_stmt, 1, SQL_C_CHAR, &getValue[0],
```



```
MESSAGE_BUFFER_LEN, NULL));  
    RETURN_IF_NOT_SUCCESS_I(i, SQLGetData(h_stmt, 2, SQL_C_CHAR, &getValue[1],  
MESSAGE_BUFFER_LEN, NULL));  
    i++;  
}  
printf("\nComplete!\n");  
RETURN_IF_NOT_SUCCESS(SQLRowCount(h_stmt, &RowCount));  
RETURN_IF_NOT(i, RowCount);  
SQLCloseCursor(h_stmt);  
/* Final step. Delete the table and restore the environment. */  
RETURN_IF_NOT_SUCCESS(execute_cmd(sql_drop));  
RETURN_IF_NOT_SUCCESS(commit_exec());  
end_unit_test();  
}
```

NOTE

In the preceding example, the **number** column is defined. When the SQLBindParameter API is called, the performance of binding SQL_NUMERIC is higher than that of SQL_LONG. If char is used, the data type needs to be converted when data is inserted to the database server, causing a performance bottleneck.

The running result is as follows:

```
Complete!
```

5.4.4 ODBC Interface Reference

The ODBC interface is a set of API functions provided to users. This chapter describes its common interfaces. For details on other interfaces, see "ODBC Programmer's Reference" at MSDN ([https://msdn.microsoft.com/en-us/library/windows/desktop/ms714177\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/desktop/ms714177(v=vs.85).aspx)).

5.4.4.1 SQLAllocEnv

In ODBC 3.x, SQLAllocEnv (an ODBC 2.x function) was deprecated and replaced by SQLAllocHandle. For details, see [SQLAllocHandle](#).

5.4.4.2 SQLAllocConnect

In ODBC 3.x, SQLAllocConnect (a function in ODBC 2.x) was deprecated and replaced by SQLAllocHandle. For details, see [SQLAllocHandle](#).

5.4.4.3 SQLAllocHandle

Description

Allocates environment, connection, statement, or descriptor handles. This function replaces the deprecated ODBC 2.x functions SQLAllocEnv, SQLAllocConnect, and SQLAllocStmt.

Prototype

```
SQLRETURN SQLAllocHandle(SQLSMALLINT HandleType,  
                          SQLHANDLE InputHandle,  
                          SQLHANDLE *OutputHandlePtr);
```

Parameter

Table 5-37 SQLAllocHandle parameters

Keyword	Description
HandleType	Type of handle to be allocated by SQLAllocHandle. The value must be one of the following: • SQL_HANDLE_ENV (environment handle) • SQL_HANDLE_DBC (connection handle) • SQL_HANDLE_STMT (statement handle) • SQL_HANDLE_DESC (descriptor handle) The handle application sequence is: SQL_HANDLE_ENV > SQL_HANDLE_DBC > SQL_HANDLE_STMT . The handle applied later depends on the handle applied prior to it.
InputHandle	Existing handle to use as a context for the new handle being allocated. • If HandleType is set to SQL_HANDLE_ENV, this parameter is set to SQL_NULL_HANDLE. • If HandleType is set to SQL_HANDLE_DBC, this parameter value must be an environment handle. • If HandleType is set to SQL_HANDLE_STMT or SQL_HANDLE_DESC, this parameter value must be a connection handle.
OutputHandlePtr	Output parameter: Pointer to a buffer that stores the returned handle in the newly allocated data structure.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLAllocHandle returns **SQL_ERROR** when it is used to allocate a non-environment handle, it sets **OutputHandlePtr** to **SQL_NULL_HDBC**, **SQL_NULL_HSTMT**, or **SQL_NULL_HDESC**. The application can then call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to the value of **InputHandle**, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.4 SQLAllocStmt

In ODBC 3.x, SQLAllocStmt (a function in ODBC 2.x) was deprecated and replaced by SQLAllocHandle. For details, see [SQLAllocHandle](#).

5.4.4.5 SQLBindCol

Description

Binds application data buffers to columns in a result set.

Prototype

```
SQLRETURN SQLBindCol(SQLHSTMT StatementHandle,  
SQLUSMALLINT ColumnNumber,  
SQLSMALLINT TargetType,  
SQLPOINTER TargetValuePtr,  
SQLLEN BufferLength,  
SQLLEN *StrLen_or_IndPtr);
```

Parameters

Table 5-38 SQLBindCol parameters

Keyword	Description
StatementHandle	Statement handle.
ColumnNumber	Number of the column to be bound. The column number starts with 0 and increases in ascending order. Column 0 is the bookmark column. If no bookmark column is set, column numbers start with 1.
TargetType	C data type in the buffer.
TargetValuePtr	Output parameter: pointer to the buffer bound with the column. The SQLFetch function returns data in the buffer. If this parameter is a null pointer, StrLen_or_IndPtr is a valid value.
BufferLength	Length of the TargetValuePtr buffer in bytes.
StrLen_or_IndPtr	Output parameter: pointer to the length or indicator of the buffer. If the value is null, no length or indicator is used.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.

- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLBindCol returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.6 SQLBindParameter

Description

Binds parameter markers in an SQL statement to a buffer.

Prototype

```
SQLRETURN SQLBindParameter(SQLHSTMT StatementHandle,  
    SQLUSMALLINT ParameterNumber,  
    SQLSMALLINT InputOutputType,  
    SQLSMALLINT ValueType,  
    SQLSMALLINT ParameterType,  
    SQLULEN ColumnSize,  
    SQLSMALLINT DecimalDigits,  
    SQLPOINTER ParameterValuePtr,  
    SQLLEN BufferLength,  
    SQLLEN *StrLen_or_IndPtr);
```

Parameters

Table 5-39 SQLBindParameter

Keyword	Description
StatementHandle	Statement handle.
ParameterNumber	Parameter marker number, starting with 1 and increasing in ascending order.
InputOutputType	Input/output type of the parameter.
ValueType	C data type of the parameter.
ParameterType	SQL data type of the parameter.

Keyword	Description
ColumnSize	Size of the column or expression of the corresponding parameter marker.
DecimalDigits	Decimal digit of the column or the expression of the corresponding parameter marker.
ParameterValuePtr	Pointer to the storage parameter buffer.
BufferLength	Length of the ParameterValuePtr buffer in bytes.
StrLen_or_IndPtr	Pointer to the length or indicator of the buffer. If the value is null, no length or indicator is used.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If `SQLBindParameter` returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.7 SQLColAttribute

Description

Returns the descriptor information about a column in the result set.

Prototype

```
SQLRETURN SQLColAttribute(SQLHSTMT StatementHandle,
                           SQLUSMALLINT ColumnNumber,
                           SQLUSMALLINT FieldIdentifier,
                           SQLPOINTER CharacterAttributePtr,
                           SQLSMALLINT BufferLength,
                           SQLSMALLINT *StringLengthPtr,
                           SQLLEN *NumericAttributePtr);
```

Parameters

Table 5-40 SQLColAttribute parameters

Keyword	Description
StatementHandle	Statement handle.
ColumnNumber	Column number of the field to be queried, starting with 1 and increasing in ascending order.
FieldIdentifier	Field identifier of ColumnNumber in IRD.
CharacterAttributePtr	Output parameter: pointer to the buffer that returns the FieldIdentifier value.
BufferLength	- This parameter indicates the length of the buffer if FieldIdentifier is an ODBC-defined field and CharacterAttributePtr points to a string or a binary buffer. - Ignore this parameter if FieldIdentifier is an ODBC-defined field and CharacterAttributePtr points to an integer.
StringLengthPtr	Output parameter: pointer to a buffer in which the total number of valid bytes (for string data) is stored in *CharacterAttributePtr . Ignore the value of BufferLength if the data is not a string.
NumericAttributePtr	Output parameter: pointer to an integer buffer in which the value of FieldIdentifier in the ColumnNumber row of the IRD is returned.

Return Value

- SQL_SUCCESS** indicates that the call succeeded.
- SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLColAttribute returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.8 SQLConnect

Description

Establishes a connection between a driver and a data source. After the connection is established, the connection handle can be used to access all information about the data source, including its application operating status, transaction processing status, and error information.

Prototype

```
SQLRETURN SQLConnect(SQLHDBC ConnectionHandle,  
    SQLCHAR *ServerName,  
    SQLSMALLINT NameLength1,  
    SQLCHAR *UserName,  
    SQLSMALLINT NameLength2,  
    SQLCHAR *Authentication,  
    SQLSMALLINT NameLength3);
```

Parameters

Table 5-41 SQLConnect parameters

Keyword	Description
ConnectionHandle	Connection handle, obtained from SQLAllocHandle.
ServerName	Name of the data source to connect.
NameLength1	Length of ServerName .
UserName	Username of the database in the data source.
NameLength2	Length of UserName .
Authentication	User password of the database in the data source.
NameLength3	Length of Authentication .

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

- **SQL_STILL_EXECUTING** indicates that the statement is being executed.

Precautions

If SQLConnect returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_DBC** and **ConnectionHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.9 SQLDisconnect

Description

Closes the connection associated with a database connection handle.

Prototype

```
SQLRETURN SQLDisconnect(SQLHDBC ConnectionHandle);
```

Parameter

Table 5-42 SQLDisconnect parameters

Keyword	Description
ConnectionHandle	Connection handle, obtained from SQLAllocHandle.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLDisconnect returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_DBC** and **ConnectionHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.10 SQLExecDirect

Description

Executes a prepared statement specified in this parameter. SQLExecDirect is the fastest method for executing only one SQL statement at a time.

Prototype

```
SQLRETURN SQLExecDirect(SQLHSTMT StatementHandle,  
                        SQLCHAR *StatementText,  
                        SQLINTEGER TextLength);
```

Parameter

Table 5-43 SQLExecDirect parameters

Keyword	Description
StatementHandle	Statement handle, obtained from SQLAllocHandle.
StatementText	SQL statement to be executed. Multiple statements cannot be executed at a time.
TextLength	Length of StatementText .

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_NEED_DATA** indicates that parameters provided before executing the SQL statement are insufficient.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.
- **SQL_STILL_EXECUTING** indicates that the statement is being executed.
- **SQL_NO_DATA** indicates that the SQL statement does not return a result set.

Precautions

If SQLExecDirect returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the

SQLSTATE value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.11 SQLExecute

Description

If a statement contains a parameter marker, the SQLExecute function uses the current value of the parameter marker to execute a prepared SQL statement.

Prototype

```
SQLRETURN SQLExecute(SQLHSTMT StatementHandle);
```

Parameter

Table 5-44 SQLExecute parameter

Keyword	Description
StatementHandle	Statement handle to be executed.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_NEED_DATA** indicates that parameters provided before executing the SQL statement are insufficient.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_NO_DATA** indicates that the SQL statement does not return a result set.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.
- **SQL_STILL_EXECUTING** indicates that the statement is being executed.

Precautions

If SQLExecute returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.12 SQLFetch

Description

Advances the cursor to the next row of the result set and retrieves any bound columns.

Prototype

```
SQLRETURN SQLFetch(SQLHSTMT StatementHandle);
```

Parameter

Table 5-45 SQLFetch parameters

Keyword	Description
StatementHandle	Statement handle, obtained from SQLAllocHandle.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_NO_DATA** indicates that the SQL statement does not return a result set.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.
- **SQL_STILL_EXECUTING** indicates that the statement is being executed.

Precautions

If SQLFetch returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.13 SQLFreeStmt

In ODBC 3.x, SQLFreeStmt (a function in ODBC 2.x) was deprecated and replaced with SQLFreeHandle. For details, see [SQLFreeHandle](#).

5.4.4.14 SQLFreeConnect

In ODBC 3.x, SQLFreeConnect (an ODBC 2.x function) was deprecated and replaced by SQLFreeHandle. For details, see [SQLFreeHandle](#).

5.4.4.15 SQLFreeHandle

Description

SQLFreeHandle is used to release resources associated with a specific environment, connection, or statement handle. It replaces the ODBC 2.x functions: SQLFreeEnv, SQLFreeConnect, and SQLFreeStmt.

Prototype

```
SQLRETURN SQLFreeHandle(SQLSMALLINT HandleType,  
                        SQLHANDLE Handle);
```

Parameter

Table 5-46 SQLFreeHandle parameters

Keyword	Description
HandleType	Type of handle to be freed by SQLFreeHandle. The value must be one of the following: • SQL_HANDLE_ENV • SQL_HANDLE_DBC • SQL_HANDLE_STMT • SQL_HANDLE_DESC If HandleType is not one of the preceding values, SQLFreeHandle returns SQL_INVALID_HANDLE .
Handle	Name of the handle to be freed.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLFreeHandle returns **SQL_ERROR**, the handle is still valid.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.16 SQLFreeEnv

In ODBC 3.x, SQLFreeEnv (an ODBC 2.x function) was deprecated and replaced by SQLFreeHandle. For details, see [SQLFreeHandle](#).

5.4.4.17 SQLPrepare

Description

Prepares an SQL statement to be executed.

Note that ODBC does not support the kernel reuse plan when sending prepared statements. As a result, a new plan needs to be generated for each execution, causing high CPU usage. If services have requirements on plan reuse, you are advised to use the JDBC client.

Prototype

```
SQLRETURN SQLPrepare(SQLHSTMT StatementHandle,  
                      SQLCHAR *StatementText,  
                      SQLINTEGER TextLength);
```

Parameter

Table 5-47 SQLPrepare parameters

Keyword	Description
StatementHandle	Statement handle.
StatementText	SQL text string.
TextLength	Length of StatementText .

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

- **SQL_STILL_EXECUTING** indicates that the statement is being executed.

Precautions

If SQLPrepare returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.18 SQLGetData

Description

SQLGetData is used to retrieve data for a single column in the result set. It can be called for many times to retrieve data of variable lengths.

Prototype

```
SQLRETURN SQLGetData(SQLHSTMT StatementHandle,  
SQLUSMALLINT Col_or_Param_Num,  
SQLSMALLINT TargetType,  
SQLPOINTER TargetValuePtr,  
SQLLEN BufferLength,  
SQLLEN *StrLen_or_IndPtr);
```

Parameter

Table 5-48 SQLGetData parameters

Keyword	Description
StatementHandle	Statement handle, obtained from SQLAllocHandle.
Col_or_Param_Num	Column number for which the data retrieval is requested. The column number starts with 1 and increases in ascending order. The number of the bookmark column is 0.
TargetType	C data type in the TargetValuePtr buffer. If TargetType is SQL_ARD_TYPE , the driver uses the data type of the SQL_DESC_CONCISE_TYPE field in ARD. If it is SQL_C_DEFAULT , the driver selects a default data type according to the source SQL data type.
TargetValuePtr	Output parameter: pointer to the pointer that points to the buffer where the data is located.
BufferLength	Size of the buffer pointed to by TargetValuePtr .

Keyword	Description
StrLen_or_IndPtr	Output parameter: pointer to the buffer where the length or identifier value is returned.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_NO_DATA** indicates that the SQL statement does not return a result set.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.
- **SQL_STILL_EXECUTING** indicates that the statement is being executed.

Precautions

If `SQLGetData` returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_STMT** and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.19 SQLGetDiagRec

Description

Returns the current values of multiple columns in a diagnostic record that contains error, warning, and status information.

Prototype

```
SQLRETURN SQLGetDiagRec(SQLSMALLINT  HandleType,  
                        SQLHANDLE      Handle,  
                        SQLSMALLINT     RecNumber,  
                        SQLCHAR         *SQLState,  
                        SQLINTEGER       *NativeErrorPtr,  
                        SQLCHAR         *MessageText,  
                        SQLSMALLINT     BufferLength  
                        SQLSMALLINT     *TextLengthPtr);
```

Parameters

Table 5-49 SQLGetDiagRec parameters

Keyword	Description
HandleType	Handle-type identifier that describes the type of handle for which diagnostics are desired. The value must be one of the following: • SQL_HANDLE_ENV • SQL_HANDLE_DBC • SQL_HANDLE_STMT • SQL_HANDLE_DESC
Handle	Handle for the diagnostic data structure. Its type is indicated by HandleType . If HandleType is set to SQL_HANDLE_ENV , Handle may indicate a shared or non-shared environment handle.
RecNumber	Status record from which the application seeks information. Status records are numbered from 1.
SQLState	Output parameter: pointer to a buffer that saves the 5-character SQLSTATE code pertaining to RecNumber .
NativeErrorPtr	Output parameter: pointer to a buffer that saves the native error code.
MessageText	Pointer to a buffer that saves text strings of diagnostic information.
BufferLength	Length of MessageText .
TextLengthPtr	Output parameter: pointer to the buffer, the total number of bytes in the returned MessageText . If the number of bytes available to return is greater than BufferLength , then the diagnostics information text in MessageText is truncated to BufferLength minus the length of the null termination character.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

SQLGetDiagRec does not release diagnostic records for itself. It uses the following return values to report execution results:

- **SQL_SUCCESS** indicates that the function successfully returns diagnostic information.
- **SQL_SUCCESS_WITH_INFO** indicates that the **MessageText** buffer is too small to hold the requested diagnostic information. No diagnostic records are generated.
- **SQL_INVALID_HANDLE** indicates that the handle indicated by **HandType** and **Handle** is an invalid handle.
- **SQL_ERROR** indicates that the value of **RecNumber** is less than or equal to 0 or the value of **BufferLength** is less than 0.

If an ODBC function returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call SQLGetDiagRec to obtain the **SQLSTATE** value. The possible **SQLSTATE** values are listed as follows:

Table 5-50 SQLSTATE values

SQLSATATE	Error	Description
HY000	General error.	An error occurred for which there is no specific SQLSTATE.
HY001	Memory allocation error.	The driver is unable to allocate memory required to support execution or completion of the function.
HY008	Operation canceled.	SQLCancel is called to terminate the statement execution, but the StatementHandle function is still called.
HY010	Function sequence error.	The function is called prior to sending data to data parameters or columns being executed.
HY013	Memory management error.	The function fails to be called. The error may be caused by low memory conditions.
HYT01	Connection timeout.	The timeout period expired before the application was able to connect to the data source.
IM001	Function not supported by the driver.	The called function is not supported by the StatementHandle driver.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.20 SQLSetConnectAttr

Description

Sets connection attributes.

Prototype

```
SQLRETURN SQLSetConnectAttr(SQLHDBC ConnectionHandle,  
                             SQLINTEGER Attribute,  
                             SQLPOINTER ValuePtr,  
                             SQLINTEGER StringLength);
```

Parameters

Table 5-51 SQLSetConnectAttr parameters

Keyword	Description
ConnectionHandle	Connection handle.
Attribute	Attribute to set.
ValuePtr	Pointer to the Attribute value. ValuePtr depends on the Attribute value, and can be a 32-bit unsigned integer value or a null-terminated string. If the ValuePtr parameter is a driver-specific value, it may be a signed integer.
StringLength	If ValuePtr points to a string or a binary buffer, StringLength is the length of *ValuePtr . If ValuePtr points to an integer, StringLength is ignored.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLSetConnectAttr returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to **SQL_HANDLE_DBC** and **ConnectionHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.21 SQLSetEnvAttr

Description

SQLSetEnvAttr is used to set environment attributes.

Prototype

```
SQLRETURN SQLSetEnvAttr(SQLHENV EnvironmentHandle,  
                        SQLINTEGER Attribute,  
                        SQLPOINTER ValuePtr,  
                        SQLINTEGER StringLength);
```

Parameter

Table 5-52 SQLSetEnvAttr parameters

Keyword	Description
EnvironmentHandle	Environment handle.
Attribute	Environment attribute to be set. The value must be one of the following: • SQL_ATTR_ODBC_VERSION: ODBC version • SQL_CONNECTION_POOLING: connection pool attribute • SQL_OUTPUT_NTS: string type returned by the driver
ValuePtr	Pointer to the Attribute value. ValuePtr depends on the Attribute value, and can be a 32-bit integer value or a null-terminated string.
StringLength	If ValuePtr points to a string or a binary buffer, StringLength is the length of *ValuePtr . If ValuePtr points to an integer, StringLength is ignored.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLSetEnvAttr returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), set **HandleType** and **Handle** to **SQL_HANDLE_ENV** and **EnvironmentHandle**, and obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples

For details, see [Typical Application Development Examples](#).

5.4.4.22 SQLSetStmtAttr

Description

Sets attributes related to a statement.

Prototype

```
SQLRETURN SQLSetStmtAttr(SQLHSTMT StatementHandle,  
                           SQLINTEGER Attribute,  
                           SQLPOINTER ValuePtr,  
                           SQLINTEGER StringLength);
```

Parameters

Table 5-53 SQLSetStmtAttr parameters

Keyword	Description
StatementHandle	Statement handle.
Attribute	Attribute to set.
ValuePtr	Pointer to the Attribute value. ValuePtr depends on the Attribute value, and can be a 32-bit unsigned integer value or a pointer to a null-terminated string, a binary buffer, or a driver-specified value. If the ValuePtr parameter is a driver-specific value, it may be a signed integer.
StringLength	If ValuePtr points to a string or a binary buffer, StringLength is the length of *ValuePtr . If ValuePtr points to an integer, StringLength is ignored.

Return Value

- **SQL_SUCCESS** indicates that the call succeeded.
- **SQL_SUCCESS_WITH_INFO** indicates that some warning information is displayed.
- **SQL_ERROR** indicates major errors, such as memory allocation and connection failures.
- **SQL_INVALID_HANDLE** indicates that invalid handles were called. This value may also be returned by other APIs.

Precautions

If SQLSetStmtAttr returns **SQL_ERROR** or **SQL_SUCCESS_WITH_INFO**, the application can call [SQLGetDiagRec](#), with **HandleType** and **Handle** set to

SQL_HANDLE_STMT and **StatementHandle**, respectively, to obtain the **SQLSTATE** value. The **SQLSTATE** value provides the detailed function calling information.

Examples


For details, see [Typical Application Development Examples](#).

5.5 Development Based on libpq

5.5.1 Development Process

[Figure 5-4](#) shows the libpq-based development process.

Figure 5-4 libpq-based development process



5.5.2 Development Procedure

5.5.2.1 Obtaining a Driver Package, a Dependency Library, and Header Files

Obtain the library and header files of libpq from **GaussDB-Kernel-Database version number-OS information-64bit-Libpq.tar.gz** in the driver package. The

required header file is stored in the **include** folder, and the **lib** folder contains the required libpq library file. Programs that use libpq must include the header file **libpq-fe.h** and must connect to the libpq library.

NOTE

In addition to **libpq-fe.h**, the **include** folder contains the header files **postgres_ext.h**, **gs_thread.h**, and **gs_threadlocal.h** by default. These three header files are the dependency files of **libpq-fe.h**.

5.5.2.2 Connecting to a Database

Connecting to a database is the first step in developing an application using libpq. In this case, you can use functions, such as **PQconnectdb** and **PQsetdbLogin**, to connect to the database server. These functions return a connection object. You need to save the connection object for subsequent database operations.

The following uses the development source program **testlibpq.c** as an example (for details about the complete example, see [Establishing a Database Connection, Executing SQL Statements, and Returning Results](#)).

```
/*
 * Note: The source program testlibpq.c provides basic and common application scenarios of libpq.
 * The PQconnectdb, PQexec, PQntuples, and PQfinish APIs provided by libpq are used to establish database
 * connections, execute SQL statements, obtain returned results, and clear resources.
 * Header file. For details about user-defined functions, see the complete example.
 */

/* The values of variables such as user and passwd must be read from environment variables or
 * configuration files. Environment variables need to be configured as required. If no environment variable is
 * used, a character string can be directly assigned. */
const char conninfo[1024];
PGconn *conn;
PGresult *res;
int nFields;
int i,j;
char *passwd = getenv("EXAMPLE_PASSWD_ENV");
char *port = getenv("EXAMPLE_PORT_ENV");
char *host = getenv("EXAMPLE_HOST_ENV");
char *username = getenv("EXAMPLE_USERNAME_ENV");
char *dbname = getenv("EXAMPLE_DBNAME_ENV");

/*
 * This value is used when the user provides the value of the conninfo character string in the command line.
 * Otherwise, the environment variables or all other connection parameters
 * use the default values.
 */
if (argc > 1)
    strcpy(conninfo, argv[1]);
else
    sprintf(conninfo,
        "dbname=%s port=%s host=%s application_name=test connect_timeout=5 sslmode=allow user=%s"
        "password=%s",
        dbname, port, host, username, passwd);

/* Connect to the database. */
conn = PQconnectdb(conninfo);

/* Check whether the backend connection has been successfully established. */
if (PQstatus(conn) != CONNECTION_OK)
{
    fprintf(stderr, "Connection to database failed: %s",
        PQerrorMessage(conn));
    exit_nicely(conn);
}
```

5.5.2.3 Executing SQL Statements

You can use the function **PQexec** to execute SQL query statements through the connection object. This includes querying data (SELECT), inserting data, updating data, and deleting data. If multiple SQL statements are executed concurrently as a transaction, use the transaction control function. For example, run SQL statements such as BEGIN, COMMIT, and ROLLBACK to control the start, committing, and rollback of a transaction. Handle the errors after executing the SQL query statements.

The following is an example (for details about the complete example, see [Establishing a Database Connection, Executing SQL Statements, and Returning Results](#)).

```
/*
 * After the connection is successful
 * Since a cursor is used in the test case, a transaction block is required.
 * Put all data in one "select * from pg_database".
 * PQexec() is too simple and is not recommended.
 */

/* Start a transaction block. */
res = PQexec(conn, "BEGIN");
if (PQresultStatus(res) != PGRES_COMMAND_OK)
{
    fprintf(stderr, "BEGIN command failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}

/*
 * PQclear PGresult should be executed when it is no longer needed, to avoid memory leakage.
 */
PQclear(res);

/*
 * Fetch data from the pg_database system catalog.
 */
res = PQexec(conn, "DECLARE myportal CURSOR FOR select * from pg_database");
if (PQresultStatus(res) != PGRES_COMMAND_OK)
{
    fprintf(stderr, "DECLARE CURSOR failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}
PQclear(res);

res = PQexec(conn, "FETCH ALL in myportal");
if (PQresultStatus(res) != PGRES_TUPLES_OK)
{
    fprintf(stderr, "FETCH ALL failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}
```

5.5.2.4 Processing Data in a Result Set

libpq provides functions, such as **PQntuples**, **PQnfields**, and **PQfname**, to help you properly parse and process the results of SELECT queries.

The following is an example (for details about the complete example, see [Establishing a Database Connection, Executing SQL Statements, and Returning Results](#)).

```
/* First, print out the attribute name. */
nFields = PQnfields(res);
for (i = 0; i < nFields; i++)
    printf("%-15s", PQfname(res, i));
printf("\n\n");

/* Print lines. */
for (i = 0; i < PQntuples(res); i++)
{
    for (j = 0; j < nFields; j++)
        printf("%-15s", PQgetvalue(res, i, j));
    printf("\n");
}

/* Release the memory of the result object to avoid memory leakage. */
PQclear(res);
```

5.5.2.5 Closing a Connection

libpq usually uses the function **PQfinish** to close the connection to the database. If your application has established multiple connections, make sure that each connection is closed correctly.

The following is an example (for details about the complete example, see [Establishing a Database Connection, Executing SQL Statements, and Returning Results](#)).

```
/* Close the portal. We do not need to check for errors. */
res = PQexec(conn, "CLOSE myportal");
PQclear(res);

/* End the transaction. */
res = PQexec(conn, "END");
PQclear(res);

/* Close the database connection and clean up the database. */
PQfinish(conn);
```

5.5.3 Typical Application Development Examples

5.5.3.1 Establishing a Database Connection, Executing SQL Statements, and Returning Results

NOTE

To compile the libpq source program by running **gcc**, use the **-I** *directory* option to provide the installation location of header files. (Sometimes the compiler looks for the default directory, so this option can be ignored.) Example:

```
gcc -I (Directory where header files are located) -L (Directory where the libpq library is located) -o
testlibpq testlibpq.c -lpq
```

Run the following command:

```
./testlibpq.c
```

If a makefile is used, add the following options to variables **CPPFLAGS**, **LDFLAGS**, and **LIBS**:

```
CPPFLAGS += -I (Directory of header files)
```

```
LDFLAGS += -L (Directory of the libpq library)
```

```
LIBS += -lpq
```

Example:

```
CPPFLAGS += -I$(GAUSSHOME)/include/libpq
```

```
LDFLAGS += -L$(GAUSSHOME)/lib
```

```
/*
 * testlibpq.c
```



```
* Note: testlibpq.c source program provides basic and common application scenarios of libpq.
* The PQconnectdb, PQexec, PQntuples, and PQfinish APIs provided by libpq are used to establish database
connections, execute SQL statements, obtain returned results, and clear resources.
*/
#include <stdio.h>
#include <stdlib.h>
#include <libpq-fe.h>
#include <string.h>

static void
exit_nicely(PGconn *conn)
{
    PQfinish(conn);
    exit(1);
}

int main(int argc, char **argv)
{
    /* The values of variables such as user and passwd must be read from environment variables or
configuration files. Environment variables need to be configured as required. If no environment variable is
used, a character string can be directly assigned. */
    const char conninfo[1024];
    PGconn *conn;
    PGresult *res;
    int nFields;
    int i,j;
    char *passwd = getenv("EXAMPLE_PASSWD_ENV");
    char *port = getenv("EXAMPLE_PORT_ENV");
    char *host = getenv("EXAMPLE_HOST_ENV");
    char *username = getenv("EXAMPLE_USERNAME_ENV");
    char *dbname = getenv("EXAMPLE_DBNAME_ENV");

    /*
    * This value is used when the user provides the value of the conninfo character string in the command
line.
    * Otherwise, the environment variables or the default values
    * are used for all other connection parameters.
    */
    if (argc > 1)
        strcpy(conninfo, argv[1]);
    else
        sprintf(conninfo,
            "dbname=%s port=%s host=%s application_name=test connect_timeout=5 sslmode=allow user=%s
password=%s",
            dbname, port, host, username, passwd);

    /* Connect to the database. */
    conn = PQconnectdb(conninfo);

    /* Check whether the backend connection has been successfully established. */
    if (PQstatus(conn) != CONNECTION_OK)
    {
        fprintf(stderr, "Connection to database failed: %s",
            PQerrorMessage(conn));
        exit_nicely(conn);
    }

    /*
    * After the connection is successful
    * Since a cursor is used in the test case, a transaction block is required.
    * Put all data in one "select * from pg_database".
    * PQexec() is too simple and is not recommended.
    */

    /* Start a transaction block. */
    res = PQexec(conn, "BEGIN");
    if (PQresultStatus(res) != PGRES_COMMAND_OK)
    {
        fprintf(stderr, "BEGIN command failed: %s", PQerrorMessage(conn));
```

```
PQclear(res);
exit_nicely(conn);
}

/*
 * PQclear PGresult should be executed when it is no longer needed, to avoid memory leakage.
 */
PQclear(res);

/*
 * Fetch data from the pg_database system catalog.
 */
res = PQexec(conn, "DECLARE myportal CURSOR FOR select * from pg_database");
if (PQresultStatus(res) != PGRES_COMMAND_OK)
{
    fprintf(stderr, "DECLARE CURSOR failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}
PQclear(res);

res = PQexec(conn, "FETCH ALL in myportal");
if (PQresultStatus(res) != PGRES_TUPLES_OK)
{
    fprintf(stderr, "FETCH ALL failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}

/* First, print out the attribute name. */
nFields = PQnfields(res);
for (i = 0; i < nFields; i++)
    printf("%-15s", PQfname(res, i));
printf("\n\n");

/* Print lines. */
for (i = 0; i < PQntuples(res); i++)
{
    for (j = 0; j < nFields; j++)
        printf("%-15s", PQgetvalue(res, i, j));
    printf("\n");
}

/* Release the memory of the result object to avoid memory leakage. */
PQclear(res);

/* Close the portal. We do not need to check for errors. */
res = PQexec(conn, "CLOSE myportal");
PQclear(res);

/* End the transaction. */
res = PQexec(conn, "END");
PQclear(res);

/* Close the database connection and clean up the database. */
PQfinish(conn);

return 0;
}
```

In the following command output, **user_name** indicates the username of a database administrator, which varies according to the actual environment:

datname	datdba	encoding	datcollate	datctype	datistemplate	datallowconn
datconnlimit	datlastsysoid	datfrozensid	dattablespace	datcompatibilitydatacl		datfrozensid64
datminmxid	dattimezone	dattype				
template_pdb	10	7	en_US.UTF-8	en_US.UTF-8	t	t
12837	0	1663	A	3	2	PRC
templatea	10	7	en_US.UTF-8	en_US.UTF-8	t	f
						-1

12837	0	1663	A	{=c/user_name,user_name=CTc/user_name}	41372
2	PRC	D			
template1	10	7	en_US.UTF-8	en_US.UTF-8	t t -1
12837	0	1663	A	{=c/user_name,user_name=CTc/user_name}	40414
2	PRC	D			
templatem	10	7	en_US.UTF-8	en_US.UTF-8	t t -1
12837	0	1663	M	{=c/user_name,user_name=CTc/user_name}	55146
2	PRC	D			
template0	10	7	en_US.UTF-8	en_US.UTF-8	t f -1
12837	0	1663	A	{=c/user_name,user_name=CTc/user_name}	39935
2	PRC	D			
postgres	10	7	en_US.UTF-8	en_US.UTF-8	f t -1
12837	0	1663	A	40893	2 PRC D

5.5.3.2 Executing Prepared Statements

```

/*
 * testlibpq2.c Test PQprepare
 * PQprepare creates a prepared statement with specified parameters for PQexecPrepared to execute the
 * prepared statement.
 */
#include <stdio.h>
#include <stdlib.h>
#include <libpq-fe.h>
#include <string.h>
int main(int argc, char * argv[])
{
    /* The values of variables such as user and passwd must be read from environment variables or
    configuration files. Environment variables need to be configured as required. If no environment variable is
    used, a character string can be directly assigned. */
    PGconn *conn;
    PGresult * res;
    ConnStatusType pgstatus;
    char connstr[1024];
    char cmd_sql[2048];
    int nParams = 0;
    int paramLengths[5];
    int paramFormats[5];
    Oid paramTypes[5];
    char * paramValues[5];
    int i, cnt;
    char cid[32];
    int k;
    char *passwd = getenv("EXAMPLE_PASSWD_ENV");
    char *port = getenv("EXAMPLE_PORT_ENV");
    char *hostaddr = getenv("EXAMPLE_HOST_ENV");
    char *username = getenv("EXAMPLE_USERNAME_ENV");
    char *dbname = getenv("EXAMPLE_DBNAME_ENV");

    /* Use PQconnectdb to connect to the database. connstr indicates the detailed connection information.*/
    sprintf(connstr,
            "hostaddr=%s dbname=%s port=%s user=%s password=%s",
            hostaddr, dbname, port, username, passwd);
    conn = PQconnectdb(connstr);
    pgstatus = PQstatus(conn);
    if (pgstatus == CONNECTION_OK)
    {
        printf("Connect database success!\n");
    }
    else
    {
        printf("Connect database fail:%s\n",PQerrorMessage(conn));
        return -1;
    }

    /* Create table t01. */
    res = PQexec(conn, "DROP TABLE IF EXISTS t01;CREATE TABLE t01(a int, b int);INSERT INTO t01
    values(1, 23);");
    if (PQresultStatus(res) != PGRES_COMMAND_OK)

```

```
{
    printf("Command failed: %s.\n", PQerrorMessage(conn));
    PQfinish(conn);
    return -1;
}

/* cmd_sql query */
sprintf(cmd_sql, "SELECT b FROM t01 WHERE a = $1");
/*Parameter corresponding to $1 in cmd_sql*/
paramTypes[0] = 23;
/* PQprepare creates a prepared statement with given parameters. */
res = PQprepare(conn,
                "pre_name",
                cmd_sql,
                1,
                paramTypes);
if( PQresultStatus(res) != PGRES_COMMAND_OK )
{
    printf("Failed to prepare SQL : %s\n: %s\n",cmd_sql, PQerrorMessage(conn));
    PQfinish(conn);
    return -1;
}
PQclear(res);
paramValues[0] = cid;
for (k=0; k<2; k++)
{
    sprintf(cid, "%d", 1);
    paramLengths[0] = 6;
    paramFormats[0] = 0;
    /*Execute the prepared statement.*/
    res = PQexecPrepared(conn,
                        "pre_name",
                        1,
                        paramValues,
                        paramLengths,
                        paramFormats,
                        0);

    if( (PQresultStatus(res) != PGRES_COMMAND_OK ) && (PQresultStatus(res) != PGRES_TUPLES_OK))
    {
        printf("%s\n",PQerrorMessage(conn));
        PQclear(res);
        PQfinish(conn);
        return -1;
    }
    cnt = PQntuples(res);
    printf("return %d rows\n", cnt);
    for (i=0; i<cnt; i++)
    {
        printf("row %d: %s\n", i, PQgetvalue(res, i, 0));
    }
    PQclear(res);
}
/* The execution is complete. Close the connection. */
PQfinish(conn);
return 0;
}
```

The command output is as follows:

```
Connect database success!
NOTICE: table "t01" does not exist, skipping
return 1 rows
row 0: 23
return 1 rows
row 0: 23
```

5.5.3.3 Binding Parameters and Returning a Binary Result

```
/*
 * testlibpq3.c
```

```
* Test PQexecParams.
* PQexecParams executes a command with parameters and requests the query result in binary format.
* Before running this example, run the following command to populate a database:
*
*
* CREATE TABLE test1 (i int4, t text);
*
* INSERT INTO test1 values (2, 'ho there');
*
* The expected output is as follows:
*
*
* tuple 0: got
* i = (4 bytes) 2
* t = (8 bytes) 'ho there'
*
*/
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <sys/types.h>
#include <libpq-fe.h>

/* for ntohl/htonl */
#include <netinet/in.h>
#include <arpa/inet.h>

static void
exit_nicely(PGconn *conn)
{
    PQfinish(conn);
    exit(1);
}

/*
* This function is used to print out the query results. The results are in binary format
* and fetched from the table created in the comment above.
*/
static void
show_binary_results(PGresult *res)
{
    int    i;
    int    i_fnum,
           t_fnum;

    /* Use PQfnumber to avoid assumptions about field order in the result. */
    i_fnum = PQfnumber(res, "i");
    t_fnum = PQfnumber(res, "t");

    for (i = 0; i < PQntuples(res); i++)
    {
        char    *iptr;
        char    *tptr;
        int      ival;

        /* Obtain the field value. (Ignore the possibility that they may be null.) */
        iptr = PQgetvalue(res, i, i_fnum);
        tptr = PQgetvalue(res, i, t_fnum);

        /*
        * The binary representation of INT4 is the network byte order,
        * which is better to be replaced with the local byte order.
        */
        ival = ntohl(*(uint32_t *) iptr);

        /*
        * The binary representation of TEXT is text. Since libpq can append a zero byte to it,
        * and think of it as a C string.
        */
    }
}
```

```
    */
    printf("tuple %d: got\n", i);
    printf(" i = (%d bytes) %d\n",
        PQgetlength(res, i, i_fnum), ival);
    printf(" t = (%d bytes) '%s'\n",
        PQgetlength(res, i, t_fnum), tptr);
    printf("\n\n");
}
}

int
main(int argc, char **argv)
{
    /* The values of variables such as user and passwd must be read from environment variables or
    configuration files. Environment variables need to be configured as required. If no environment variable is
    used, a character string can be directly assigned. */
    const char conninfo[1024];
    PGconn    *conn;
    PGresult   *res;
    const char *paramValues[1];
    int        paramLengths[1];
    int        paramFormats[1];
    uint32_t    binaryIntVal;
    char        *passwd = getenv("EXAMPLE_PASSWD_ENV");
    char        *port = getenv("EXAMPLE_PORT_ENV");
    char        *hostaddr = getenv("EXAMPLE_HOST_ENV");
    char        *username = getenv("EXAMPLE_USERNAME_ENV");
    char        *dbname = getenv("EXAMPLE_DBNAME_ENV");

    /*
    * If the user provides a parameter on the command line,
    * The value of this parameter is a conninfo character string. Otherwise,
    * Use environment variables or default values.
    */
    if (argc > 1)
        strcpy(conninfo, argv[1]);
    else
        sprintf(conninfo,
            "dbname=%s port=%s host=%s application_name=test connect_timeout=5 sslmode=allow user=%s",
            dbname, port, hostaddr, username, passwd);

    /* Connect to the database. */
    conn = PQconnectdb(conninfo);

    /* Check whether the connection to the server was successfully established. */
    if (PQstatus(conn) != CONNECTION_OK)
    {
        fprintf(stderr, "Connection to database failed: %s",
            PQerrorMessage(conn));
        exit_nicely(conn);
    }

    res = PQexec(conn, "drop table if exists test1;CREATE TABLE test1 (i int4, t text);");
    if (PQresultStatus(res) != PGRES_COMMAND_OK)
    {
        fprintf(stderr, "command failed: %s", PQerrorMessage(conn));
        PQclear(res);
        exit_nicely(conn);
    }

    PQclear(res);

    res = PQexec(conn, "INSERT INTO test1 values (2, 'ho there');");
    if (PQresultStatus(res) != PGRES_COMMAND_OK)
    {
        fprintf(stderr, "command failed: %s", PQerrorMessage(conn));
        PQclear(res);
    }
}
```

```
    exit_nicely(conn);
}

PQclear(res);

/* Convert the integer value "2" to the network byte order. */
binaryIntVal = htonl((uint32_t) 2);

/* Set the parameter array for PQexecParams. */
paramValues[0] = (char *) &binaryIntVal;
paramLengths[0] = sizeof(binaryIntVal);
paramFormats[0] = 1; /* Binary */

/* PQexecParams executes a command with parameters. */
res = PQexecParams(conn,
    "SELECT * FROM test1 WHERE i = $1::int4",
    1, /* One parameter */
    NULL, /* Enable the backend to deduce the parameter type. */
    paramValues,
    paramLengths,
    paramFormats,
    1); /* Binary result is required. */

if (PQresultStatus(res) != PGRES_TUPLES_OK)
{
    fprintf(stderr, "SELECT failed: %s", PQerrorMessage(conn));
    PQclear(res);
    exit_nicely(conn);
}
/* Output the binary result.*/
show_binary_results(res);

PQclear(res);

/* Close the database connection and clean up the database. */
PQfinish(conn);

return 0;
}
```

The command output is as follows:

```
tuple 0: got
i = (4 bytes) 2
t = (8 bytes) 'ho there'
```

5.5.4 libpq API Reference

5.5.4.1 Database Connection Control Functions

Database connection control functions control the connections to database servers. An application can connect to multiple servers at a time. For example, a client connects to multiple databases. Each connection is represented by a `PGconn` object, which is obtained from the function `PQconnectdb`, `PQconnectdbParams`, or `PQsetdbLogin`. Note that these functions will always return a non-null object pointer, unless memory allocation fails. In that case, a null pointer will be returned. The interface for establishing a connection is stored in the `PGconn` object. The `PQstatus` function can be called to check the return value for a successful connection.

5.5.4.1.1 PQconnectdbParams

Description

Establishes a new connection with the database server.

Prototype

```
PGconn *PQconnectdbParams(const char * const *keywords,  
                           const char * const *values,  
                           int expand_dbname);
```

Parameter

Table 5-54 PQconnectdbParams parameters

Keyword	Description
keywords	An array of strings, each of which is a keyword.
values	Value assigned to each keyword.
expand_dbname	When expand_dbname is non-zero, the dbname keyword value can be recognized as a connection string. Only dbname that first appears is expanded in this way, and any subsequent dbname value is treated as a database name.

Return Value

PGconn * points to the object pointer that contains a connection. The memory is applied for by the function internally.

Precautions

This function establishes a new database connection using the parameters taken from two NULL-terminated arrays. Unlike PQsetdbLogin, the parameter set can be extended without changing the function signature. Therefore, use of this function (or its non-blocking analogs PQconnectStartParams and PQconnectPoll) is preferred for new application programming.

5.5.4.1.2 PQconnectdb

Description

Establishes a new connection with the database server.

Prototype

```
PGconn *PQconnectdb(const char *conninfo);
```


Parameter

Table 5-55 PQconnectdb parameter

Keyword	Description
conninfo	Connection string. For details about the fields in the string, see Connection Parameters .

Return Value

PGconn * points to the object pointer that contains a connection. The memory is applied for by the function internally.

Precautions

- This function opens a new connection with the parameters specified by **conninfo**.
- The input parameter can be empty, indicating that all default parameters can be used. It can also contain one or more parameters separated by spaces or it can contain a URL.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.1.3 PQconninfoParse

Function

PQconninfoParse is used to return parsed connection options based on the connection.

Prototype

```
PQconninfoOption* PQconninfoParse(const char* conninfo, char** errmsg);
```

Parameters

Table 5-56

Keyword	Parameter Description
conninfo	Passed string. This parameter can be left empty. In this case, the default value is used. It can contain one or more values separated by spaces or contain a URL.
errmsg	Error information.

Return Value

PQconninfoOption pointers

5.5.4.1.4 PQconnectStart

Description

Establishes a non-blocking connection with the database server.

Prototype

```
PGconn* PQconnectStart(const char* conninfo);
```

Parameters

Table 5-57

Keyword	Description
conninfo	String of connection information. This parameter can be left empty. In this case, the default value is used. It can contain one or more values separated by spaces or contain a URI.

Return Value

PGconn pointers

5.5.4.1.5 PQerrorMessage

Function

PQerrorMessage is used to return error information on a connection.

Prototype

```
char* PQerrorMessage(const PGconn* conn);
```

Parameter

Table 5-58

Keyword	Parameter Description
conn	Connection handle.

Return Value

char pointers

Example

For details, see [Typical Application Development Examples](#).

5.5.4.1.6 PQsetdbLogin

Description

Establishes a new connection with the database server.

Prototype

```
PGconn *PQsetdbLogin(const char *pghost,  
                    const char *pgport,  
                    const char *pgoptions,  
                    const char *pgtty,  
                    const char *dbName,  
                    const char *login,  
                    const char *pwd);
```

Parameters

Table 5-59 PQsetdbLogin parameters

Keyword	Description
pghost	Name of the host to be connected. For details, see the host field described in Connection Parameters .
pgport	Port number of the host server. For details, see the port field described in Connection Parameters .
pgoptions	Command-line options to be sent to the server during running. For details, see the options field described in Connection Parameters .
pgtty	Ignored. (This option declares the output direction of server logs.)
dbName	Name of the database to be connected. For details, see the dbname field described in Connection Parameters .
login	Username for connection. For details, see the user field described in Connection Parameters .
pwd	Password used for authentication during connection. For details, see the password field described in Connection Parameters .

Return Value

PGconn * points to the object pointer that contains a connection. The memory is applied for by the function internally.

Precautions

- This function is the predecessor of PQconnectdb with a fixed set of parameters. When an undefined parameter is called, its default value is used. Write NULL or an empty string for any one of the fixed parameters that is to be defaulted.
- If the **dbName** value contains an equal sign (=) or a valid prefix in the connection URL, it is taken as a conninfo string and passed to PQconnectdb, and the remaining parameters are consistent with PQconnectdbParams parameters.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.1.7 PQfinish

Function

Closes the connection to the server and releases the memory used by the PGconn object.

Prototype

```
void PQfinish(PGconn *conn);
```

Parameter

Table 5-60 PQfinish parameter

Keyword	Parameter Description
conn	Points to the object pointer that contains the connection information.

Precautions

If the server connection attempt fails (as indicated by PQstatus), the application should call PQfinish to release the memory used by the PGconn object. The PGconn pointer must not be used again after PQfinish has been called.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.1.8 PQreset

Description

Resets the communication port to the server.

Prototype

```
void PQreset(PGconn *conn);
```

Parameter

Table 5-61 PQreset parameter

Keyword	Description
conn	Points to the object pointer that contains the connection information.

Precautions

This function will close the connection to the server and attempt to establish a new connection to the same server by using all the parameters previously used. This function is applicable to fault recovery after a connection exception occurs.

5.5.4.1.9 PQstatus

Description

Returns the connection status.

Prototype

```
ConnStatusType PQstatus(const PGconn *conn);
```

Parameter

Table 5-62 PQ status parameter

Keyword	Description
conn	Points to the object pointer that contains the connection information.

Return Value

ConnStatusType indicates the connection status. The enumerated values are as follows:

```
CONNECTION_STARTED  
Waiting for the connection to be established.  
  
CONNECTION_MADE  
Connection succeeded; waiting to send  
  
CONNECTION_AWAITING_RESPONSE  
Waiting for a response from the server.  
  
CONNECTION_AUTH_OK
```

Authentication received; waiting for backend startup to complete.

CONNECTION_SSL_STARTUP
Negotiating SSL encryption.

CONNECTION_SETENV
Negotiating environment-driven parameter settings.

CONNECTION_OK
Normal connection.

CONNECTION_BAD
Failed connection.

Precautions

The connection status can be one of the preceding values. After the asynchronous connection procedure is complete, only two of them, **CONNECTION_OK** and **CONNECTION_BAD**, can return. **CONNECTION_OK** indicates that the connection to the database is normal. **CONNECTION_BAD** indicates that the connection attempt fails. Generally, the normal state remains until PQfinish is called. However, a communication failure may cause the connection status to become to **CONNECTION_BAD** before the connection procedure is complete. In this case, the application can attempt to call PQreset to restore the communication.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.2 Database Statement Execution Functions

After the connection to the database server is successfully established, you can use the functions described in this section to execute SQL queries and commands.

5.5.4.2.1 PQclear

Function

PQclear is used to release the storage associated with PGresult. Any query result should be released by PQclear when it is no longer needed.

Prototype

```
void PQclear(PGresult *res);
```

Parameters

Table 5-63 PQclear parameter

Keyword	Parameter Description
res	Object pointer that contains the query result.

Precautions

PGresult is not automatically released. That is, it does not disappear when a new query is submitted or even if you close the connection. To delete it, you must call PQclear. Otherwise, memory leakage occurs.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.2.2 PQexec

Description

Commits a command to the server and waits for the result.

Prototype

```
PGresult *PQexec(PGconn *conn, const char *command);
```

Parameter

Table 5-64 PQexec parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
command	Query string to be executed.

Return Value

PGresult indicates the object pointer that contains the query result.

Precautions

The PQresultStatus function should be called to check the return value for any errors (including the value of a null pointer, in which **PGRES_FATAL_ERROR** will be returned). The PQerrorMessage function can be called to obtain more information about such errors.

NOTICE

The command string can contain multiple SQL commands separated by semicolons (;). Multiple queries sent in a PQexec call are processed in one transaction, unless there are specific BEGIN/COMMIT commands in the query string to divide the string into multiple transactions. Note that the returned PGresult structure describes only the result of the last command executed from the string. If a command fails, the string processing stops and the returned PGresult describes the error condition.

Example

For details, see [Typical Application Development Examples](#).

5.5.4.2.3 PQexecParams

Description

Runs a command to bind parameters.

Prototype

```
PGresult* PQexecParams(PGconn* conn,
    const char* command,
    int nParams,
    const Oid* paramTypes,
    const char* const* paramValues,
    const int* paramLengths,
    const int* paramFormats,
    int resultFormat);
```

Parameter

Table 5-65 PQexecParams parameters

Keyword	Description
conn	Connection handle.
command	SQL text string.
nParams	Number of parameters to be bound.
paramTypes	Types of parameters to be bound.
paramValues	Values of parameters to be bound.
paramLengths	Parameter lengths.
paramFormats	Parameter formats (text or binary).
resultFormat	Result format (text or binary).

Return Value

PGresult pointers

5.5.4.2.4 PQexecParamsBatch

Description

Runs a command to bind batches of parameters.

Prototype

```
PGresult* PQexecParamsBatch(PGconn* conn,
    const char* command,
```



```
int nParams,  
int nBatch,  
const Oid* paramTypes,  
const char* const* paramValues,  
const int* paramLengths,  
const int* paramFormats,  
int resultFormat);
```

Parameters

Table 5-66 PQexecParamsBatch parameters

Keyword	Description
conn	Connection handle.
command	SQL text string.
nParams	Number of parameters to be bound.
nBatch	Number of batch operations.
paramTypes	Types of parameters to be bound.
paramValues	Values of parameters to be bound.
paramLengths	Parameter lengths.
paramFormats	Parameter formats (text or binary).
resultFormat	Result format (text or binary).

Return Value

PGresult pointers.

5.5.4.2.5 PQexecPrepared

Function

PQexecPrepared is used to send a request to execute a prepared statement with given parameters and wait for the result.

Prototype

```
PGresult* PQexecPrepared(PGconn* conn,  
const char* stmtName,  
int nParams,  
const char* const* paramValues,  
const int* paramLengths,  
const int* paramFormats,  
int resultFormat);
```

Parameter

Table 5-67 PQexecPrepared parameters

Keyword	Parameter Description
conn	Connection handle.
stmtName	<i>stmt</i> name, which can be set to "" or NULL to reference an unnamed statement. Otherwise, it must be the name of an existing prepared statement.
nParams	Parameter quantity.
paramValues	Actual values of parameters.
paramLengths	Actual data lengths of parameters.
paramFormats	Parameter formats (text or binary).
resultFormat	Return result format (text or binary).

Return Value

PGresult pointers

5.5.4.2.6 PQexecPreparedBatch

Description

Sends a request to execute a prepared statement with batches of given parameters and waits for the result.

Prototype

```
PGresult* PQexecPreparedBatch(PGconn* conn,  
                              const char* stmtName,  
                              int nParams,  
                              int nBatchCount,  
                              const char* const* paramValues,  
                              const int* paramLengths,  
                              const int* paramFormats,  
                              int resultFormat);
```

Parameter

Table 5-68 PQexecPreparedBatch parameters

Keyword	Description
conn	Connection handle.
stmtName	<i>stmt</i> name, which can be set to "" or NULL to reference an unnamed statement. Otherwise, it must be the name of an existing prepared statement.

Keyword	Description
nParams	Parameter quantity.
nBatchCount	Number of batches.
paramValues	Actual values of parameters.
paramLengths	Actual data lengths of parameters.
paramFormats	Parameter formats (text or binary).
resultFormat	Return result format (text or binary).

Return Value

PGresult pointers

5.5.4.2.7 PQprepare

Description

Commits a request to create a prepared statement with given parameters and waits for completion.

Prototype

```
PGresult *PQprepare(PGconn *conn,  
    const char *stmtName,  
    const char *query,  
    int nParams,  
    const Oid *paramTypes);
```

Parameters

Table 5-69 PQprepare parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
stmtName	Name of stmt to be executed.
query	Query string to be executed.
nParams	Parameter quantity.
paramTypes	Array of the parameter type.

Return Value

PGresult indicates the object pointer that contains the query result.

Precautions

- PQprepare creates a prepared statement for later execution with PQexecPrepared. This function allows commands to be repeatedly executed, without being parsed and planned each time they are executed. PQprepare is supported only in protocol 3.0 or later. It will fail when protocol 2.0 is used.
- This function creates a prepared statement named **stmtName** from the query string, which must contain an SQL command. **stmtName** can be "" to create an unnamed statement. In this case, any pre-existing unnamed statement will be automatically replaced. Otherwise, this is an error if the statement name has been defined in the current session. If any parameters are used, they are referred to in the query as \$1, \$2, and so on. **nParams** is the number of parameters for which types are pre-specified in the array paramTypes[]. (The array pointer can be **NULL** when **nParams** is 0.) paramTypes[] specifies the data types to be assigned to the parameter symbols by OID. If **paramTypes** is **NULL**, or any element in the array is 0, the server assigns a data type to the parameter symbol in the same way as it does for an untyped literal string. In addition, the query can use parameter symbols whose numbers are greater than **nParams**. Data types of these symbols will also be inferred.

NOTICE

You can also execute the **SQLPREPARE** statement to create a prepared statement that is used with PQexecPrepared. Although there is no libpq function of deleting a prepared statement, the **SQL DEALLOCATE** statement can be used for this purpose.

Examples

For details, see [Typical Application Development Examples](#).

5.5.4.3 Asynchronous Command Processing Functions

The PQexec function is adequate for committing commands in common, synchronous applications. However, it has several defects, which may be important to some users:

- PQexec waits for the end of the command, but the application may have other work to do (for example, maintaining a user interface). In this case, PQexec would not want to be blocked to wait for the response.
- As the client application is suspended while waiting for the result, it is difficult for the application to determine whether to cancel the ongoing command.
- PQexec can return only one PGresult structure. If the committed command string contains multiple SQL commands, all the PGresult structures except the last PGresult are discarded by PQexec.
- PQexec always collects the entire result of the command and caches it in a PGresult. Although this mode simplifies the error handling logic for applications, it is impractical for results that contain multiple rows.

Applications that do not want to be restricted by these limitations can use the following functions that PQexec is built from: PQsendQuery and PQgetResult. The

functions PQsendQueryParams, PQsendPrepare, and PQsendQueryPrepared can also be used with PQgetResult.

5.5.4.3.1 PQsendQuery

Description

Commits a command to the server without waiting for the result. If the query is successful, **1** is returned. Otherwise, **0** is returned.

Prototype

```
int PQsendQuery(PGconn *conn, const char *command);
```

Parameter

Table 5-70 PQsendQuery parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
command	Query string to be executed.

Return Value

int indicates the execution result. **1** indicates successful execution and **0** indicates an execution failure. The failure cause is stored in **conn->errorMessage**.

Precautions

After PQsendQuery is successfully called, call PQgetResult one or more times to obtain the results. PQsendQuery cannot be called again (on the same connection) until PQgetResult returns a null pointer, indicating that the command execution is complete.

5.5.4.3.2 PQsendQueryParams

Description

Sends a command with separated parameters to the server without waiting for the result.

Prototype

```
int PQsendQueryParams(PGconn *conn,  
                      const char *command,  
                      int nParams,  
                      const Oid *paramTypes,  
                      const char * const *paramValues,  
                      const int *paramLengths,  
                      const int *paramFormats,  
                      int resultFormat);
```

Parameter

Table 5-71 PQsendQueryParams parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
command	Query string to be executed.
nParams	Parameter quantity.
paramTypes	Parameter type.
paramValues	Parameter value.
paramLengths	Parameter length.
paramFormats	Parameter format.
resultFormat	Result format.

Return Value

int indicates the execution result. **1** indicates successful execution and **0** indicates an execution failure. The failure cause is stored in **conn->errorMessage**.

Precautions

This function is equivalent to PQsendQuery. The only difference is that query parameters can be specified separately from the query string. The parameter processing of this function is similar to that of PQexecParams. It cannot work on connections using protocol v2.0, and it allows only one command to appear in the query string.

5.5.4.3 PQsendPrepare

Description

PQsendPrepare is used to send a request to create a prepared statement with given parameters, without waiting for completion.

Prototype

```
int PQsendPrepare(PGconn *conn,  
                 const char *stmtName,  
                 const char *query,  
                 int nParams,  
                 const Oid *paramTypes);
```

Parameters

Table 5-72 PQsendPrepare parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
stmtName	Name of prepare to be executed.
query	Query string to be executed.
nParams	Parameter quantity.
paramTypes	Array of the parameter type.

Return Value

int indicates the execution result. **1** indicates successful execution and **0** indicates an execution failure. The failure cause is stored in **conn->errorMessage**.

Precautions

PQsendPrepare is an asynchronous version of PQprepare. If it can dispatch a request, **1** is returned. Otherwise, **0** is returned. After a successful calling, call PQgetResult to check whether the server successfully created the prepared statement. PQsendPrepare parameters are handled in the same way as PQprepare parameters. Like PQprepare, it cannot work on connections using protocol v2.0.

5.5.4.3.4 PQsendQueryPrepared

Description

Sends a request to execute a prepared statement with given parameters, without waiting for the result.

Prototype

```
int PQsendQueryPrepared(PGconn *conn,  
                        const char *stmtName,  
                        int nParams,  
                        const char * const *paramValues,  
                        const int *paramLengths,  
                        const int *paramFormats,  
                        int resultFormat);
```

Parameters

Table 5-73 PQsendQueryPrepared parameters

Keyword	Description
conn	Points to the object pointer that contains the connection information.
stmtName	Name of prepare to be executed.
nParams	Parameter quantity.
paramValues	Parameter value.
paramLengths	Parameter length.
paramFormats	Parameter format.
resultFormat	Result format.

Return Value

int indicates the execution result. **1** indicates successful execution and **0** indicates an execution failure. The failure cause is stored in **conn->errorMessage**.

Precautions

This function is similar to PQsendQueryParams, but the command to be executed is specified by naming a previously-prepared statement, instead of providing a query string. The parameters of this function are processed in the same way as PQexecPrepared. Like PQexecPrepared, it cannot work on connections using protocol v2.0.

5.5.4.3.5 PQflush

Description

Tries to flush any queued output data to the server.

Prototype

```
int PQflush(PGconn *conn);
```

Parameter

Table 5-74 PQflush parameter

Keyword	Description
conn	Points to the object pointer that contains the connection information.

Return Value

int indicates the execution result. If the operation is successful (or the send queue is empty), **0** is returned. If the operation fails, **-1** is returned. If all data in the send queue fails to be sent, **1** is returned. (This case occurs only when the connection is non-blocking.) The failure cause is stored in **conn->errorMessage**.

Precautions

Call PQflush after sending any command or data over a non-blocking connection. If **1** is returned, wait for the socket to become read- or write-ready. If the socket becomes write-ready, call PQflush again. If the socket becomes read-ready, call PQconsumeInput and then call PQflush again. Repeat the operation until the value **0** is returned for PQflush. It is necessary to check read-ready and use PQconsumeInput to exhaust input because the server may prevent attempts to send data (such as NOTICE messages) to the client and does not read the client's data until the client reads its data. Once PQflush returns **0**, wait for the socket to be read-ready and then read the response as described above.

5.5.4.4 Functions for Canceling Query Processing

A client application can use the functions described in this section to cancel a command that is still being processed by the server.

5.5.4.4.1 PQgetCancel

Description

Creates a data structure that contains the information required to cancel a command issued through a specific database connection.

Prototype

```
PGcancel *PQgetCancel(PGconn *conn);
```

Parameter

Table 5-75 PQgetCancel parameter

Keyword	Description
conn	Points to the object pointer that contains the connection information.

Return Value

PGcancel points to the object pointer that contains the cancel information.

Precautions

PQgetCancel creates a PGcancel object for a given PGconn connection object. If the given connection object (**conn**) is NULL or an invalid connection, it will return

NULL. The PGcancel object is an opaque structure that cannot be directly accessed by applications. It can be passed only to PQcancel or PQfreeCancel.

5.5.4.4.2 PQfreeCancel

Description

Releases the data structure created by PQgetCancel.

Prototype

```
void PQfreeCancel(PGcancel *cancel);
```

Parameter

Table 5-76 PQfreeCancel parameter

Keyword	Description
cancel	Points to the object pointer that contains the cancel information.

Precautions

PQfreeCancel releases a data object previously created by PQgetCancel.

5.5.4.4.3 PQcancel

Description

Requests the server to abandon processing of the current command.

Prototype

```
int PQcancel(PGcancel *cancel, char *errbuf, int errbufsize);
```

Parameter

Table 5-77 PQcancel parameters

Keyword	Description
cancel	Points to the object pointer that contains the cancel information.
errbuf	Buffer for storing error information.
errbufsize	Size of the buffer for storing error information.

Return Value

int indicates the execution result. **1** indicates successful execution and **0** indicates an execution failure. The failure cause is stored in **errbuf**.

Precautions

- Successful sending does not guarantee that the request will have any effect. If the cancellation is valid, the current command is terminated early and an error is returned. If the cancellation fails (for example, because the server has processed the command), no result is returned.
- If **errbuf** is a local variable in a signal handler, you can safely call PQcancel from the signal handler. For PQcancel, the PGcancel object is read-only, so it can also be called from a thread that is separate from the thread that is operating the PGconn object.

5.5.4.5 Functions for Processing Database Result

5.5.4.5.1 PQgetvalue

Description

Returns a single field value of one row of a PGresult. Row and column numbers start from 0. The caller should not release the result directly. The result will be released when the associated PGresult handle is passed to PQclear.

Prototype

```
char *PQgetvalue(const PGresult *res,  
                 int row_number,  
                 int column_number);
```

Parameter

Table 5-78 PQgetvalue parameters

Keyword	Description
res	Operation result handle.
row_number	Number of rows.
column_number	Number of columns.

Return Value

For data in text format, the value returned by PQgetvalue is a null-terminated string representation of the field value.

For binary data, the value is a binary representation determined by the typsend and typreceive functions of the data type.

If this field is left blank, an empty string is returned.

5.5.4.5.2 PQfname

Description

Returns the column name associated with the given column number. Column numbers start from 0. The caller should not release the result directly. The result will be released when the associated PGresult handle is passed to PQclear.

Prototype

```
char *PQfname(const PGresult *res,  
              int column_number);
```

Parameter

Table 5-79 PQfname parameters

Keyword	Description
res	Operation result handle.
column_number	Number of columns.

Return Value

char pointers

5.5.4.5.3 PQnfields

Description

Returns the number of columns (fields) in each row of the query result.

Prototype

```
int PQnfields(const PGresult *res);
```

Parameter

Table 5-80 PQnfields parameters

Keyword	Description
res	Operation result handle.

Return Value

Value of the int type

Example

For details, see [Typical Application Development Examples](#).

5.5.4.5.4 PQntuples

Function

PQntuples is used to return the number of rows (tuples) in the query result. An overflow may occur if the return value is out of the value range allowed in a 32-bit OS.

Prototype

```
int PQntuples(const PGresult *res);
```

Parameter

Table 5-81 PQntuples parameters

Keyword	Parameter Description
res	Operation result handle.

Return Value

Value of the int type

Examples

For details, see [Typical Application Development Examples](#).

5.5.4.5.5 PQresultStatus

Description

Returns the result status of a command.

Prototype

```
ExecStatusType PQresultStatus(const PGresult *res);
```

Parameter

Table 5-82 PQresultStatus parameter

Keyword	Description
res	Object pointer that contains the query result.

Return Value

PQresultStatus indicates the command execution status. The enumerated values are as follows:

PQresultStatus can return one of the following values:

PGRES_EMPTY_QUERY

The string sent to the server was empty.

PGRES_COMMAND_OK

A command that does not return data was successfully executed.

PGRES_TUPLES_OK

A query (such as SELECT or SHOW) that returns data was successfully executed.

PGRES_COPY_OUT

Copy Out (from the server) data transfer started.

PGRES_COPY_IN

Copy In (to the server) data transfer started.

PGRES_BAD_RESPONSE

The response from the server cannot be understood.

PGRES_NONFATAL_ERROR

A non-fatal error (notification or warning) occurred.

PGRES_FATAL_ERROR

A fatal error occurred.

PGRES_COPY_BOTH

Copy In/Out (to and from the server) data transfer started. This state occurs only in streaming replication.

PGRES_SINGLE_TUPLE

PQresult contains a result tuple from the current command. This state occurs in a single-row query.

Precautions

- Note that the SELECT command that happens to retrieve zero rows still returns **PGRES_TUPLES_OK**. **PGRES_COMMAND_OK** is used for commands that can never return rows (such as INSERT or UPDATE, without return clauses). The result status **PGRES_EMPTY_QUERY** might indicate a bug in the client software.
- The result status **PGRES_NONFATAL_ERROR** will never be returned directly by PQexec or other query execution functions. Instead, such results will be passed to the notice processor.

Example

For details, see [Typical Application Development Examples](#).

5.6 Psycopg-based Development

Psycopg is a Python API used to execute SQL statements and provides a unified access API for GaussDB. Applications can perform data operations based on psycopg. Psycopg2 is the encapsulation of libpq and is implemented using the C language, which is efficient and secure. It provides cursors on both clients and servers, asynchronous communication and notification, and the COPY TO and COPY FROM functions. It supports multiple types of Python out-of-the-box and

adapts to GaussDB data types. Through the flexible object adaptation system, you can extend and customize the adaptation. Psycopg2 is compatible with Unicode.

GaussDB supports the psycopg2 feature and allows psycopg2 to be connected in SSL mode.

Table 5-83 Platforms supported by Psycopg

OS	Platform	Python Version
EulerOS 2.5	Arm64 x86_64	3.8.5
EulerOS 2.9	Arm64 x86_64	3.7.4
EulerOS 2.10, Kylin v10, UnionTech20	Arm64 x86_64	3.7.9
EulerOS 2.11 and SUSE 12.5	Arm64 x86_64	3.9.11

NOTICE

During psycopg2 compilation, OpenSSL of GaussDB is linked. OpenSSL of GaussDB may be incompatible with OpenSSL of the OS. If incompatibility occurs, for example, "version 'OPENSSL_1_1_1f' not found" is displayed, use the environment variable `LD_LIBRARY_PATH` to isolate the OpenSSL provided by the OS and the OpenSSL on which GaussDB depends.


For example, when the application software **client.py** that calls psycopg2 is executed, the environment variable is explicitly assigned to the application software.

```
export LD_LIBRARY_PATH=/path/to/gaussdb/libs:$LD_LIBRARY_PATH python client.py
```

In the preceding command, **/path/to/psycopg2/lib** indicates the directory where the OpenSSL library on which the GaussDB depends is located. Change it as required.

5.6.1 Development Process

Figure 5-5 Application development process based on psycopg2



5.6.2 Psycopg Package

Connecting to a Database

Step 1 Prepare related drivers and dependency libraries. Obtain the package **GaussDB-Kernel_Database version number_OS information_64bit_Python.tar.gz** from the driver package.

After the decompression, the following folders are generated:

- **psycopg2**: **psycopg2** library file
- **lib**: **lib** library file

Step 2 Load the driver.

- Before using the driver, perform the following operations:
 - a. Decompress the driver package of the corresponding version.
`tar zxvf xxxx-Python.tar.gz`

- b. Copy **psycopg2** to the **site-packages** folder in the Python installation directory as the **root** user.

```
su root  
cp psycopg2 $(python3 -c 'import site; print(site.getsitepackages()[0])') -r
```
 - c. Change the **psycopg2** directory permission to **755**.

```
chmod 755 $(python3 -c 'import site; print(site.getsitepackages()[0])')/psycopg2 -R
```
 - d. Add the **psycopg2** directory to the environment variable **\$PYTHONPATH** and validate it.

```
export PYTHONPATH=$(python3 -c 'import site; print(site.getsitepackages()[0])'):$PYTHONPATH
```
 - e. For non-database users, configure the **lib** directory in **LD_LIBRARY_PATH** after decompression.

```
export LD_LIBRARY_PATH=path/to/lib:$LD_LIBRARY_PATH
```
- Load a database driver before creating a database connection.

```
import psycopg2
```

Step 3 Connect to a database.

Connect to the database in non-SSL mode.

1. Use the `psycopg2.connect` function to obtain the connection object.
2. Use the connection object to create a cursor object.

Connect to the database in SSL mode.

When you use `psycopy2` to connect to the GaussDB server, you can enable SSL to encrypt the communication between the client and server. To enable SSL, you must have the server certificate, client certificate, and private key files. For details on how to obtain these files, see related documents and commands of OpenSSL.

1. Use the `.ini` file (the **configparser** package of Python can parse this type of configuration file) to save the configuration information about the database connection.
2. Add SSL connection parameters **sslmode**, **sslcert**, **sslkey**, and **sslrootcert** to the connection options.
 - a. **sslmode**: For details about the options, see [Table 5-84](#).
 - b. **sslcert**: client certificate path.
 - c. **sslkey**: client key path.
 - d. **sslrootcert**: root certificate path.
3. Use the `psycopg2.connect` function to obtain the connection object.
4. Use the connection object to create a cursor object.

 CAUTION

To use SSL to connect to the database, ensure that the Python interpreter is compiled in the mode of generating a dynamic link library (.so) file. You can perform the following steps to check the connection mode of the Python interpreter:

1. Run the **import ssl** command in the Python interpreter to import SSL.
2. Run the **ps ux** command to query the PID of the Python interpreter. Assume that the PID is *********.
3. In the shell CLI, run the **pmap -p ***** | grep ssl** command and check whether the command output contains the path related to **libssl.so**. If it does, the Python interpreter is compiled in dynamic link mode.

Table 5-84 sslmode options

sslmode	Enable SSL Encryption	Description
disable	No	SSL connection is not enabled.
allow	Possible	If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified.
prefer	Possible	If the database supports SSL connection, SSL connection is preferred. However, authenticity of the database server will not be verified.
require	Yes	SSL connection is required, but only data is encrypted. However, authenticity of the database server will not be verified.
verify-ca	Yes	SSL connection is required, and the validity of the server CA must be verified.
verify-full	Yes	The SSL connection must be enabled, which is not supported by GaussDB currently.

Step 4 Run SQL statements.

1. Construct an operation statement and use %s as a placeholder. During execution, psycopg2 will replace the placeholder with the parameter value. You can add the RETURNING clause to obtain the automatically generated column values.
2. Use the cursor.execute method to execute one row of SQL statement, and use the cursor.executemany method to execute multiple rows of SQL statements.

Step 5 Process the result set.

1. `cursor.fetchone()`: fetches the next row in a query result set and returns a sequence. If no data is available, null is returned.
2. `cursor.fetchall()`: fetches all remaining rows in a query result and returns a list. An empty list is returned when no rows are available.

 NOTE

For database-specific data types, such as `tinyint`, the corresponding columns in the query result are character strings.

Step 6 Close the connection.

After you complete required data operations in a database, close the database connection. Call the close method such as `connection.close()` to close the connection.

 CAUTION

This method closes the database connection and does not automatically call `commit()`. If you close the database connection without calling `commit()`, changes will be lost.

----End

5.6.3 Example: Common Operations

```
import psycopg2
import os

# Obtain the username and password from environment variables.
user = os.getenv('user')
password = os.getenv('password')

# Create a connection object.
conn=psycopg2.connect(database="database", user=user, password=password, host="localhost", port=port)
cur=conn.cursor() # Create a pointer object.

# Create a connection object (using SSL).
conn = psycopg2.connect(dbname="database", user=user, password=password, host="localhost", port=port,
                        sslmode="verify-ca", sslcert="client.crt", sslkey="client.key", sslrootcert="cacert.pem")
Note: If sslcert, sslkey, and sslrootcert are not set, the following files in the .postgresql directory of the
current user are used by default: client.crt, client.key, and root.crt.

# Create a table.
cur.execute("CREATE TABLE student(id integer,name varchar,sex varchar);")

# Insert data.
cur.execute("INSERT INTO student(id,name,sex) VALUES(%s,%s,%s)",(1,'Aspirin','M'))
cur.execute("INSERT INTO student(id,name,sex) VALUES(%s,%s,%s)",(2,'Taxol','F'))
cur.execute("INSERT INTO student(id,name,sex) VALUES(%s,%s,%s)",(3,'Dixheral','M'))

# Insert data in batches.
stus = ((4,'John','M'),(5,'Alice','F'),(6,'Peter','M'))
cur.executemany("INSERT INTO student(id,name,sex) VALUES(%s,%s,%s)",stus)

# Obtain the result.
cur.execute("SELECT * FROM student")
results=cur.fetchall()
print (results)
```

```
# Perform a commit.
conn.commit()

# Insert a data record.
cur.execute("INSERT INTO student(id,name,sex) VALUES(%s,%s,%s)",(7,'Lucy','F'))

# Perform a rollback.
conn.rollback()

# Close the connection.
cur.close()
conn.close()

Common connection modes of psycopg2
1. conn = psycopg2.connect(dbname="dbname", user=user, password=password, host="localhost",
port=port)
2. conn = psycopg2.connect(f"dbname=dbname user={user} password={password} host=localhost
port=port")
3. Using logs
import logging
import psycopg2
from psycopg2.extras import LoggingConnection
import os

# Obtain the username and password from environment variables.
user = os.getenv('user')
password = os.getenv('password')

logging.basicConfig(level=logging.DEBUG) # Log level
logger = logging.getLogger(__name__)

db_settings = {
    "user": user,
    "password": password,
    "host": "localhost",
    "database": "dbname",
    "port": port
}

# LoggingConnection records all SQL statements by default. You can filter unnecessary or sensitive SQL
statements. Example of filtering password-related SQL statements:
class SelfLoggingConnection(LoggingConnection):

    def filter(self, msg, curs):
        if db_settings['password'] in msg.decode():
            return b'queries containing the password will not be recorded'
        return msg

conn = psycopg2.connect(connection_factory=SelfLoggingConnection, **db_settings)
conn.initialize(logger)
```

CAUTION

- By default, **LoggingConnection** records all SQL information and does not anonymize sensitive information. You can use the filter function to define the output SQL content.
- The log function is an additional function provided by psycopg2 for developers to explicitly debug full SQL statements. By default, the log function is not used. This function prints SQL statements before psycopg2 executes SQL statements. However, the SQL statements can be printed only when the log level is **DEBUG**. This function is not a default function. It is used only when there are special requirements. You are advised not to use this function unless there are special requirements. For details, visit <https://www.psycopg.org/docs/extras.html?highlight=loggingconnection>.

5.6.4 Psycopg API Reference

Psycopg APIs are a set of methods provided for users. This section describes some common APIs.

5.6.4.1 psycopg2.connect()

Description

This method creates a database session and returns a new connection object.

Prototype

```
import os
conn=psycopg2.connect(dbname="test",user=os.getenv('user'),password=os.getenv('password'),host="127.0.0.1",port="5432")
```

Parameter

Table 5-85 psycopg2.connect parameters

Keyword	Description
dbname	Database name
user	Username
password	Password
host	Database IP address. The default type is Unix socket.
port	Connection port number. The default value is 5432 .
sslmode	SSL mode, which is used for SSL connection.
sslcert	Path of the client certificate, which is used for SSL connection.
sslkey	Path of the client key, which is used for SSL connection.

Keyword	Description
sslrootcert	Path of the root certificate, which is used for SSL connection.
hostaddr	IP address of the database
connect_timeout	Client connection timeout interval
client_encoding	Encoding format of the client
application_name	Value of application_name
fallback_application_name	Rollback value of application_name
keepalives	Specifies whether to enable the TCP connection on the client. The default value is 1 , indicating that the TCP connection is enabled. The value 0 indicates that the TCP connection is disabled. If the UDS connection is used, ignore this parameter.
options	Specifies the command line options sent to the server when the connection starts.
keepalives_idle	Describes inactivity before keepalive messages are sent to the server. If keepalive messages are disabled, ignore this parameter.
keepalives_interval	Specifies whether keepalive messages that are not confirmed by the server need to be resent. If keepalive messages are disabled, ignore this parameter.
keepalives_count	Specifies the number of TCP connections that may be lost before the client is disconnected from the server.
replication	Ensures that the connection uses the replication protocol instead of the common protocol.
requiressl	Supports the SSL mode.
sslcompression	Specifies the SSL compression. If this parameter is set to 1 , the data sent through the SSL connection is compressed. If this parameter is set to 0 , compression is disabled. If no SSL connection is established, ignore this parameter.
sslcrld	Specifies the path of the CRL, which is used to check whether the SSL server certificate is available.
requirepeer	Specifies the OS username of the server.

Return Value

Connection object (for connecting to a database instance)

Examples

For details, see [Example: Common Operations](#).

5.6.4.2 connection.cursor()

Function

This method returns a new cursor object.

Prototype

```
cursor(name=None, cursor_factory=None, scrollable=None, withhold=False)
```

Parameter

Table 5-86 connection.cursor parameters

Keyword	Description
name	Cursor name. The default value is None .
cursor_factory	Creates a non-standard cursor. The default value is None .
scrollable	Sets the SCROLL option. The default value is None .
withhold	Sets the HOLD option. The default value is False .

Return Value

Cursor object (used for cursors that are programmed using Python in the entire database)

Examples

For details, see [Example: Common Operations](#).

5.6.4.3 cursor.execute(query,vars_list)

Description

This method executes the parameterized SQL statements (that is, placeholders instead of SQL literals). The psycopg2 module supports placeholders marked with %s.

Prototype

```
cursor.execute(query,vars_list)
```

Parameters

Table 5-87 cursor.execute parameters

Keyword	Description
query	SQL statement to be executed.
vars_list	Variable list, which matches the %s placeholder in the query.

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.6.4.4 cursor.executemany(query,vars_list)

Description

This method executes an SQL command against all parameter sequences or mappings found in the sequence SQL.

Prototype

```
cursor.executemany(query,vars_list)
```

Parameter

Table 5-88 cursor.executemany parameters

Keyword	Description
query	SQL statement that you want to execute.
vars_list	Variable list, which matches the %s placeholder in the query.

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.6.4.5 connection.commit()

Function

This method commits the currently pending transaction to the database.

CAUTION

By default, Psycopg opens a transaction before executing the first command. If **commit()** is not called, the effect of any data operation will be lost.

Prototype

```
connection.commit()
```

Parameter

None

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.6.4.6 connection.rollback()

Function

This method rolls back the current pending transaction.

CAUTION

If you close the connection using **close()** but do not commit the change using **commit()**, an implicit rollback will be performed.

Prototype

```
connection.rollback()
```

Parameter

None

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.6.4.7 cursor.fetchone()

Description

This method extracts the next row of the query result set and returns a tuple.

Prototype

```
cursor.fetchone()
```

Parameter

None

Return Value

A single tuple is the first result in the result set. If no more data is available, **None** is returned.

Examples

For details, see [Example: Common Operations](#).

5.6.4.8 cursor.fetchall()

Function

This method obtains all the (remaining) rows of the query result and returns them as a list of tuples.

Prototype

```
cursor.fetchall()
```

Parameter

None

Return Value

Tuple list, which contains all results of the result set. An empty list is returned when no rows are available.

Examples

For details, see [Example: Common Operations](#).

5.6.4.9 cursor.close()

Function

This method closes the cursor of the current connection.

Prototype

```
cursor.close()
```

Parameter

None

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.6.4.10 connection.close()

Function

This method closes the database connection.



CAUTION

This method closes the database connection and does not automatically call **commit()**. If you just close the database connection without calling **commit()** first, changes will be lost.

Prototype

```
connection.close()
```

Parameter

None

Return Value

None

Examples

For details, see [Example: Common Operations](#).

5.7 Development Based on the Go Driver

5.7.1 Go Driver Package, Environment Class, and Driver Class

Go Driver Package

Obtain the **GaussDB-Kernel-*VxxxRxxxCxx-OS information*-64bit-Go.tar.gz** package from the driver package. Decompress the package to obtain the Go driver source code package.

NOTICE

The Go driver package provided by the database depends on Go 1.13 or later.

Environment Class

- **Configure the Go environment.**
- You need to configure the following parameters in the environment variables:
 - **GO111MODULE**: Set **GO111MODULE** to **on** when installing the Go driver by importing a file online. If you do not want to reconstruct the **go mod** project, set **GO111MODULE** to **off** and manually download the dependency package. The dependency package must be at the same level as the driver root directory and service code.
 - **GOPROXY**: When importing data online, you need to configure the path that contains the Go driver package.
 - You can configure other Go environment variables based on your scenario parameters.

Run the **go env** command to view the Go environment variable configuration result and check whether the Go version is 1.13 or later.

- **Install the Go driver.**
 - Obtain the Go driver package **GaussDB-Kernel_*Database version number_OS information*_64bit_Go.tar.gz** from the driver package and save it to the local PC. Decompress the package to obtain the Go driver source code package.
 - Go to the root path of the Go driver code and run the **go mod tidy** command to download related dependencies. You need to configure **GOPATH=\${Path for storing the Go driver dependency package}** in the environment variables.
 - If the dependencies have been downloaded to the localhost, you can add a line "Replace the Go driver package with the local Go driver package address through replace" to **go.mod**, indicating that all import Go driver packages in the code are stored in the local path and the dependencies are not downloaded from the proxy.

 CAUTION

When you run the **go mod tidy** command to download dependencies, some of them may be downloaded as an earlier version. If the earlier version has vulnerabilities, you can change the dependency version in the **go.mod** file and update the dependency to the version after the vulnerability is fixed to avoid risks.

Driver Class

When creating a database connection, you need to enter the database driver name **opengauss**.


NOTICE

- The Go driver provided by the database does not adapt to mature ORM frameworks (such as XORM) in the industry. As such, the driver name input during database connection creation must be compatible with Postgres and PostgreSQL.
 - The Go driver of the database cannot coexist with that of PostgreSQL.
-

5.7.2 Development Process

The Go driver of the database complies with the rule of the Go language third-party library. You only need to import the driver to the application program and save the driver code in the directory specified by *GOPATH*.

Figure 5-6 Application development process based on Go



5.7.3 Connecting to the Database

When you call the standard SQL API **open** of the Go language to create a database connection, a connected object is returned to transfer the driver name and description string.

Function Prototype

The Go driver provides the following method to generate a database connected object:

```
func Open(driverName, dataSourceName string) (*DB, error)
```

Parameter description:

- **driverName** indicates the driver name. The database driver name is **opengauss**, which is compatible with Postgres.
- **dataSourceName** indicates the data source to be connected. The value can be in DSN or URL format.
 - DSN format: key1 = value1 key2 = value2... Different groups of keywords are separated by space. The space on the left and right of the equal sign (=) is optional.
 - URL format: driverName://[userspec@][hostspec]/[dbname][?paramspec]

In the preceding information, **driverName** indicates the driver name. The database driver name is **opengauss**, which is compatible with Postgres and PostgreSQL.

userspec indicates user[:password]. When a URL is used for connection, the password cannot contain separators in the URL string. If the password contains separators, the DSN format is recommended.

hostspec indicates [host][:port][, ...].

dbname indicates the database name. Note: The initial user cannot be used for remote login. **paramspec** indicates name=value[&...].

NOTICE

- In the DSN format, if there are multiple IP addresses:
 - When the value of **num(ip)** is the same as that of **num(port)**, the IP address matches the port number.
 - When the value of **num(ip)** is greater than that of **num(port)**, the IP address that cannot match the port number matches the first port number. For example, the matching condition of **host = ip1, ip2, ip3 port = port1, port2** is **ip1:port1, ip2:port2, ip3:port1**.
 - When the value of **num(ip)** is smaller than that of **num(port)**, the extra port numbers are discarded. For example, the mapping result of **host = ip1, ip2, ip3 port = port1, port2, port3, port4** is **ip1:port1, ip2:port2, ip3:port3**.
- In the URL format, if there are multiple IP addresses:
 - In the URL, *ip:port* must appear in pairs, that is, the value of **num(ip)** is the same as that of **num(port)**. Use commas (,) to separate multiple pairs. Example: **opengauss://user:password@ip1:port1, ip2:port2, ip3:port3/postgres**.
 - The URL contains only multiple IP addresses. The port number is specified by the environment variable or uses the default value **5432**. For example, in the case of **opengauss://user:password@ip1, ip2, ip3/postgres**, if the environment variable is set as **PGPORT = "port1, port2"**, the mapping is **ip1:port1, ip2:port2, ip3:port1**. If the environment variable is not set, the mapping is **ip1:5432, ip2:5432, ip3:5432**.

Parameters

Table 5-89 Database connection parameters

Parameter	Parameters
host	IP address of the host server, which can also be specified by the environment variable <code>\${PGHOST}</code>
port	Port number of the host server, which can also be specified by the environment variable <code>\${PGPORT}</code>
dbname	Database name, which can also be specified by the environment variable <code>\${PGDATABASE}</code>
user	Username to be connected, which can also be specified by the environment variable <code>\${PGUSER}</code>
password	Password of the user to be connected
connect_timeout	Timeout interval for connecting to the server, which can also be specified by the environment variable <code>\${PGCONNECT_TIMEOUT}</code>

sslmode	SSL encryption mode, which can also be specified by the environment variable <code>PGSSLMODE</code> Value range: • disable: SSL connection is disabled. • allow: If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified. • prefer: If the database supports SSL connection, SSL connection is preferred. However, authenticity of the database server will not be verified. • require: SSL connection is required and data is encrypted. However, authenticity of the database server will not be verified. • verify-ca: SSL connection is required, and whether the server certificate is issued by a trusted CA is verified. • verify-full: SSL connection is required, and whether the server certificate is issued by a trusted CA and whether the host name of the server is the same as that in the certificate are verified.
sslkey	Key location of the client certificate. If SSL connection is required and this parameter is not specified, you can set the environment variable <code>PGSSLKEY</code> to specify the location.
sslcert	File name of the client SSL certificate, which can also be specified by the environment variable <code>PGSSLCERT</code>
sslrootcert	Name of the file that contains the SSL CA certificate, which can also be specified by the environment variable <code>PGSSLROOTCERT</code>
sslcrll	File name of the SSL CRL. If a certificate listed in this file exists, the server certificate authentication will be rejected and the connection will fail. The value can also be specified by the environment variable <code>PGSSLCRL</code>.
sslpassword	Passphrase used to decrypt a key into plaintext. If this parameter is specified, the SSL key is an encrypted file. Currently, the SSL key supports DES encryption and AES encryption. NOTE The DES encryption algorithm has lower security and poses security risks. Therefore, you are advised to use a more secure encryption algorithm.

disable_prepared_binary_result	The value of this parameter is a string. If it is set to yes, the connection should not use the binary format when the query results are received from prepared statements. This parameter is used only for debugging. Value range: yes and no.
binary_parameters	Specifies whether []byte is always sent in binary format. The value is a string. Value range: yes and no. If this parameter is set to yes, you are advised to bind parameters based on [] byte to reduce internal type conversion.
target_session_attrs	Connection type of the database, which can also be specified by the environment variable <code>\${PGTARGETSESSIONATTRS}</code>. This parameter is used to identify the primary and standby nodes. There are six value options, namely, any, master, slave, preferSlave, read-write, and read-only. The default value is any. • any: attempts to connect to any DN in the URL connection string. • master: attempts to connect to a primary node in the URL connection string. If the primary node cannot be found, an exception is thrown. • slave: attempts to connect to a standby node in the URL connection string. If the standby node cannot be found, an exception is thrown. • preferSlave: attempts to connect to a standby DN (if available) in the URL connection string. Otherwise, it connects to the primary DN. • read-write: specifies that only the primary node can be connected. • read-only: specifies that only the standby node can be connected.
loggerLevel	Log level, which is used to print debugging information. The value can also be specified by the environment variable <code>\${PGLOGGERLEVEL}</code>. The value can be trace, debug, info, warn, error, or none, in descending order of priority.
application_name	Name of the Go driver that is being connected. The default value is go-driver. You are advised not to configure this parameter.

RuntimeParams	Value of the GUC parameter of the set type that is run by default when a session is connected. for example, search_path , application_name , or timezone . For details about the parameters, see the default settings of the client connection. You can run the SHOW command to check whether the parameters are set successfully.

Example 1:

```
// Single IP address and port (ip:port) are used as an example. In this example, the username and password
// are stored in environment variables. Before running this example, set environment variables in the local
// environment (set the environment variable name based on the actual situation).
func main() {
    hostip := os.Getenv("GOHOSTIP") // GOHOSTIP indicates the IP address written into the environment
    // variable.
    port := os.Getenv("GOPORT") // GOPORT indicates the port number written into the environment
    // variable.
    username := os.Getenv("GOUSRNAME") // GOUSRNAME indicates the username for writing environment
    // variables.
    passwd := os.Getenv("GOPASSWD") // GOPASSWD indicates the user password written into the
    // environment variable.
    str := "host=" + hostip + " port=" + port + " user=" + username + " password=" + passwd + "
    dbname=postgres sslmode=disable" // DSN connection string
    //str := "opengauss://" + username + ":" + passwd + "@" + hostip + ":" + port + "/postgres?
    //sslmode=disable" // URL connection string
    db, err := sql.Open("opengauss", str)
    if err != nil {
        log.Fatal(err)
    }
    defer db.Close()

    err = db.Ping()
    if err != nil {
        log.Fatal(err)
    }

    sqls := []string{
        "drop table if exists testExec",
        "create table testExec(f1 int, f2 varchar(20), f3 number, f4 timestamp, f5 boolean)",
        "insert into testExec values(1, 'abcdefg', 123.3, '2022-02-08 10:30:43.31 +08', true)",
        "insert into testExec values(:f1, :f2, :f3, :f4, :f5)",
    }

    inF1 := []int{2, 3, 4, 5, 6}
    inF2 := []string{"hello world", "huawei", "beijing", "nanjing", "yanjiusuo"}
    inF3 := []float64{641.43, 431.54, 5423.52, 665537.63, 6503.1}
    inF4 := []time.Time{
        time.Date(2022, 2, 8, 10, 35, 43, 623431, time.Local),
        time.Date(2022, 2, 10, 19, 11, 54, 353431, time.Local),
        time.Date(2022, 2, 12, 6, 11, 15, 636431, time.Local),
        time.Date(2022, 2, 14, 4, 51, 22, 747653, time.Local),
        time.Date(2022, 2, 16, 13, 45, 55, 674636, time.Local),
    }
    inF5 := []bool{false, true, false, true, true}

    for _, s := range sqls {
        if strings.Contains(s, ":f") {
            for i, _ := range inF1 {
                _, err := db.Exec(s, inF1[i], inF2[i], inF3[i], inF4[i], inF5[i])
                if err != nil {
                    log.Fatal(err)
                }
            }
        }
    }
}
```

```
} else {
    _, err = db.Exec(s)
    if err != nil {
        log.Fatal(err)
    }
}
}

var f1 int
var f2 string
var f3 float64
var f4 time.Time
var f5 bool
err = db.QueryRow("select * from testExec").Scan(&f1, &f2, &f3, &f4, &f5)
if err != nil {
    log.Fatal(err)
} else {
    fmt.Printf("f1:%v, f2:%v, f3:%v, f4:%v, f5:%v\n", f1, f2, f3, f4, f5)
}

row, err := db.Query("select * from testExec where f1 > :1", 1)
if err != nil {
    log.Fatal(err)
}
defer row.Close()

for row.Next() {
    err = row.Scan(&f1, &f2, &f3, &f4, &f5)
    if err != nil {
        log.Fatal(err)
    } else {
        fmt.Printf("f1:%v, f2:%v, f3:%v, f4:%v, f5:%v\n", f1, f2, f3, f4, f5)
    }
}
}
```

Example 2:

// Multiple IP addresses and ports (ip:port) are used as an example. In this example, the username and password are stored in environment variables. Before running this example, set environment variables in the local environment (set the environment variable name based on the actual situation).

```
func main() {
    ctx := context.Background()
    ctx2SecondTimeout, cancelFunc2SecondTimeout := context.WithTimeout(ctx, 2*time.Second)
    defer cancelFunc2SecondTimeout()

    hostip1 := os.Getenv("GOHOSTIP1") // GOHOSTIP1 indicates the IP address written into the environment variable.
    hostip2 := os.Getenv("GOHOSTIP2") // GOHOSTIP2 indicates the IP address written into the environment variable.
    hostip3 := os.Getenv("GOHOSTIP3") // GOHOSTIP3 indicates the IP address written into the environment variable.
    port1 := os.Getenv("GOPORT1")     // GOPORT1 indicates the port number written into the environment variable.
    port2 := os.Getenv("GOPORT2")     // GOPORT2 indicates the port number written into the environment variable.
    usrname := os.Getenv("GOUSRNAME") // GOUSRNAME indicates the username for writing environment variables.
    passwd := os.Getenv("GOPASSWD")   // GOPASSWDW indicates the user password written into the environment variable.

    str := "host="+hostip1+","+hostip2+","+hostip3+" port="+port1+","+port2+" user="+usrname+"
    password="+passwd+" dbname=postgres sslmode=disable" // DSN connection string.
    //str := "opengauss://" +usrname+":"+passwd+"@"+hostip1+":"+port1+","+hostip2+":"+port2+","+hostip3+"/
    postgres?sslmode=disable" // URL connection string.
    db, err := sql.Open("opengauss", str)
    if err != nil {
        log.Fatal(err)
    }
}
```

```
defer db.Close()

// Ping database connection with 2 second timeout
err = db.PingContext(ctx2SecondTimeout)
if err != nil {
    log.Fatal(err)
}

sqls := []string{
    "drop table if exists testExecContext",
    "create table testExecContext(f1 int, f2 varchar(20), f3 number, f4 timestampz, f5 boolean)",
    "insert into testExecContext values(1, 'abcdefg', 123.3, '2022-02-08 10:30:43.31 +08', true)",
    "insert into testExecContext values(:f1, :f2, :f3, :f4, :f5)",
}

inF1 := []int{2, 3, 4, 5, 6}
intF2 := []string{"hello world", "Huawei", "Beijing 2022 Winter Olympics", "nanjing", "Research Center"}
intF3 := []float64{641.43, 431.54, 5423.52, 665537.63, 6503.1}
intF4 := []time.Time{
    time.Date(2022, 2, 8, 10, 35, 43, 623431, time.Local),
    time.Date(2022, 2, 10, 19, 11, 54, 353431, time.Local),
    time.Date(2022, 2, 12, 6, 11, 15, 636431, time.Local),
    time.Date(2022, 2, 14, 4, 51, 22, 747653, time.Local),
    time.Date(2022, 2, 16, 13, 45, 55, 674636, time.Local),
}
intF5 := []bool{false, true, false, true, true}

for _, s := range sqls {
    if strings.Contains(s, ":") {
        for i, _ := range inF1 {
            _, err := db.ExecContext(ctx2SecondTimeout, s, inF1[i], intF2[i], intF3[i], intF4[i], intF5[i])
            if err != nil {
                log.Fatal(err)
            }
        }
    } else {
        _, err = db.ExecContext(ctx2SecondTimeout, s)
        if err != nil {
            log.Fatal(err)
        }
    }
}

var f1 int
var f2 string
var f3 float64
var f4 time.Time
var f5 bool
err = db.QueryRowContext(ctx2SecondTimeout, "select * from testExecContext").Scan(&f1, &f2, &f3, &f4, &f5)
if err != nil {
    log.Fatal(err)
} else {
    fmt.Printf("f1:%v, f2:%v, f3:%v, f4:%v, f5:%v\n", f1, f2, f3, f4, f5)
}

row, err := db.QueryContext(ctx2SecondTimeout, "select * from testExecContext where f1 > :1", 1)
if err != nil {
    log.Fatal(err)
}
defer row.Close()

for row.Next() {
    err = row.Scan(&f1, &f2, &f3, &f4, &f5)
    if err != nil {
        log.Fatal(err)
    } else {
        fmt.Printf("f1:%v, f2:%v, f3:%v, f4:%v, f5:%v\n", f1, f2, f3, f4, f5)
    }
}
```

```
}  
}
```

5.7.4 Connecting to the Database (Using SSL)

The Go driver supports SSL connections to the database. After the SSL mode is enabled, if the Go driver connects to the database server in SSL mode, the Go driver uses the standard TLS 1.3 protocol by default, and the TLS version must be 1.2 or later. This section describes how applications configure the client in SSL mode through Go. Contact an administrator for details about the configuration on the server. To enable SSL connections, you must have the server certificate, client certificate, and private key files. For details on how to obtain these files, see related documents and commands of OpenSSL.

NOTE

In SSL-based certificate authentication mode, you do not need to specify the user password in the connection string.

Configuring the Client

Upload the certificate files **client.key**, **client.crt**, and **ca.crt** generated on the server to the client. Contact an administrator for details about the configuration on the server.

Example 1:

```
// Mutual authentication is used as an example. In this example, the username and password are stored in  
// environment variables. Before running this example, set environment variables in the local environment (set  
// the environment variable name based on the actual situation).  
func main() {  
    hostip := os.Getenv("GOHOSTIP") // GOHOSTIP indicates the IP address written into the environment  
    // variable.  
    port := os.Getenv("GOPORT") // GOPORT indicates the port number written into the environment  
    // variable.  
    username := os.Getenv("GOUSRNAME") // GOUSRNAME indicates the username for writing environment  
    // variables.  
    passwd := os.Getenv("GOPASSWD") // GOPASSWD indicates the user password written into the  
    // environment variable.  
    sslpasswd := os.Getenv("GOSSLPASSWD") // GOSSLPASSWD indicates the passphrase written into the  
    // environment variable.  
    dsnStr := "host=" + hostip + " port=" + port + " user=" + username + " password=" + passwd + "  
    sslcert=certs/client.crt sslkey=certs/client.key sslpassword=" + sslpasswd  
    parameters := []string{  
        "sslmode=require",  
        "sslmode=verify-ca sslrootcert=certs/ca.crt",  
    }  
  
    for _, param := range parameters {  
        db, err := sql.Open("opengauss", dsnStr+param)  
        if err != nil {  
            log.Fatal(err)  
        }  
    }  
  
    var f1 int  
    err = db.QueryRow("select 1").Scan(&f1)  
    if err != nil {  
        log.Fatal(err)  
    } else {  
        fmt.Printf("RESULT: select 1: %d\n", f1)  
    }  
  
    db.Close()
```

```
}  
}
```

Example 2:

```
// For example, verify sslpassword. In this example, the username and password are stored in environment  
// variables. Before running this example, set environment variables in the local environment (set the  
// environment variable name based on the actual situation).  
func main() {  
    hostip := os.Getenv("GOHOSTIP") // GOHOSTIP indicates the IP address written into the environment  
    // variable.  
    port := os.Getenv("GOPORT") // GOPORT indicates the port number written into the environment  
    // variable.  
    username := os.Getenv("GOUSRNAME") // GOUSRNAME indicates the username for writing environment  
    // variables.  
    passwd := os.Getenv("GOPASSWD") // GOPASSWD indicates the user password written into the  
    // environment variable.  
    dsnStr := "host=" + hostip + " port=" + port + " user=" + username + " password=" + passwd + "  
    dbname=postgres"  
  
    sslpasswd := os.Getenv("GOSSLPASSWD") // GOSSLPASSWD indicates the passphrase written into the  
    // environment variable.  
    connStrs := []string{  
        " sslmode=verify-ca sslcert=certs/client_rsa.crt sslkey=certs/client_rsa.key sslpassword=" + sslpasswd + "  
    sslrootcert=certs/cacert_rsa.pem",  
        " sslmode=verify-ca sslcert=certs/client_ecdsa.crt sslkey=certs/client_ecdsa.key sslpassword=" + sslpasswd + "  
    sslrootcert=certs/cacert_ecdsa.pem",  
    }  
    for _, connStr := range connStrs {  
        db, err := sql.Open("opengauss", dsnStr+connStr)  
        if err != nil {  
            log.Fatal(err)  
        }  
        var f1 int  
        err = db.QueryRow("select 1").Scan(&f1)  
        if err != nil {  
            if !strings.HasPrefix(err.Error(), "connect failed.") {  
                log.Fatal(err)  
            }  
        }  
        db.Close()  
    }  
}
```

5.7.5 Go API Reference

5.7.5.1 sql.Open

The following table describes sql.Open.

Method	Description	Return Value
Open(driverName, dataSourceName string)	Opens a database based on a specified database driver and the dedicated data source of the driver.	*DB and error

For details about the **driverName** and **dataSourceName** parameters, see [Connecting to the Database](#).

5.7.5.2 type DB

The following table describes type DB.

Method	Description	Return Value
(db *DB)Begin()	Starts a transaction. The isolation level of the transaction is determined by the driver.	*Tx and error
(db *DB)BeginTx(ctx context.Context, opts *TxOptions)	Starts a transaction with a specified transaction isolation level. A specified context is used until the transaction is committed or rolled back. If the context is canceled, the SQL package rolls back the transaction.	*Tx and error
(db *DB)Close()	Closes the database and releases all the opened resources.	error
(db *DB)Exec(query string, args ...interface{})	Performs an operation that does not return rows of data.	Result and error
(db *DB)ExecContext(ctx context.Context, query string, args ...interface{})	Performs an operation that does not return rows of data in a specified context.	Result and error
(db *DB)Ping()	Checks whether the database connection is still valid and establishes a connection if necessary.	error
(db *DB)PingContext(ctx context.Context)	Checks whether the database connection is still valid in a specified context and establishes a connection if necessary.	error
(db *DB)Prepare(query string)	Creates a prepared statement for subsequent queries or executions.	*Stmt and error
(db *DB)PrepareContext(ctx context.Context, query string)	Creates a prepared statement for subsequent queries or executions in a specified context.	*Stmt and error
(db *DB)Query(query string, args ...interface{})	Executes a query and returns multiple rows of data.	*Rows and error
(db *DB)QueryContext(ctx context.Context, query string, args ...interface{})	Executes a query and returns multiple rows of data in a specified context.	*Rows and error

(db *DB)QueryRow(query string, args ...interface{})	Executes a query that returns only one row of data.	*Row
(db *DB)QueryRowContext(ctx context.Context, query string, args ...interface{})	Executes a query that returns only one row of data in a specified context.	*Row

Parameters

Parameter	Description
ctx	Specified context
query	Executed SQL statement
args	Parameter that needs to be bound to the executed SQL statement. Binding by location and binding by name are supported. For details, see the following example.
opts	Transaction isolation level and transaction access mode. The value of the transaction isolation level (opts.Isolation) can be sql.LevelReadUncommitted, sql.LevelReadCommitted, sql.LevelRepeatableRead, or sql.LevelSerializable. The value of the transaction access mode (opts.ReadOnly) can be true (read-only) or false (read/write).

Example:

```
// In this example, the username and password are stored in environment variables. Before running this
// example, set environment variables in the local environment (set the environment variable names based on
// the actual situation).
func main() {
    hostip := os.Getenv("GOHOSTIP") // GOHOSTIP indicates the IP address written into the environment
    // variable.
    port := os.Getenv("GOPORT") // GOPORT indicates the port number written into the environment
    // variable.
    username := os.Getenv("GOUSRNAME") // GOUSRNAME indicates the username for writing environment
    // variables.
    passwd := os.Getenv("GOPASSWD") // GOPASSWD indicates the user password written into the
    // environment variable.
    str := "host=" + hostip + " port=" + port + " user=" + username + " password=" + passwd + "
    dbname=postgres sslmode=disable"
    db, err := sql.Open("opengauss", str)
    if err != nil {
        log.Fatal(err)
    }
    defer db.Close()

    err = db.Ping()
    if err != nil {
        log.Fatal(err)
    }
}
```

```
// Binding by location
_, err = db.Exec("insert into test(id, name) values(:1, :2)", 1, "Zhang San")
if err != nil {
    log.Fatal(err)
}

// Binding by name
_, err = db.Exec("insert into test(id, name) values(:id, :name)", sql.Named("id", 1), sql.Named("name",
"Zhang San"))
if err != nil {
    log.Fatal(err)
}
}
```

5.7.5.3 type Stmt

The following table describes type Stmt.

Method	Description	Return Value
(s *Stmt)Close()	Closes a specified prepared statement.	error
(s *Stmt)Exec(args ...interface{})	Executes a prepared statement with specified parameters and returns a Result value.	Result and error
(s *Stmt)ExecContext(ctx context.Context, args ...interface{})	Executes a prepared statement with specified parameters in a specified context and returns a Result value.	Result and error
(s *Stmt)Query(args ...interface{})	Executes a prepared statement with specified parameters and returns *Rows as the query result.	*Rows and error
(s *Stmt)QueryContext(ctx context.Context, args ...interface{})	Executes a prepared statement with specified parameters in a specified context and returns *Rows as the query result.	*Rows and error
(s *Stmt)QueryRow(args ...interface{})	Executes a prepared statement with specified parameters and returns *Row as the result.	*Row
(s *Stmt)QueryRowContext (ctx context.Context, args ...interface{})	Executes a prepared statement with specified parameters in a specified context and returns *Row as the result.	*Row

Parameters

Parameter	Description
ctx	Specified context
query	Executed SQL statement
args	Parameter that needs to be bound to the executed SQL statement. Binding by location and binding by name are supported. For details, see the example in type DB.

5.7.5.4 type Tx

The following table describes type Tx.

Method	Description	Return Value
(tx *Tx)Commit()	Commits a transaction.	error
(tx *Tx)Exec(query string, args ...interface{})	Performs an operation that does not return rows of data.	Result and error
(tx *Tx)ExecContext(ctx context.Context, query string, args ...interface{})	Performs an operation that does not return rows of data in a specified context.	Result and error
(tx *Tx)Prepare(query string)	Creates a prepared statement for subsequent queries or executions. The returned statement is executed within a transaction and cannot be used when the transaction is committed or rolled back.	*Stmt and error

(tx *Tx)PrepareContext(ctx context.Context, query string)	Creates a prepared statement for subsequent queries or executions. The returned statement is executed within a transaction and cannot be used when the transaction is committed or rolled back. The specified context will be used in the preparation phase, not in the transaction execution phase. The statement returned by this method will be executed in the transaction context.	*Stmt and error
(tx *Tx)Query(query string, args ...interface{})	Executes a query that returns rows of data.	*Rows and error
(tx *Tx)QueryContext(ctx context.Context, query string, args ...interface{})	Executes a query that returns rows of data in a specified context.	*Rows and error
(tx *Tx)QueryRow(query string, args ...interface{})	Executes a query that returns only one row of data.	*Row
(tx *Tx)QueryRowContext(ctx context.Context, query string, args ...interface{})	Executes a query that returns only one row of data in a specified context.	*Row
(tx *Tx) Rollback()	Rolls back a transaction.	error
(tx *Tx)Stmt(stmt *Stmt)	Returns a transaction-specific prepared statement for an existing statement. Example: str, err := db.Prepare("insert into t1 values(:1, :2)") tx, err := db.Begin() res, err := tx.Stmt(str).Exec(1, "aaa")	*Stmt
(tx *Tx)StmtContext(ctx context.Context, stmt *Stmt)	Returns a transaction-specific prepared statement for an existing statement in a specified context.	*Stmt

Parameter Description

Parameter	Description
ctx	Specified context
query	Executed SQL statement
args	Parameter that needs to be bound to the executed SQL statement. Binding by location and binding by name are supported. For details, see the example in type DB.
stmt	Existing prepared statement, which is generally the prepared statement returned by the PREPARE statement

5.7.5.5 type Rows

The following table describes type Rows.

Method	Description	Return Value
(rs *Rows)Close()	Closes Rows to stop the iteration of the dataset.	error
(rs *Rows)ColumnTypes()	Returns column information.	[]*ColumnType and error
(rs *Rows)Columns()	Returns the name of each column.	[]string and error
(rs *Rows)Err()	Returns any errors that occur during iteration.	error
(rs *Rows)Next()	Prepares the next data row to be read with the Scan method. If there is an additional result set, true is returned. Otherwise, false is returned.	bool
(rs *Rows)Scan(dest ...interface{})	Copies the columns of the current iterated row of data to the value specified by dest .	error
(rs *Rows)NextResultSet() bool	Specifies whether there is an additional result set.	Bool

Parameter

Parameter	Description
dest	The column to be queried needs to be copied to the value specified by this parameter.

5.7.5.6 type Row

The following table describes type Row.

Method	Description	Return Value
(r *Row)Scan(dest ...interface{})	Copies the columns in the current row of data to the value specified by dest .	error
(r *Row)Err()	Returns errors that occur during execution.	error

Parameter Description

Parameter	Description
dest	The column to be queried needs to be copied to the value specified by this parameter.

5.7.5.7 type ColumnType

The following table describes type ColumnType.

Method	Description	Return Value
(ci *ColumnType)DatabaseTypeName()	Returns the name of the column-type database system. If an empty string is returned, driver-type names are not supported.	error
(ci *ColumnType)DecimalSize()	Returns the scale and precision of the decimal type. If the value of ok is false , the specified type is unavailable or not supported.	precision, scale int64, ok bool

(ci *ColumnType)Length()	Returns the length of the data column type. If the value of ok is false , the specified type does not have a length.	length int64, ok bool
(ci *ColumnType)ScanType()	Returns a Go type that can be used for scanning by using Rows.Scan.	reflect.Type
(ci *ColumnType)Name()	Returns the name of a data column.	string

5.7.5.8 type Result

The following table describes type Result.

Method	Description	Return Value
(res Result)RowsAffected()	Returns the number of rows affected by the INSERT, DELETE, UPDATE, SELECT, MOVE, FETCH, and COPY operations.	int64 and error

5.8 Appendix

5.8.1 JDBC

5.8.1.1 Data Type Mapping

The relationships among data types, Java variable types, and JDBC type indexes are as follows (A: Oracle-compatible; B: MySQL-compatible):

Table 5-90 Mapping for JDBC Data Types

Compatibility Mode	Gauss Kernel Data Type	Java Variable Type	JDBC Type Index
A/B	oid	java.lang.Long	java.sql.Types.BIGINT
A/B	numeric	java.math.BigDecimal	java.sql.Types.NUMERIC
A/B	tinyint	java.lang.Integer	java.sql.Types.TINYINT

Compatibility Mode	Gauss Kernel Data Type	Java Variable Type	JDBC Type Index
A/B	smallint	java.lang.Integer	java.sql.Types.SMALLINT
A/B	bigint	java.lang.Long	java.sql.Types.BIGINT
A/B	float4	java.lang.Float	java.sql.Types.REAL
A/B	float8	java.lang.Double	java.sql.Types.DOUBLE
A/B	char	java.lang.String	java.sql.Types.CHAR
A/B	character	java.lang.String	java.sql.Types.CHAR
A/B	bpchar	java.lang.String	java.sql.Types.CHAR
A/B	character varying	java.lang.String	java.sql.Types.VARCHAR
A/B	varchar	java.lang.String	java.sql.Types.VARCHAR
A/B	text	java.lang.String	java.sql.Types.VARCHAR
A/B	name	java.lang.String	java.sql.Types.VARCHAR
A/B	bytea	byte[]	java.sql.Types.BINARY
A/B	blob	java.sql.Blob	java.sql.Types.BLOB
A/B	clob	java.sql.Clob	java.sql.Types.CLOB
A/B	bool	java.lang.Boolean	java.sql.Types.BIT
B	date	java.sql.Date	java.sql.Types.DATE
A/B	time	java.sql.Time	java.sql.Types.TIME
A/B	timetz	java.sql.Time	java.sql.Types.TIME
A/B	timestamp	java.sql.Timestamp	java.sql.Types.TIMESTAMP
A/B	smalldatetime	java.sql.Timestamp	java.sql.Types.TIMESTAMP
A/B	timestampz	java.sql.Timestamp	java.sql.Types.TIMESTAMP
A/B	refcursor	java.sql.ResultSet	java.sql.Types.REF_CURSOR java.sql.Types.OTHER

5.8.1.2 Log Management

The GaussDB JDBC driver uses log records to help solve problems when the GaussDB JDBC driver is used in applications. GaussDB JDBC supports the following log management methods:

1. Use the SLF4J log framework for interconnecting with applications.
2. Use the JdkLogger log framework for interconnecting with applications.

SLF4J and JdkLogger are mainstream frameworks for Java application log management in the industry. For details about how to use these frameworks, see the official documents (SLF4J: <http://www.slf4j.org/manual.html>; JdkLogger: <https://docs.oracle.com/javase/8/docs/technotes/guides/logging/overview.html>).

Method 1: Use the SLF4J log framework for interconnecting with applications.

When a connection is set up, **logger=Slf4JLogger** is configured in the URL.

The SLF4J may be implemented by using Log4j or Log4j2. When the Log4j is used to implement the SLF4J, the following JAR packages need to be added: **log4j-*.jar**, **slf4j-api-*.jar**, and **slf4j-log4j-*.jar** (* varies according to versions), and configuration file **log4j.properties**. In addition, the code for reading the configuration file must be added to the main program. If the Log4j2 is used to implement the SLF4J, you need to add the following JAR packages: **log4j-api-*.jar**, **log4j-core-*.jar**, **log4j-slf4j18-impl-*.jar**, and **slf4j-api-*-alpha1.jar** (* varies according to versions), and configuration file **log4j2.xml**.

This method supports log management and control. The SLF4J can implement powerful log management and control functions through related configurations in files. This method is recommended.

CAUTION

This method depends on the general SLF4J APIs, such as **org.slf4j.LoggerFactory.getLogger(String name)**, **org.slf4j.Logger.debug(String var1)**, **org.slf4j.Logger.info(String var1)**, **org.slf4j.Logger.warn(String warn)**, and **org.slf4j.Logger.warn(String warn)**. If these APIs are changed, logs cannot be recorded.

Example:

```
// Note: The path of the log4j.properties or log4j2.xml file must be specified during calling.

// Use the log4j.properties file to configure the attributes of the log recorder, such as the log level and
// output target. These properties are loaded into the application by calling the configure method of
// PropertyConfigurator and passing the path of the log4j.properties file.
//PropertyConfigurator.configure("log4j.properties");

// Create the Log4j2 configuration source, specify the path of the log4j2.xml configuration file, and
// initialize the log system.
//ConfigurationSource source = new ConfigurationSource(new FileInputStream("log4j2.xml"));
//Configurator.initialize(null, source);

public static Connection GetConnection(String username, String passwd){

    String sourceURL = "jdbc:opengauss://$ip:$port/database?logger=Slf4JLogger";
    Connection conn = null;

    try{
        // Create a connection.
        conn = DriverManager.getConnection(sourceURL,username,passwd);
        System.out.println("Connection succeed!");
    }catch (Exception e){
        e.printStackTrace();
    }
```

```
        return null;
    }
    return conn;
}
```

log4j.properties configuration file:

```
log4j.logger.com.huawei.opengauss.jdbc=ALL, log_gsjdbc

# Default file output configuration
log4j.appender.log_gsjdbc=org.apache.log4j.RollingFileAppender
log4j.appender.log_gsjdbc.Append=true
log4j.appender.log_gsjdbc.File=gsjdbc.log
log4j.appender.log_gsjdbc.Threshold=TRACE
log4j.appender.log_gsjdbc.MaxFileSize=10MB
log4j.appender.log_gsjdbc.MaxBackupIndex=5
log4j.appender.log_gsjdbc.layout=org.apache.log4j.PatternLayout
log4j.appender.log_gsjdbc.layout.ConversionPattern=%d %p %t %c - %m%n
log4j.appender.log_gsjdbc.File.Encoding = UTF-8
```

log4j2.xml configuration file:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration status="OFF">
    <appenders>
        <Console name="Console" target="SYSTEM_OUT">
            <PatternLayout pattern="%d %p %t %c - %m%n"/>
        </Console>
        <File name="FileTest" fileName="test.log">
            <PatternLayout pattern="%d %p %t %c - %m%n"/>
        </File>
        <!-- JDBC driver log file output configuration. Log rewinding is supported. When the log size exceeds
        10 MB, a new file is created. The new file is named in the format of yyyy-mm-dd-file ID. -->
        <RollingFile name="RollingFileJdbc" fileName="gsjdbc.log" filePattern="%d{yyyy-MM-dd}-%i.log">
            <PatternLayout pattern="%d %p %t %c - %m%n"/>
            <Policies>
                <SizeBasedTriggeringPolicy size="10 MB"/>
            </Policies>
        </RollingFile>
    </appenders>
    <loggers>
        <root level="all">
            <appender-ref ref="Console"/>
            <appender-ref ref="FileTest"/>
        </root>
        <!-- JDBC driver logs. The log level is all. All logs can be viewed and exported to the gsjdbc.log file. -->
        <logger name="com.huawei.opengauss.jdbc" level="all" additivity="false">
            <appender-ref ref="RollingFileJdbc"/>
        </logger>
    </loggers>
</configuration>
```

NOTICE

- In the configuration file, you can specify the configuration paths for JDBC and upper-layer services to store JDBC logs and service logs to different files.
- When multiple GaussDB data sources are configured in the same project, if **logger** is set to **Slf4JLogger** in some services, it will affect the logging configuration of each other. In this case, you are advised to configure the connection strings in a unified manner.

Method 2: Use the JdkLogger log framework for interconnecting with applications. You can use the **logging.properties** configuration file or configure the parameters in the sample code.

1. The default Java logging framework stores its configurations in a file named **logging.properties**. Java installs the global configuration file in the folder in the Java installation directory. The **logging.properties** file can also be created and stored with a single project.

logging.properties configuration file:

```
# Specify the processing program as a file.
handlers= java.util.logging.FileHandler

# Specify the default global log level.
.level= ALL

# Specify the log output control standard.
java.util.logging.FileHandler.level=ALL
java.util.logging.FileHandler.pattern = gsjdbc.log
java.util.logging.FileHandler.limit = 500000
java.util.logging.FileHandler.count = 30
java.util.logging.FileHandler.formatter = java.util.logging.SimpleFormatter
java.util.logging.FileHandler.append=false
```

2. Configure the parameters in the sample code:

```
System.setProperty("java.util.logging.FileHandler.pattern","jdbc.log");
FileHandler fileHandler = new FileHandler(System.getProperty("java.util.logging.FileHandler.pattern"));
Formatter formatter = new SimpleFormatter();
fileHandler.setFormatter(formatter);
Logger logger = Logger.getLogger("com.huawei.opengauss.jdbc");
logger.addHandler(fileHandler);
logger.setLevel(Level.ALL);
logger.setUseParentHandlers(false);
```

Link Trace Function

The GaussDB JDBC driver provides the application-to-database link trace function to associate discrete SQL statements on the database side with application requests. This function requires application developers to implement the `com.huawei.opengauss.jdbc.log.Tracer` API class and specify the full name of the API implementation class in the URL.

URL example:

```
String URL = "jdbc:opengauss://$ip:$port/database?traceInterfaceClass=xxx.xxx.xxx.OpenGaussTraceImpl";
```

The `com.huawei.opengauss.jdbc.log.Tracer` API class is defined as follows:

```
public interface Tracer {
    // Retrieves the value of traceId.
    String getTraceId();
}
```

The following is an example of the `com.huawei.opengauss.jdbc.log.Tracer` API implementation class:

```
import com.huawei.opengauss.jdbc.log.Tracer;

public class OpenGaussTraceImpl implements Tracer {
    private static MDC mdc = new MDC();

    private final String TRACE_ID_KEY = "traceId";

    public void set(String traceId) {
        mdc.put(TRACE_ID_KEY, traceId);
    }

    public void reset() {
        mdc.clear();
    }

    @Override
    public String getTraceId() {
```

```
        return mdc.get(TRACE_ID_KEY);
    }
}
```

The following is an example of context mapping which is used to store trace IDs generated for different requests:

```
import java.util.HashMap;

public class MDC {
    static final private ThreadLocal<HashMap<String, String>> threadLocal = new ThreadLocal<>();

    public void put(String key, String val) {
        if (key == null || val == null) {
            throw new IllegalArgumentException("key or val cannot be null");
        } else {
            if (threadLocal.get() == null) {
                threadLocal.set(new HashMap<>());
            }
            threadLocal.get().put(key, val);
        }
    }

    public String get(String key) {
        if (key == null) {
            throw new IllegalArgumentException("key cannot be null");
        } else if (threadLocal.get() == null) {
            return null;
        } else {
            return threadLocal.get().get(key);
        }
    }

    public void clear() {
        if (threadLocal.get() == null) {
            return;
        } else {
            threadLocal.get().clear();
        }
    }
}
```

Take the service **traceld** as an example. The prerequisite is that the table **test_trace_id (id int,name varchar(20))** is created.

```
String traceld = UUID.randomUUID().toString().replaceAll("-", "");
OpenGaussTraceImpl openGaussTrace = new OpenGaussTraceImpl();
openGaussTrace.set(traceld);
Connection con = DriverManager.getConnection(url, user, password);
pstmt = con.prepareStatement("select * from test_trace_id where id = ?");
pstmt.setInt(1, 1);
pstmt.execute();
pstmt = con.prepareStatement("insert into test_trace_id values(?,?)");
pstmt.setInt(1, 2);
pstmt.setString(2, "test");
pstmt.execute();
openGaussTrace.reset();
```

NOTE

- When the link trace function is used, the link function at the application layer is guaranteed by services.
- The application must expose the API for obtaining **traceld** to the JDBC and configure the API implementation class to the JDBC connection string.
- SQL statements of the same request must use the same **traceld**.
- The value of **traceld** transferred by the application cannot exceed 32 bytes. Otherwise, the extra bytes will be truncated.

5.8.1.3 Troubleshooting

5.8.1.3.1 Incorrect batchMode Settings

Symptom

Set the URL parameters **batchMode** to **on** and **reWriteBatchedInserts** to **true**. After JDBC is used to insert data in batches, a message is displayed, indicating that the number of binding parameters is inconsistent with the number of parameters required by the statement.

bind message supplies * parameters, but prepared statement "" requires *

Example 1:

```
// conn is a created connection object. The URL parameters for creating the connection contain
// &batchMode=on&reWriteBatchedInserts=true.
// Bind parameters in batches and then execute the statement. The number of bound parameters does not
// match the number of columns in the rewritten INSERT statement. As a result, an exception is reported.
// java.sql.BatchUpdateException: bind message supplies 3 parameters, but prepared statement "" requires 6
PreparedStatement stmt = conn.prepareStatement("insert into test_tbl values (?, ?, ?)");

stmt.setInt(1, 1);
stmt.setString(2, "aaa");
stmt.setString(3, "bbbb");
stmt.addBatch();

stmt.setInt(1, 1);
stmt.setString(2, "aaa");
stmt.setString(3, "bbbb");
stmt.addBatch();

stmt.setInt(1, 1);
stmt.setString(2, "aaa");
stmt.setString(3, "bbbb");
stmt.addBatch();

stmt.executeBatch();
```

Cause Analysis

When **reWriteBatchedInserts** is set to **true**, the batch statement combines multiple SQL statements into one. As a result, the number of reserved parameter columns in the statement changes. If **batchMode** is set to **on**, parameters are bound based on the SQL statements before combination. As a result, the number of bound parameters is inconsistent with the number of parameters required by the statement.

Solution

If **reWriteBatchedInserts** is set to **true**, **batchMode** is set to **off**.

5.8.1.3.2 Error Is Reported or Connection Is Blocked in SSL Mode

Symptom

When JDBC establishes a connection in SSL mode, a strong random number is obtained on the client. During the connection establishment, the following error information may be reported.

Scenario 1: Error report.

```
"Thread-0" #18 prio=5 os_prio=0 tid=0x00007f2ad0385000 nid=0x5429 runnable [0x00007f2aa069b000]
java.lang.Thread.State: RUNNABLE
    at java.io.FileInputStream.readBytes(Native Method)
    at java.io.FileInputStream.read(FileInputStream.java:255)
    at sun.security.provider.NativePRNG$RandomIO.readFully(NativePRNG.java:424)
    at sun.security.provider.NativePRNG$RandomIO.ensureBufferValid(NativePRNG.java:526)
    at sun.security.provider.NativePRNG$RandomIO.implNextBytes(NativePRNG.java:545)
    - locked <0x000000067273a950> (a java.lang.Object)
    at sun.security.provider.NativePRNG$RandomIO.access$400(NativePRNG.java:331)
    at sun.security.provider.NativePRNG$Blocking.engineNextBytes(NativePRNG.java:268)
    at java.security.SecureRandom.nextBytes(SecureRandom.java:468)
    at java.security.SecureRandom.next(SecureRandom.java:491)
    at java.util.Random.nextInt(Random.java:329)
    at sun.security.ssl.SSLContextImpl.engineInit(SSLContextImpl.java:106)
    at javax.net.ssl.SSLContext.init(SSLContext.java:282)
    at org.postgresql.ssl.LibPQFactory.<init>(LibPQFactory.java:175)
    at org.postgresql.core.SocketFactoryFactory.getSslSocketFactory(SocketFactoryFactory.java:62)
    at org.postgresql.ssl.MakeSSL.convert(MakeSSL.java:33)
    at org.postgresql.core.v3.ConnectionFactoryImpl.enableSSL(ConnectionFactoryImpl.java:723)
    at org.postgresql.core.v3.ConnectionFactoryImpl.tryConnect(ConnectionFactoryImpl.java:203)
    at org.postgresql.core.v3.ConnectionFactoryImpl.openConnectionImpl(ConnectionFactoryImpl.java:330)
    at org.postgresql.core.ConnectionFactory.openConnection(ConnectionFactory.java:58)
    at org.postgresql.jdbc.PgConnection.<init>(PgConnection.java:357)
```

Scenario 2: The connection is blocked. If **loginTimeout** is set in the connection string, the message "Connection attempt timed out" is displayed. If this parameter is not set, the connection is blocked.

Cause Analysis

The random number generation on the client is too slow to meet product requirements. The entropy source is insufficient. As a result, the service fails to be started. This problem exists in some Linux environments.

Solution

Solution 1: Start the haveged service on the client and increase the entropy value of the system entropy pool to improve the speed of reading random numbers. The startup command is as follows:

```
systemctl start haveged
```

Solution 2: Adjust the JDK configuration on the client.

Open the `$JAVA_PATH/jre/lib/security/java.security` file and modify the following configuration items:

```
securerandom.source=file:/dev/.urandom
securerandom.strongAlgorithms=NativePRNGNonBlocking:SUN
```

NOTICE

The essence of solution 2 is to use pseudo-random numbers instead of strong random numbers to reduce the entropy value to be consumed. All applications that use the JDK on the client are affected. Pseudo-random numbers are used to instead of strong random numbers.

5.8.2 libpq

5.8.2.1 Connection Parameters

Table 5-91 Connection parameters

Parameter	Description
host	Name of the host to connect to. If the host name starts with a slash (/), UDS communications instead of TCP/IP communications are used. The value is the directory where the socket file is stored. If host is not specified, the default behavior is to connect to the UDS in the /tmp directory (or the socket directory specified during database installation). On a machine without a UDS, the default behavior is to connect to localhost. You can specify multiple host names by using a character string separated by commas (.). Multiple host names can be specified.
hostaddr	IP address of the host to connect to. The value is in standard IPv4 address format, for example, 172.28.40.9. If the machine supports IPv6, IPv6 address can also be used. If a non-null string is specified, TCP/IP communications are used. You can specify multiple IP addresses by using a character string separated by commas (.). Multiple IP addresses can be specified. Replacing host with hostaddr can prevent applications from querying host names, which may be important for applications with time constraints. However, a host name is required for GSSAPI or SSPI authentication methods. Therefore, the following rules are used: 1. If host is specified but hostaddr is not, a query for the host name will be executed. 2. If hostaddr is specified but host is not, the value of hostaddr is the server network address. If the host name is required for authentication, the connection attempt fails. 3. If both host and hostaddr are specified, the value of hostaddr is the server network address. The value of host is ignored unless it is required by authentication, in which case it is used as the host name. NOTICE • If host is not the server name in the network address specified by hostaddr, the authentication may fail. • If neither host nor hostaddr is specified, libpq will use a local UDS for connection. If the machine does not have a UDS, it will attempt to connect to localhost.
port	Port number of the host server, or the socket file name extension for UDS connections. You can specify multiple port numbers by using a character string separated by commas (.). Multiple port numbers can be specified.
user	Name of the user to be connected. By default, the username is the same as the OS name of the user running the application.
dbname	Database name. The default value is the same as the username.

Parameter	Description
password	Password to be used if the server requires password authentication.
connect_timeout	Maximum timeout period of the connection, in seconds (written as a decimal integer string). The value 0 or null indicates infinity. You are advised not to set the connection timeout period to a value less than 2 seconds.
client_encoding	Client encoding for the connection. In addition to the values accepted by the corresponding server options, you can use auto to determine the correct encoding from the current environment in the client (the <i>LC_CTYPE</i> environment variable in the Unix system).
tty	This parameter can be ignored. (This parameter was used to specify the location to which the debugging output of the server was sent).
options	Adds command-line options to send to the server at runtime.
application_name	Current user identity.
fallback_application_name	Specifies a backup value for the application_name parameter. This value is used if no value is set for application_name through a connection parameter or the <i>PGAPPNAME</i> environment variable. In a common tool program, if you set a default name but do not want the default name to be overwritten by the user, you can specify a backup value.
keepalives	Specifies whether TCP keepalive is enabled on the client side. The default value is 1 , indicating that the function is enabled. The value 0 indicates that the function is disabled. Ignore this parameter for UDS connections.
keepalives_idle	The number of seconds of inactivity after which TCP should send a keepalive message to the server. The value 0 indicates that the default value is used. Ignore this parameter for UDS connections or if keepalive is disabled.
keepalives_interval	The number of seconds after which a TCP keepalive message that is not acknowledged by the server should be retransmitted. The value 0 indicates that the default value is used. Ignore this parameter for UDS connections or if keepalive is disabled.
keepalives_count	Controls the number of times that keepalive messages are sent through TCP. The value 0 indicates that the default value is used. Ignore this parameter for UDS connections or if keepalive is disabled.

Parameter	Description
tcp_user_timeout	Specifies the maximum duration for which transmitted data can remain unacknowledged before the TCP connection is forcibly closed on an OS that supports the TCP_USER_TIMEOUT socket option. The value 0 indicates that the default value is used. Ignore this parameter for UDS connections.
rw_timeout	Sets the read and write timeout interval of the client connection. When the timeout is triggered on the libpq and the connection is closed, the running services delivered by the libpq to the database are forcibly terminated. This capability is controlled by the GUC parameter check_disconnect_query . If this parameter is set to on , the capability is supported. If this parameter is set to off , the capability is not supported.
sslmode	Specifies whether to enable SSL encryption. • disable: SSL connection is disabled. • allow: If the database server requires SSL connection, SSL connection can be enabled. However, authenticity of the database server will not be verified. • prefer: If the database supports SSL connection, SSL connection is preferred. However, authenticity of the database server will not be verified. • require: SSL connection is required and data is encrypted. However, authenticity of the database server will not be verified. • verify-ca: SSL connection is required. Currently, Windows ODBC does not support cert authentication. • verify-full: SSL connection is required. Currently, Windows ODBC does not support cert authentication.
sslcompression	If the value is 1 (default value), the data transmitted over the SSL connection is compressed (this requires that the OpenSSL version be 0.9.8 or later). If the value is 0 , compression will be disabled (this requires OpenSSL 1.0.0 or later). If a connection without SSL is established, this parameter is ignored. If the OpenSSL version in use does not support this parameter, it will also be ignored. Compression takes up CPU time, but it increases throughput when the bottleneck is the network. If CPU performance is a limiting factor, disabling compression can improve response time and throughput.
sslcert	This parameter specifies the file name of the client SSL certificate. If no SSL connection is established, this parameter is ignored.
sslkey	This parameter specifies the location of the key used for the client certificate. It can specify a key obtained from an external "engine" (the engine is a loadable module of OpenSSL). The description of an external engine should consist of a colon-separated engine name and an engine-related key identifier. If no SSL connection is established, this parameter is ignored.

Parameter	Description
sslrootcert	This parameter specifies the name of a file that contains the SSL Certificate Authority (CA) certificate. If the file exists, the system authenticates the server certificate issued by one of these authorities.
sslcr	This parameter specifies the file name of the SSL Certificate Revocation List (CRL). If a certificate listed in this file exists, the server certificate authentication will be rejected.
requirepeer	This parameter specifies the OS user of the server, for example, requirepeer=postgres . When a UDS connection is established, if this parameter is set, the client checks whether the server process is running under the specified username at the beginning of the connection. If not, the connection will be interrupted by an error. This parameter can be used to provide server authentication similar to that of the SSL certificate on TCP/IP connections. (Note that if the UDS is in /tmp or another public writable location, any user can start a server for listening on the location. Use this parameter to ensure that you are connected to a server that is run by a trusted user.) This option is supported only on platforms that implement the peer authentication method.
krbsrvname	This parameter specifies the Kerberos service name used for GSSAPI authentication. For successful Kerberos authentication, this value must match the service name specified in the server configuration.
gsslib	This parameter specifies the GSS library used for GSSAPI authentication. It is used only in the Windows OS. If the value is gssapi , libpq is forced to use the GSSAPI library to replace the default SSPI for authentication.
service	This parameter specifies the name of the service for which the additional parameter is used. It specifies a service name in pg_service.conf that holds the additional connection parameters. This allows the application to specify only one service name so that the connection parameters can be centrally maintained.
authtype	authtype is no longer used, so it is marked as a parameter not to be displayed. It is retained in an array so as not to reject the conninfo string from old applications that might still try to set it.
remote_node_name	Specifies the name of the remote node connected to the local node.
localhost	Specifies the localhost in a connection channel.
localport	Specifies the local port in a connection channel.

Parameter	Description
fencedUdfRPCMode	Specifies whether the fenced UDF RPC protocol uses UDSs or special socket file names. The default value is 0 , indicating that the UDS mode is used and the file type is <code>.s.PGSQL.%d</code> . To use the fenced UDF mode, set this parameter to 1 . In this case, the file type is <code>.s.fencedMaster_unixdomain</code> .
replication	Specifies whether the connection should use replication protocols instead of common protocols. Protocols with this parameter configured are internal protocols used for PostgreSQL replication connections and tools such as pg_basebackup, while they can also be used by third-party applications. The following values, which are case-insensitive, are supported: • true, on, yes, and 1: Specify that the physical replication mode is connected. • database: Specifies that the logical replication mode and the database specified by dbname are connected. • false, off, no, and 0: Specify that the connection is a regular connection, which is the default behavior. In physical or logical replication mode, only simple query protocols can be used.
backend_version	Specifies the backend version to be passed to the remote end.
prototype	Sets the current protocol level. The default value is PROTO_TCP .
connection_info	The value of connection_info is a JSON character string consisting of driver_name, driver_version, driver_path, and os_user. If the value is not null, use connection_info and ignore connectionExtraInf. If the value is null, a connection information string related to libpq is generated. When connectionExtraInf is set to false, the value of connection_info consists of only driver_name and driver_version.
connectionExtraInf	Specifies whether the value of connection_info contains extension information. The default value is 0 . If the value contains other information, set this parameter to 1 .

Parameter	Description
target_session_attrs	Specifies the type of the host to be connected. The connection is successful only when the host type is the same as the configured value. The rules for setting target_session_attrs are as follows: • any (default value): All types of hosts can be connected. • read-write: The connection is set up only when the connected host is readable and writable. • read-only: Only readable hosts can be connected. • primary: Only the primary node in the primary/standby system can be connected. • standby: Only the standby node in the primary/standby system can be connected. • prefer-standby: The system first attempts to find a standby node for connection. If all hosts in the hosts list fail to be connected, try the any mode.

5.8.3 Parameters Related to Log Output

To control the output of log files and better understand the operating status of the database, modify specific configuration parameters in the **postgresql.conf** file in the instance data directory.

[Table 5-92](#) describes the adjustable configuration parameters.

Table 5-92 Configuration parameters

Parameter	Description	Value Range	Remarks
client_min_messages	Level of messages to be sent to clients.	• DEBUG5 • DEBUG4 • DEBUG3 • DEBUG2 • DEBUG1 • LOG • NOTICE • WARNING • ERROR • FATAL • PANIC Default value: NOTICE	Messages of the set level or lower will be sent to clients. The lower the level is, the fewer the messages will be sent.

Parameter	Description	Value Range	Remarks
log_min_messages	Level of messages to be recorded in server logs.	• DEBUG5 • DEBUG4 • DEBUG3 • DEBUG2 • DEBUG1 • INFO • NOTICE • WARNING • ERROR • LOG • FATAL • PANIC Default value: WARNING	Messages higher than the set level will be recorded in logs. The higher the level is, the fewer the server logs will be recorded.
log_min_error_statement	Level of SQL error statements to be recorded in server logs.	• DEBUG5 • DEBUG4 • DEBUG3 • DEBUG2 • DEBUG1 • INFO • NOTICE • WARNING • ERROR • FATAL • PANIC Default value: ERROR	SQL error statements of the set level or higher will be recorded in server logs. Only a system administrator is allowed to modify this parameter.

Parameter	Description	Value Range	Remarks
log_min_duration_statement	Minimum execution duration of a statement. If the execution duration of a statement is equal to or longer than the set milliseconds, the statement and its duration will be recorded in logs. Enabling this function can help you track the query attempts to be optimized.	INT type Default value: 30min Unit: millisecond	The value -1 indicates that the function is disabled. Only a system administrator is allowed to modify this parameter.
log_connections/ log_disconnections	Specifies whether to record a server log message when each session is connected or disconnected.	• on: The system records a log server when each session is connected or disconnected. • off: The system does not record a log server when each session is connected or disconnected. Default value: off	-
log_duration	Specifies whether to record the duration of each executed statement.	• on: The system records the duration of each executed statement. • off: The system does not record the duration of each executed statement. Default value: off	Only a system administrator is allowed to modify this parameter.

Parameter	Description	Value Range	Remarks
log_statement	SQL statements to be recorded in logs.	• none: The system does not record any SQL statements. • ddl: The system records data definition statements. • mod: The system records data definition statements and data operation statements. • all: The system records all statements. Default value: none	Only a system administrator is allowed to modify this parameter.
log_hostname	Specifies whether to record host names.	• on: The system records host names. • off: The system does not record host names. Default value: off	By default, connection logs only record the IP addresses of connected hosts. With this function, the host names will also be recorded. This parameter has an impact on viewing audit results, GS_SESSION_MEMORY_DETAIL , PG_STAT_ACTIVITY , and the parameter log_line_prefix .

[Table 5-93](#) describes the preceding parameter levels.

Table 5-93 Description of log level parameters

Level	Description
DEBUG[1-5]	Provides information that can be used by developers. Level 1 is the lowest level whereas level 5 is the highest level.
INFO	Provides information about users' hidden requests, for example, information about the VACUUM VERBOSE process.

Level	Description
NOTICE	Provides information that may be important to users, for example, truncations of long identifiers or indexes created as a part of a primary key.
WARNING	Provides warning information for users, for example, COMMIT out of transaction blocks.
ERROR	Reports an error that causes a command to terminate.
LOG	Reports information that administrators may be interested in, for example, the activity levels of check points.
FATAL	Reports the reason that causes a session to terminate.
PANIC	Reports the reason that causes all sessions to terminate.


6 SQL Optimization

The aim of SQL optimization is to maximize the utilization of resources, including CPU, memory, and disk I/O. All optimization methods are intended for resource utilization. To maximize resource utilization is to run SQL statements as efficiently as possible to achieve the highest performance at a lower cost. For example, when performing a typical point query, you can use a Seq Scan and a filter (that is, read every tuple and point query conditions for match). You can also use an Index Scan, which can be implemented at a lower cost but achieve the same effect.

You can determine a proper database deployment solution and table definition based on hardware resources and service characteristics. This is the basis of meeting performance requirements. The following performance tuning sections assume that you have finished installation based on a proper database solution in the software installation guide and performed database design based on the guide for database design and development.

6.1 Query Execution Process

The process from receiving SQL statements to the statement execution by the SQL engine is shown in [Figure 6-1](#) and described in [Table 6-1](#). The texts in red are steps where database administrators can optimize queries.

Figure 6-1 Execution process of query-related SQL statements by the SQL engine**Table 6-1** Execution process of query-related SQL statements by the SQL engine

Step	Description
1. Perform syntax and lexical parsing.	Converts the input SQL statements from the string data type to the formatted structure stmt based on the specified SQL statement rules.
2. Perform semantic parsing.	Converts the formatted structure obtained from the previous step into objects that can be recognized by the database.
3. Rewrite the query statements.	Converts the output of the previous step into the structure that optimizes the query execution.
4. Optimize the query.	Determines the execution mode of SQL statements (the execution plan) based on the result obtained from the previous step and the internal database statistics. For details about how the internal database statistics and GUC parameters affect the query optimization (execution plan), see Optimizing Queries Using Statistics and Optimizing Queries Using GUC parameters .
5. Perform the query.	Executes the SQL statements based on the execution path specified in the previous step. Selecting a proper underlying storage mode improves the query execution efficiency.

Optimizing Queries Using Statistics

The GaussDB optimizer is a typical Cost-based Optimization (CBO). By using CBO, the database calculates the number of tuples and the execution cost for each step under each execution plan based on the number of table tuples, column width, null record ratio, and characteristic values, such as distinct, MCV, and HB values, and certain cost calculation methods. The database then selects the execution plan that takes the lowest cost for the overall execution or for the return of the first tuple. These characteristic values are the statistics, which is the core for optimizing a query. Accurate statistics helps the planner select the most appropriate query plan. Generally, you can collect statistics of a table or that of some columns in a table using **ANALYZE**. You are advised to periodically execute **ANALYZE** or execute it immediately after you modified most contents in a table.

Note that DDL operations may cause statistics to change, resulting in plan changes. After DDL operations are performed on a table, check whether statistics need to be collected again.

Optimizing Queries Using GUC parameters

Optimizing queries aims to select an efficient execution mode.

Take the following SQL statement as an example:

```
select count(1)
from customer inner join store_sales on (ss_customer_sk = c_customer_sk);
```

During execution of **customer inner join store_sales**, GaussDB supports nested loop, merge join, and hash join. The optimizer estimates the result set size and the execution cost of each join mode based on the statistics on the **customer** and **store_sales** tables. It then compares the costs and selects the one costing the least.

As mentioned above, the execution cost is calculated based on certain methods and statistics. If the actual execution cost cannot be accurately estimated, you need to optimize the execution plan by setting GUC parameters. For example, the **random_page_cost** parameter indicates the optimizer's calculation of the cost of a non-sequentially-fetched disk page. The default value is **4**. When the random read speed of a machine disk (for example, an SSD) is high, you can decrease the value of this parameter. After the change, the cost of index scan is reduced, and the index scan mode is preferred when a plan is generated.

Optimizing Queries by Rewriting SQL Statements

Besides the preceding methods that improve the performance of the execution plan generated by the SQL engine, database administrators can also enhance SQL statement performance by rewriting SQL statements while retaining the original service logic based on the execution mechanism of the database and abundant practices.

This requires that database administrators know the customer services well and have professional knowledge of SQL statements. Below chapters will describe some common SQL rewriting scenarios.

6.2 Introduction to the SQL Execution Plan

6.2.1 Overview

The SQL execution plan is a node tree, which displays detailed procedure when GaussDB runs an SQL statement. A database operator indicates one step.

You can run the **EXPLAIN** command to view the execution plan generated for each query by an optimizer. The output of **EXPLAIN** has one row for each execution node, showing the basic node type and the cost estimation that the optimizer made for the execution of this node. For details, see [Figure 6-2](#).

Figure 6-2 SQL execution plan example

```
openGauss=# explain select *from t1, t2 where t1.c1=t2.c2;
               QUERY PLAN
-----
Hash Join (cost=58.35..355.67 rows=23091 width=16)
-> Seq Scan on t1 (cost=0.00..31.49 rows=2149 width=8)
-> Hash (cost=31.49..31.49 rows=2149 width=8)
    -> Seq Scan on t2 (cost=0.00..31.49 rows=2149 width=8)
(5 rows)
```

- Nodes at the bottom level are scan nodes. They scan tables and return raw rows. The types of scan nodes (sequential scans and index scans) vary depending on the table access methods. Objects scanned by the bottom layer nodes may not be row-store data (not directly read from a table), such as **VALUES** clauses and functions that return rows, which have their own types of scan nodes.
- If the query requires join, aggregation, sorting, or other operations on the raw rows, there will be other nodes above the scan nodes to perform these operations. In addition, there is more than one way to perform these operations, so different types of execution nodes may be displayed here.
- The first row (the upper-layer node) estimates the total execution cost of the execution plan. Such an estimate indicates the value that the optimizer tries to minimize.

Execution Plan Display Format

GaussDB provides four display formats: normal, pretty, summary, and run.

- normal: indicates that the default printing format is used.
- pretty: indicates that the new plan display format improved by GaussDB is used. The new format contains a plan node ID, directly and effectively analyzing performance.
- summary: indicates that the printing information analysis is added based on the pretty format.
- run: indicates that the information based on the summary format is exported as a CSV file for further analysis.

An example of an execution plan using the pretty format is as follows:

```
explain select * from t1,t2 where t1.c1=t2.c2;
id | operation | E-rows | E-width | E-costs
-----+-----+-----+-----+-----
1 | -> Hash Join (2,3) | 23091 | 16 | 58.353..355.674
2 | -> Seq Scan on t1 | 2149 | 8 | 0.000..31.490
3 | -> Hash | 2149 | 8 | 31.490..31.490
4 | -> Seq Scan on t2 | 2149 | 8 | 0.000..31.490
(4 rows)

Predicate Information (identified by plan id)
-----
1 --Hash Join (2,3)
Hash Cond: (t1.c1 = t2.c2)
(2 rows)
```

You can change the display format of execution plans by setting the GUC parameter **explain_perf_mode**. Later examples use the pretty format by default.

Execution Plan Information

In addition to setting different display formats for an execution plan, you can use different **EXPLAIN** syntax to display execution plan information in detail. The following lists the common EXPLAIN syntax. For more EXPLAIN syntax, see [EXPLAIN](#).

- **EXPLAIN *statement***: only generates an execution plan and does not execute. The *statement* indicates SQL statements.
- **EXPLAIN ANALYZE *statement***: generates and executes an execution plan, and displays the execution summary. Then actual execution time statistics are added to the display, including the total elapsed time expended within each plan node (in milliseconds) and the total number of rows it actually returned.
- **EXPLAIN PERFORMANCE *statement***: generates and executes the execution plan, and displays all execution information.

To measure the run time cost of each node in the execution plan, the current execution of EXPLAIN ANALYZE or EXPLAIN PERFORMANCE adds profiling overhead to query execution. Running **EXPLAIN ANALYZE** or **EXPLAIN PERFORMANCE** on a query sometimes takes more time than a normal query. The amount of overhead depends on the nature of the query, as well as the platform being used.

Therefore, if an SQL statement is not finished after being running for a long time, run the **EXPLAIN** statement to view the execution plan and then locate the fault. If the SQL statement has been properly executed, execute the EXPLAIN ANALYZE or EXPLAIN PERFORMANCE statement to check the execution plan and information to locate the fault.

6.2.2 Description

As described in [Overview](#), EXPLAIN displays the execution plan, but will not actually run SQL statements. **EXPLAIN ANALYZE** and **EXPLAIN PERFORMANCE** both will actually run SQL statements and return the execution information. This section describes the execution plan and execution information in detail.

Execution Plans

The following SQL statement is used as an example:

```
SELECT * FROM t1, t2 WHERE t1.c1 = t2.c2;
```

Run the **EXPLAIN** command and the output is as follows:

```
openGauss=# explain select *from t1, t2 where t1.c1=t2.c2;
               QUERY PLAN
-----
Hash Join (cost=58.35..355.67 rows=23091 width=16)
-> Seq Scan on t1 (cost=0.00..31.49 rows=2149 width=8)
-> Hash (cost=31.49..31.49 rows=2149 width=8)
    -> Seq Scan on t2 (cost=0.00..31.49 rows=2149 width=8)
(5 rows)
```

Interpretation of the execution plan level (vertical):

1. Layer 1: **Seq Scan on t2**

The table scan operator scans the table **t2** using **Seq Scan**. At this layer, data in the table **t2** is read from a buffer or disk, and then transferred to the upper-layer node for calculation.

2. Layer 2: **Hash**

Hash operator. It is used to calculate the hash value of the operator transferred from the lower layer for subsequent hash join operations.

3. Layer 3: **Seq Scan on t1**

The table scan operator scans the table **t1** using **Seq Scan**. At this layer, data in the table **t1** is read from a buffer or disk, and then transferred to the upper-layer node for hash join calculation.

4. Layer 4: **Hash Join**

Join operator. It is used to join data in the **t1** and **t2** tables using the hash join method and output the result data.

Keywords in the execution plan:

1. Table access modes

- Seq Scan

Scans all rows of the table in sequence.

- Index Scan

The optimizer uses a two-step plan: the child plan node visits an index to find the locations of rows matching the index condition, and then the upper plan node actually fetches those rows from the table itself. Fetching rows separately is much more expensive than reading them sequentially, but because not all pages of the table have to be visited, this is still cheaper than a sequential scan. The upper-layer planning node sorts index-identified rows based on their physical locations before reading them. This minimizes the independent capturing overhead.

If there are separate indexes on multiple columns referenced in **WHERE**, the optimizer might choose to use an **AND** or **OR** combination of the indexes. However, this requires the visiting of both indexes, so it is not

necessarily a win compared to using just one index and treating the other condition as a filter.

Index scans can be classified into the following types based on the index sorting mechanism:

- Bitmap index scan
Fetches data pages using a bitmap.
- Index scan using index_name
Uses simple index search, which fetches data from an index table in the sequence of index keys. This mode is commonly used when only a small amount of data needs to be fetched from a large data table or when the ORDER BY condition is used to match the index sequence to reduce the sorting time.
- Index-Only Scan
Scans the index that contains all required data, instead of referencing a table.
- Bitmap Heap Scan
Reads pages from bitmaps created by other operations and filters out the rows that do not meet the conditions. Bitmap heap scan can avoid random I/Os and accelerate read speed.
- TID Scan
Scans a table by a tuple ID.
- Index Ctid Scan
Scans a table based on the CTID index.
- CTE Scan
Specifies that CTE evaluates subquery operations and stores query results as a temporary table. The temporary table is scanned by the CTE Scan operator.
- Foreign Scan
Reads data from a remote data source.
- Function Scan
Obtains result sets returned by functions and returns them as the rows read from tables.
- Sample Scan
Queries and returns sampled data.
- Subquery Scan
Reads subquery results.
- Values Scan
Reads constants as part of the **VALUES** command.
- WorkTable Scan
Scans a work table. Data is read in the middle of an operation which is usually a recursive operation declared using WITH RECURSIVE.
- CStore Index Ctid Scan
Scans a table based on an index and returns a qualified TID set.

- CStore Index Heap Scan
Performs the intersection, difference, and union operations of TID sets, and obtains corresponding tuple based on set results.
- 2. Table connection modes
 - Nested Loop
A nested loop is used for queries that have a smaller data set connected. In a nested-loop join, the outer table drives the inner table and each row returned from the outer table should have a matching row in the inner table. The returned result set of all queries should be less than 10,000. The table that returns a smaller subset will work as an outer table, and indexes are recommended for connection columns of the inner table.
 - (Sonic) Hash Join
A hash join is used for large tables. The optimizer uses a hash join, in which rows of one table are entered into an in-memory hash table, after which the other table is scanned and the hash table is probed for matches to each row. Sonic and non-Sonic hash joins differ in their hash table structures, which do not affect the execution result set.
 - Merge Join
In most cases, the execution performance of a merge join is lower than that of a hash join. However, if the source data has been pre-sorted and no more sorting is needed during the merge join, its performance excels.
- 3. Operators
 - sort
Sorts the result set.
 - filter
The **EXPLAIN** output shows the **WHERE** clause being applied as a **Filter** condition attached to the **Seq Scan** plan node. This means that the plan node checks the condition for each row it scans, and returns only the ones that meet the condition. The estimated number of output rows has been reduced because of the **WHERE** clause. However, the scan will still have to visit all 10,000 rows, as a result, the cost is not decreased. It increases a bit (by 10,000 x **cpu_operator_cost**) to reflect the extra CPU time spent on checking the **WHERE** condition.
 - LIMIT
Limits the number of output execution results. If a **LIMIT** condition is added, not all rows are retrieved.
 - Append
Appends sub-operation results.
 - Aggregate
Aggregates the results generated from querying row. It can be an aggregation of statements such as GROUP BY, UNION, and SELECT DISTINCT.
 - BitmapAnd
Specifies the AND operation of a bitmap, which is used to form a bitmap that matches more complex conditions.
 - BitmapOr

- Specifies the OR operation of a bitmap, which is used to form a bitmap that matches more complex conditions.
- Gather
Gathers data of parallel threads.
- Group
Groups rows to perform the GROUP BY operation.
- GroupAggregate
Aggregates the pre-sorted rows of the GROUP BY operation.
- Hash
Hashes rows for the parent query. It is usually used to perform the JOIN operation.
- HashAggregate
Aggregates the result rows of GROUP BY by using a hash table.
- Merge Append
Merges subquery results in a way that preserves the sort order. It can be used to merge sorted rows in a table partition.
- Recursive Union
Performs a union operation on all steps of a recursive function.
- SetOp
Specifies a set operation, such as INTERSECT or EXCEPT.
- Unique
Removes duplicates from an ordered result set.
- HashSetOp
Specifies a strategy for set operations such as INTERSECT or EXCEPT. It uses Append to avoid pre-sorted input.
- LockRows
Locks problematic rows to prevent other queries from writing, but allows reading.
- Materialize
Stores subquery results in the memory so that the parent query can quickly access and obtain the subquery results.
- Result
Returns a value (such as a hard-coded value) without scanning.
- WindowAgg
Specifies a window aggregate function, which is triggered by the OVER statement.
- Merge
Performs a merge operation.
- StartWith Operator
Specifies the hierarchical query operator, which is used to perform recursive query operations.
- Index Cond
Specifies the index scan conditions.

- Cstore Index And
Performs the intersection operation of TID sets with CStore Index Heap Scan.
 - Cstore Index Or
Performs the union operation of TID sets with CStore Index Heap Scan.
4. Other keywords
- Partitioned
Indicates operations on a specific partition.
 - Partition Iterator
Partition iterator, which usually indicates that a subquery is an operation on a partition.
 - InitPlan
Indicates a non-related subplan.

Execution Information

The execution result of the following SQL statement in pretty mode is used as an example:

```
select sum(t2.c1) from t1,t2 where t1.c1=t2.c2 group by t1.c2;
```

The output of running EXPLAIN PERFORMANCE is as follows:

```
explain performance select sum(t2.c1) from t1,t2 where t1.c1=t2.c2 group by t1.c2;
id | operation | A-time | A-rows | E-rows | E-distinct | Peak Memory | A-width | E-width | E-costs
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
1 | -> HashAggregate | 0.574 | 0 | 200 | | 29KB | | 8 |
396.113..398.113
2 | -> Hash Join (3,4) | 0.358 | 0 | 18915 | 200, 200 | 12KB | | 8 |
53.763..301.538
3 | -> Seq Scan on public.t1 | 0.037 | 1 | 1945 | | 22KB | | 8 | 0.000..29.450
4 | -> Hash | 0.038 | 0 | 1945 | | 264KB | | 8 | 29.450..29.450
5 | -> Seq Scan on public.t2 | 0.029 | 30 | 1945 | | 22KB | | 8 | 0.000..29.450
(5 rows)

Predicate Information (identified by plan id)
-----
2 --Hash Join (3,4)
Hash Cond: (t1.c1 = t2.c2)
(2 rows)

Memory Information (identified by plan id)
-----
1 --HashAggregate
Peak Memory: 29KB, Estimate Memory: 64MB
2 --Hash Join (3,4)
Peak Memory: 12KB, Estimate Memory: 64MB
3 --Seq Scan on public.t1
Peak Memory: 22KB, Estimate Memory: 64MB
4 --Hash
Peak Memory: 264KB
Buckets: 32768 Batches: 1 Memory Usage: 0kB
5 --Seq Scan on public.t2
Peak Memory: 22KB, Estimate Memory: 64MB
(11 rows)

Targetlist Information (identified by plan id)
-----
```

```
1 --HashAggregate
  Output: sum(t2.c1), t1.c2
  Group By Key: t1.c2
2 --Hash Join (3,4)
  Output: t1.c2, t2.c1
3 --Seq Scan on public.t1
  Output: t1.c1, t1.c2, t1.c3
4 --Hash
  Output: t2.c1, t2.c2
5 --Seq Scan on public.t2
  Output: t2.c1, t2.c2
(11 rows)

-----
Datanode Information (identified by plan id)
-----
1 --HashAggregate
  (actual time=0.574..0.574 rows=0 loops=1)
  (Buffers: shared hit=2)
  (CPU: ex c/r=0, ex row=0, ex cyc=527797, inc cyc=8385141377087373)
2 --Hash Join (3,4)
  (actual time=0.358..0.358 rows=0 loops=1)
  (Buffers: shared hit=2)
  (CPU: ex c/r=-8385141375712241, ex row=1, ex cyc=-8385141375712241, inc cyc=8385141376559576)
3 --Seq Scan on public.t1
  (actual time=0.037..0.037 rows=1 loops=1)
  (Buffers: shared hit=1)
  (CPU: ex c/r=8385141375728512, ex row=1, ex cyc=8385141375728512, inc cyc=8385141375728512)
4 --Hash
  (actual time=0.038..0.038 rows=0 loops=1)
  (Buffers: shared hit=1)
  (CPU: ex c/r=0, ex row=0, ex cyc=-251554241295571040, inc cyc=8385141376543305)
5 --Seq Scan on public.t2
  (actual time=0.019..0.029 rows=30 loops=1)
  (Buffers: shared hit=1)
  (CPU: ex c/r=8664646089070478, ex row=30, ex cyc=259939382672114336, inc
cyc=259939382672114336)
(20 rows)

===== Query Summary =====
-----
Datanode executor start time: 0.180 ms
Datanode executor run time: 0.590 ms
Datanode executor end time: 0.051 ms
Planner runtime: 0.366 ms
Query Id: 844424930141239
Total runtime: 0.866 ms
(6 rows)
```

In the preceding example, the execution information consists of the following parts:

1. The plan is displayed as a table, which contains 11 columns: **id**, **operation**, **A-time**, **A-rows**, **E-rows**, **E-distinct**, **Peak Memory**, **E-memory**, **A-width**, **E-width**, and **E-costs**. The meanings of the plan-type columns (**id**, **operation**, or columns started with **E**) are the same as those when EXPLAIN is executed. For details, see [Execution Plans](#). The definition of **A-time**, **A-rows**, **E-distinct**, **Peak Memory**, and **A-width** are described as follows:
 - **A-time**: execution completion time of the current operator.
 - **A-rows**: number of actual output tuples of the operator
 - **E-distinct**: estimated distinct value of the hash join operator
 - **Peak Memory**: peak memory used by the operator during execution.
 - **A-width**: actual tuple width in each row of the current operator. This parameter is valid only for heavy memory operators, including (Vec)HashJoin, (Vec)HashAgg, (Vec)HashSetOp, (Vec)Sort, and

(Vec)Materialize. The (Vec)HashJoin calculation width is the width of its right subtree operator and will be displayed on the right subtree.

2. **Predicate Information (identified by plan id):**

This part displays the static information that does not change in the plan execution process, such as some join conditions and filter information.

3. **Memory Information (identified by plan id):**

This part displays the memory usage information printed by certain operators (mainly Hash and Sort), including **peak memory**, **control memory**, **operator memory**, **width**, **auto spread num**, and **early spilled**; and spill details, including **spill Time(s)**, **inner/outer partition spill num**, **temp file num**, spilled data volume, and **written disk IO** [*min*, *max*].

4. **Targetlist Information (identified by plan id):**

This part displays the target columns provided by each operator.

5. **DataNode Information (identified by plan id):**

The execution time, CPU, and buffer usage of each operator are printed in this part.

6. **=====
Query Summary
=====**

This part displays the total execution time and network traffic, including the maximum and minimum execution time in the initialization and end phases, available system memory when the current statement is executed, and estimated statement memory.

6.3 Tuning Process

You can analyze slow SQL statements to optimize them.

Procedure

- Step 1** Collect all table statistics associated with the SQL statements. In a database, statistics indicate the source data of a plan generated by a planner. If no statistics are available or out of date, the execution plan may seriously deteriorate, leading to low performance. According to past experience, about 10% performance problems occurred because no statistics are collected. For details, see [Updating Statistics](#).
- Step 2** View the execution plan to find out the cause. If the SQL statements have been running for a long period of time and not ended, run the **EXPLAIN** statement to view the execution plan and then locate the fault. If the SQL statement has been properly executed, run the **EXPLAIN ANALYZE** or **EXPLAIN PERFORMANCE** statement to check the execution plan and information to locate the fault. For details about the execution plan, see [Introduction to the SQL Execution Plan](#).
- Step 3** Review and modify a table definition. For details, see [Reviewing and Modifying a Table Definition](#).
- Step 4** For details about **EXPLAIN** or **EXPLAIN PERFORMANCE**, the reason why SQL statements are slowly located, and how to solve this problem, see [Typical SQL Optimization Methods](#).

Step 5 Generally, some SQL statements can be converted to its equivalent statements in all or certain scenarios by rewriting queries. SQL statements are simpler after they are rewritten. Some execution steps can be simplified to improve the performance. Query rewriting methods are universal in all databases. [Experience in Rewriting SQL Statements](#) describes several tuning methods by rewriting SQL statements.

----End

6.4 Updating Statistics

In a database, statistics indicate the source data of a plan generated by a planner. If no statistics are available or out of date, the execution plan may seriously deteriorate, leading to low performance.

Context

The **ANALYZE** statement collects statistic about table contents in databases, which will be stored in the **PG_STATISTIC** system catalog. Then, the query optimizer uses the statistics to work out the most efficient execution plan.

After executing batch insertions and deletions, you are advised to run the **ANALYZE** statement on the table or the entire library to update statistics. By default, 30,000 rows of statistics are sampled. That is, the default value of the GUC parameter **default_statistics_target** is **100**. If the total number of rows in the table exceeds 1,600,000, you are advised to set **default_statistics_target** to **-2**, indicating that 2% of the statistics are collected.

For an intermediate table generated during the execution of a batch script or stored procedure, you also need to run the **ANALYZE** statement.

Procedure

Run the following commands to update the statistics about a table or the entire database:

```
ANALYZE tablename;           -- Update statistics about a table.  
ANALYZE;                     -- Update statistics about the entire database.
```

NOTE

Use **EXPLAIN** to show the execution plan of each SQL statement. If **rows=10** (the default value, probably indicating that the table has not been analyzed) is displayed in the **SEQ SCAN** output of a table, run the **ANALYZE** statement for this table.

6.5 Reviewing and Modifying a Table Definition

6.5.1 Overview

To properly define a table, you must:

1. **Reduce the data volume scanned** by using the partition pruning mechanism.
2. **Minimize random I/Os** by using clustering or partial clustering.

The table definition is created during the database design and is reviewed and modified during the SQL statement optimization.

6.5.2 Selecting a Storage Model

During database design, some key factors about table design will greatly affect the subsequent query performance of the database. Table design affects data storage as well. Scientific table design reduces I/O operations and minimizes memory usage, improving the query performance.

Selecting a model for table storage is the first step of table definition. Select a proper storage model for your service based on the following table.

Storage Model	Application Scenario
Row store	Point queries (simple index-based queries that only return a few records) Scenarios requiring frequent addition, deletion, and modification operations

6.5.3 Using Partitioned Tables

Partitioning refers to splitting what is logically one large table into smaller physical pieces based on specific schemes. The table based on the logic is called a partitioned table, and a physical piece is called a partition. Data is stored in physical partitions not the logical table. A partitioned table has the following advantages over an ordinary table:

1. High query performance: You can specify partitions when querying partitioned tables, improving query efficiency.
2. High availability: If a certain partition in a partitioned table is faulty, data in the other partitions is still available.
3. Easy maintenance: To fix a partitioned table having a faulty partition, you only need to fix the partition.

Partitioned tables supported by the GaussDB database are level-1 and level-2 partitioned tables. Level-1 partitioned tables include range partitioned tables, interval partitioned tables, list partitioned tables, and hash partitioned tables. Level-2 partitioned tables include nine combinations of any two of range partitioned tables, list partitioned tables, and hash partitioned tables.

- Range partitioned table: Data in different ranges is mapped to different partitions. The range is determined by the partition key specified during the partitioned table creation. The partition key is usually a date. For example, sales data is partitioned by month.
- Interval partitioned table: a special type of range partitioned tables. Compared with range partitioned tables, interval value definition is added. When no matching partition can be found for an inserted record, a partition can be automatically created based on the interval value.
- List partitioned table: Key values contained in the data are stored in different partitions, and the data is mapped to each partition in sequence. The key

values contained in the partitions are specified when the partitioned table is created.

- Hash partitioned table: Data is mapped to each partition based on the internal hash algorithm. The number of partitions is specified when the partitioned table is created.
- Level-2 partitioned table: a partitioned table obtained by randomly combining range partitioning, list partitioning, and hash partitioning. Both level-1 and level-2 partitions can be defined in the preceding three ways.

6.5.4 Selecting a Data Type

Use the following principles to select efficient data types:

1. **Select data types that facilitate data calculation.**

Generally, the calculation of integers (including common comparison calculations, such as =, >, <, ≥, ≤, and ≠ and group by) is more efficient than that of strings and floating point numbers.

2. **Select data types with a short length.**

Data types with short length reduce both the data file size and the memory used for computing, improving the I/O and computing performance. For example, use **SMALLINT** instead of **INT**, and **INT** instead of **BIGINT**.

3. **Use the same data type for a join.**

You are advised to use the same data type for a join. To join columns with different data types, the database needs to convert them to the same type, which leads to additional performance overheads.

6.6 Typical SQL Optimization Methods

SQL optimization involves continuous analysis and trying. Queries are run before they are used for services to determine whether the performance meets requirements. If it does not, queries will be optimized by [checking the execution plan](#) and identifying the causes. Then, the queries will be run and optimized again until they meet the requirements.

6.6.1 Optimizing SQL Self-Diagnosis

Performance issues may occur when you query data or run the **INSERT**, **DELETE**, **UPDATE**, or **CREATE TABLE AS** statement. In this case, you can query the **warning** column in the [PG_CONTROL_GROUP_CONFIG](#) and [GS_SESSION_MEMORY_DETAIL](#) views to obtain reference for performance optimization.

Alarms that can trigger SQL self diagnosis depend on the settings of **resource_track_level**. If **resource_track_level** is set to **query**, alarms about the failures in collecting column statistics and pushing down SQL statements will trigger the diagnosis. If **resource_track_level** is set to **operator**, all alarms will trigger the diagnosis.

Whether an SQL plan will be diagnosed depends on the settings of **resource_track_cost**. An SQL plan will be diagnosed only if its execution cost is greater than **resource_track_cost**. You can use the **EXPLAIN** keyword to check the plan execution cost.

The SQL self-diagnosis function is affected by the **enable_analyze_check** parameter. Ensure that the function is enabled before using it.

If a large number of statements are executed, certain data may fail to be collected due to memory control. In this case, you can increase the value of **instr_unique_sql_count**.

Alarms

Currently, performance alarms will be reported when statistics about one or multiple columns are not collected.

An alarm will be reported if statistics of a single column are not collected. For details about the optimization, see [Updating Statistics](#) and [Optimizing Statistics](#).

Example alarms:

No statistics about a table are not collected.

```
Statistic Not Collect:  
  schema_test.t1
```

The statistics about a single column are not collected.

```
Statistic Not Collect:  
  schema_test.t2(c1,c2)
```

Restrictions

1. An alarm contains a maximum of 2048 characters. If the length of an alarm exceeds this value (for example, a large number of long table names and column names are displayed in the alarm when their statistics are not collected), a warning instead of an alarm will be reported.
WARNING, "Planner issue report is truncated, the rest of planner issues will be skipped"
2. If a query statement contains the **Limit** operator, alarms of operators lower than **Limit** will not be reported.

6.6.2 Optimizing Subqueries

Context

When an application runs an SQL statement to operate the database, a large number of subqueries are used because they are more clear than table join. Especially in complicated query statements, subqueries have more complete and independent semantics, which makes SQL statements clearer and easier to understand. Therefore, subqueries are widely used.

In GaussDB, subqueries can also be called sublinks based on the location of subqueries in SQL statements.

- Subquery: corresponds to a range table (RangeTblEntry) in the query parse tree. That is, a subquery is a **SELECT** statement following immediately after the **FROM** keyword.
- Sublink: corresponds to an expression in the query parsing tree. That is, a sublink is a statement in the **WHERE** or **ON** clause or in the target list.

In conclusion, a subquery is a range table and a sublink is an expression in the query parsing tree. A sublink can be found in constraint conditions and expressions. In GaussDB, sublinks can be classified into the following types:

- exist_sublink: corresponds to the **EXIST** and **NOT EXIST** statements.
- any_sublink: corresponds to the *op ANY(SELECT...)* statement. *op* can be the <, >, or = operator. **IN/NOT IN (SELECT...)** also belongs to this type.
- all_sublink: corresponds to the *op ALL(SELECT...)* statement. *op* can be the <, >, or = operator.
- rowcompare_sublink: corresponds to the **RECORD op (SELECT...)** statement.
- expr_sublink: corresponds to the **(SELECT with a single target list item...)** statement.
- array_sublink: corresponds to the **ARRAY(SELECT...)** statement.
- cte_sublink: corresponds to the **WITH(...)** query statement.

The exist_sublink and any_sublink are pulled up by the optimization engine of GaussDB. In addition, expr_sublink can also be pulled up. However, because of the flexible use of subqueries in SQL statements, complex subqueries may affect query performance. To disable the optimization of **expr_sublink**, set the GUC parameter **rewrite_rule**. Subqueries are classified into non-correlated subqueries and correlated subqueries.

- Non-correlated subqueries

The execution of a subquery is independent from attributes of the outer query. In this way, a subquery can be executed before outer queries.

Example:

```
openGauss=# explain select t1.c1,t1.c2
from t1
where t1.c1 in (
  select c2
  from t2
  where t2.c2 IN (2,3,4)
);
               QUERY PLAN
-----
Hash Join
Hash Cond: (t1.c1 = t2.c2)
-> Seq Scan on t1
    Filter: (c1 = ANY ('{2,3,4}'::integer[]))
-> Hash
    -> HashAggregate
        Group By Key: t2.c2
        -> Seq Scan on t2
            Filter: (c2 = ANY ('{2,3,4}'::integer[]))
(9 rows)
```

- Correlated subqueries

The execution of a subquery depends on some attributes (used as AND conditions of the subquery) of outer queries. In the following example, **t1.c1** in the **t2.c1 = t1.c1** condition is a correlated attribute. Such a subquery depends on outer queries and needs to be executed once for each outer query.

Example:

```
openGauss=# explain select t1.c1,t1.c2
from t1
where t1.c1 in (
```

```
select c2
from t2
where t2.c1 = t1.c1 AND t2.c2 in (2,3,4)
);
          QUERY PLAN
-----
Seq Scan on t1
  Filter: (SubPlan 1)
  SubPlan 1
    -> Seq Scan on t2
        Filter: ((c1 = t1.c1) AND (c2 = ANY ('{2,3,4}'::integer[])))
(5 rows)
```

Sublink Optimization on GaussDB

To optimize a sublink, a subquery is pulled up to join with tables in outer queries. You can run the **EXPLAIN** statement to check whether a sublink is converted into such a subplan.


Example:

```
openGauss=# explain SELECT t1.c1, t1.c2 FROM t1 WHERE t1.c1 IN(SELECT c2 FROM t2 WHERE t2.c1 =
t1.c1);
QUERY PLAN
-----
Seq Scan on t1
  Filter: (SubPlan 1)
  SubPlan 1
    -> Seq Scan on t2
        Filter: (c1 = t1.c1)
(5 rows)
```

- **Sublink-release scenarios supported by GaussDB**

- Non-correlated **IN** sublink
 - The subquery cannot contain columns in the outer query (columns in more outer queries are allowed).
 - The subquery cannot contain volatile functions.

Example:



Replace the execution plan on the right of the arrow with the following execution plan:


```
QUERY PLAN
-----
Hash Join
Hash Cond: (t1.c1 = t2.c2)
-> Seq Scan on t1
-> Hash
-> HashAggregate
Group By Key: t2.c2
-> Seq Scan on t2
Filter: (c1 = 1)
(8 rows)
```

– Correlated **EXISTS** sublink

The **WHERE** clause must contain a column in the outer query. Other parts of the subquery cannot contain the column. Other restrictions are as follows:

- The subquery must contain the **FROM** clause.
- The subquery cannot contain the **WITH** clause.
- The subquery cannot contain aggregate functions.
- The subquery cannot contain a **SET**, **SORT**, **LIMIT**, **WindowAgg**, or **HAVING** operation.
- The subquery cannot contain volatile functions.

Example:



Replace the execution plan on the right of the arrow with the following execution plan:

QUERY PLAN

```

Hash Join
Hash Cond: (t1.c1 = t2.c1)
-> Seq Scan on t1
-> Hash
    -> HashAggregate
        Group By Key: t2.c1
    -> Seq Scan on t2
(7 rows)

```

– Pulling up an equivalent correlated subquery containing aggregate functions

The **WHERE** condition of the subquery must contain a column from the outer query. Equality comparison must be performed between this column and related columns in tables of the subquery. These conditions must be connected using **AND**. Other parts of the subquery cannot contain the column. Other restrictions are as follows:

- The columns in the expression in the **WHERE** condition of the subquery must exist in tables.
- After the **SELECT** keyword of the subquery, there must be only one output column. The output column must be an aggregate function (for example, **MAX**), and the parameter (for example, **t2.c2**) of the aggregate function cannot be columns of a table (for example, **t1**) in outer queries. The aggregate function cannot be **COUNT**.

For example, the following subquery can be pulled up:

```

select * from t1 where c1 >(
  select max(t2.c1) from t2 where t2.c1=t1.c1
);

```

The following subquery cannot be pulled up because the subquery has no aggregate function:

```
select * from t1 where c1 >(
    select t2.c1 from t2 where t2.c1=t1.c1
);
```

The following subquery cannot be pulled up because the subquery has two output columns:

```
select * from t1 where (c1,c2) >(
    select max(t2.c1),min(t2.c2) from t2 where t2.c1=t1.c1
);
```

- The subquery must be a **FROM** clause.
- The subquery cannot contain a **GROUP BY**, **HAVING**, or **SET** operation.
- The subquery can only be an inner join.

For example, the following example cannot be pulled up:

```
select * from t1 where c1 >(
    select max(t2.c1) from t2 full join t3 on (t2.c2=t3.c2) where t2.c1=t1.c1
);
```

- The target list of the subquery cannot contain the function that returns a set.
- The **WHERE** condition of the subquery must contain a column from the outer query. Equality comparison must be performed between this column and related columns in tables of the subquery. These conditions must be connected using **AND**. Other parts of the subquery cannot contain the column. For example, the following innermost sublink can be pulled up:

```
select * from t3 where t3.c1=(
    select t1.c1
    from t1 where c1 >(
        select max(t2.c1) from t2 where t2.c1=t1.c1
    ));
```

If another condition is added to the subquery in the previous example, the subquery cannot be pulled up because the innermost subquery references to the column in the outer and outer query. Example:

```
select * from t3 where t3.c1=(
    select t1.c1
    from t1 where c1 >(
        select max(t2.c1) from t2 where t2.c1=t1.c1 and t3.c1>t2.c2
    ));
```

- Pulling up a sublink in the **OR** clause

If the **WHERE** condition contains an **EXIST** correlated sublink connected by **OR**:

Example:

```
select a, c from t1
where t1.a = (select avg(a) from t3 where t1.b = t3.b) or
exists (select * from t4 where t1.c = t4.c);
```

The process of pulling up the OR clause of the EXIST related subquery connected by the OR condition is as follows:

- i. Extract **opExpr** from the OR clause in the WHERE condition. The value is **t1.a = (select avg(a) from t3 where t1.b = t3.b)**.
- ii. The **op** contains a subquery. If the subquery can be pulled up, the subquery is rewritten as **select avg(a), t3.b from t3 group by t3.b**, generating the **NOT NULL** condition **t3.b is not null**. The **opExpr** is replaced with this **NOT NULL** condition. In this case, the SQL statement changes to:

```
select a, c
from t1 left join (select avg(a) avg, t3.b from t3 group by t3.b) as t3 on (t1.a = avg
and t1.b = t3.b)
where t3.b is not null or exists (select * from t4 where t1.c = t4.c);
```
- iii. Extract the **EXISTS** sublink **exists (select * from t4 where t1.c = t4.c)** from the **OR** clause to check whether the sublink can be pulled up. If it can be pulled up, the subquery is converted into **select t4.c from t4 group by t4.c**, generating the **NOT NULL** condition **t4.c is not null**. In this case, the SQL statement changes to:

```
select t1.a, t1.c from t1 left join (select avg(a) avg, t3.b from t3 group by t3.b) as t3 on
(t1.a = avg and t1.b = t3.b) left join (select t5.c from t5 group by t5.c) as t5 on (t1.c =
t5.c) where t3.b is not null or t5.c is not null;
```

- **Sublink-release scenarios not supported by GaussDB**

Except the sublinks described above, all the other sublinks cannot be pulled up. In this case, a correlated subquery is planned as the combination of subplans and broadcast. As a result, if inner tables have a large amount of data, query performance may be poor.

If a correlated subquery joins with two tables in outer queries, the subquery cannot be pulled up. You need to change the parent SQL statement into a **WITH** clause and then perform the join.

Example:

```
select distinct t1.a, t2.a
from t1 left join t2 on t1.a=t2.a and not exists (select a,b from test1 where test1.a=t1.a and
test1.b=t2.a);
```

The outer query is changed into:

```
with temp as
(
    select * from (select t1.a as a, t2.a as b from t1 left join t2 on t1.a=t2.a)
)
select distinct a,b
from temp
where not exists (select a,b from test1 where temp.a=test1.a and temp.b=test1.b);
```

- The correlated subquery (without **COUNT**) in the target list cannot be pulled up.

Example:

```
openGauss=# explain (costs off)
select (select c2 from t2 where t1.c1 = t2.c1) ssq, t1.c2
from t1
where t1.c2 > 10;
```

The execution plan is as follows:

```
openGauss=# explain (costs off)
select (select c2 from t2 where t1.c1 = t2.c1) ssq, t1.c2
from t1
where t1.c2 > 10;
QUERY PLAN
-----
Seq Scan on t1
  Filter: (c2 > 10)
```

```
SubPlan 1
-> Seq Scan on t2
    Filter: (t1.c1 = c1)
(5 rows)
```

The correlated subquery is displayed in the target list (query return list). Values need to be returned even if the condition **t1.c1=t2.c1** is not met. Therefore, use a right outer join to join **t2** and **t1** so that the SSQ can return padding values when the condition **t1.c1=t2.c1** is not met.

NOTE

ScalarSubQuery (SSQ) and Correlated-ScalarSubQuery (CSSQ) are described as follows:

- SSQ: a sublink that returns a scalar value of a single row with a single column
- CSSQ: an SSQ containing correlation conditions

The preceding SQL statement can be changed into:

```
with ssq as
(
    select * from t1 where t1.c2 >10
)
select t2.c2,ssq.c2 from t2 right join ssq on ssq.c1 = t2.c1;
```

The execution plan after the change is as follows:

```
QUERY PLAN
-----
Hash Right Join
Hash Cond: (t2.c1 = t1.c1)
-> Seq Scan on t2
-> Hash
    -> Seq Scan on t1
        Filter: (c2 > 10)
(6 rows)
```

In the preceding example, the SSQ in the target list is pulled up to right join, preventing poor performance caused by the plan involving subplans when the table (**t2**) in the subquery is too large.

- The subquery (with **COUNT**) in the target list cannot be pulled up.

Example:

```
select (select count(*) from t2 where t2.c1=t1.c1) cnt, t1.c1, t3.c1
from t1,t3
where t1.c1=t3.c1 order by cnt, t1.c1;
```

The execution plan is as follows:

```
QUERY PLAN
-----
Sort
Sort Key: ((SubPlan 1)), t1.c1
-> Hash Join
Hash Cond: (t1.c1 = t3.c1)
-> Seq Scan on t1
-> Hash
    -> Seq Scan on t3
SubPlan 1
-> Aggregate
    -> Seq Scan on t2
        Filter: (c1 = t1.c1)
(11 rows)
```

The correlated subquery is displayed in the target list (query return list). Values need to be returned even if the condition **t1.c1=t2.c1** is not met. Therefore, use left outer join to join **T1** and **T2** so that SSQ can return padding values when the condition **t1.c1=t2.c1** is not met. However,

COUNT is used, which requires that **0** is returned when the condition is not met. Therefore, **case-when NULL then 0 else count(*)** can be used.

The preceding SQL statement can be changed into:

```
with ssq as
(
  select count(*) cnt, c1 from t2 group by c1
)
select case when
  ssq.cnt is null then 0
  else ssq.cnt
end cnt, t1.c1, t3.c1
from t1 left join ssq on ssq.c1 = t1.c1,t3
where t1.c1 = t3.c1
order by ssq.cnt, t1.c1;
```

The execution plan after the change is as follows:

```
QUERY PLAN
-----
Sort
  Sort Key: ssq.cnt, t1.c1
  CTE ssq
    -> HashAggregate
      Group By Key: t2.c1
      -> Seq Scan on t2
    -> Hash Join
      Hash Cond: (t1.c1 = t3.c1)
      -> Hash Left Join
        Hash Cond: (t1.c1 = ssq.c1)
        -> Seq Scan on t1
        -> Hash
          -> CTE Scan on ssq
      -> Hash
        -> Seq Scan on t3
(15 rows)
```

– Non-equivalent correlated scenarios

Example:

```
select t1.c1, t1.c2
from t1
where t1.c1 = (select agg() from t2.c2 > t1.c2);
```

Non-equivalent correlated subqueries cannot be pulled up. You can perform join twice (one CorrelationKey and one rownum self-join) to rewrite the statement.

You can rewrite the statement in either of the following ways:

■ Subquery rewriting

```
select t1.c1, t1.c2
from t1, (
  select t1.rowid, agg() aggref
  from t1,t2
  where t1.c2 > t2.c2 group by t1.rowid
) dt /* derived table */
where t1.rowid = dt.rowid AND t1.c1 = dt.aggref;
```

■ CTE rewriting

```
WITH dt as
(
  select t1.rowid, agg() aggref
  from t1,t2
  where t1.c2 > t2.c2 group by t1.rowid
)
select t1.c1, t1.c2
from t1, derived_table
where t1.rowid = derived_table.rowid AND
t1.c1 = derived_table.aggref;
```

NOTICE

- If the AGG type is **COUNT(*)**, **0** is used for data padding when **CASE-WHEN** is not matched. If the type is not **COUNT(*)**, **NULL** is used.
- CTE rewriting works better by using sharescan.

More Optimization Examples

Example: Modify the SELECT statement by modifying the subquery to a join relationship between the main table and the parent query or modifying the subquery to improve the query performance. Ensure that the subquery to be used is semantically correct.

```
openGauss=# explain (costs off) select * from t1 where t1.c1 in (select t2.c1 from t2 where t1.c1 = t2.c2);
QUERY PLAN
-----
Seq Scan on t1
  Filter: (SubPlan 1)
  SubPlan 1
    -> Seq Scan on t2
        Filter: (t1.c1 = c2)
(5 rows)
```

In the preceding example, a subplan is used. To remove the subplan, you can modify the statement as follows:

```
openGauss=# explain (costs off) select * from t1 where exists (select t2.c1 from t2 where t1.c1 = t2.c2 and t1.c1 = t2.c1);
QUERY PLAN
-----
Hash Join
  Hash Cond: (t1.c1 = t2.c2)
  -> Seq Scan on t1
  -> Hash
    -> HashAggregate
      Group By Key: t2.c2, t2.c1
      -> Seq Scan on t2
          Filter: (c2 = c1)
(8 rows)
```

In this way, the subplan is replaced by the hash-join between the two tables, greatly improving the execution efficiency.

6.6.3 Optimizing Statistics

Context

GaussDB generates optimal execution plans based on the cost estimation. Optimizers need to estimate the number of data rows and the cost based on statistics collected using **ANALYZE**. Therefore, the statistics is vital for the estimation of the number of rows and cost. Global statistics are collected using **ANALYZE: relpages** and **reltuples** in the **pg_class** table; **stadistinct**, **stanullfrac**, **stanumbersN**, **stavaluesN**, and **histogram_bounds** in the **pg_statistic** table.

Example 1: Poor Query Performance Due to the Lack of Statistics

In most cases, the lack of statistics about tables or columns involved in the query greatly affects the query performance.

The table structure is as follows:

```
CREATE TABLE LINEITEM
(
  L_ORDERKEY      BIGINT      NOT NULL
, L_PARTKEY       BIGINT      NOT NULL
, L_SUPPKEY       BIGINT      NOT NULL
, L_LINENUMBER    BIGINT      NOT NULL
, L_QUANTITY       DECIMAL(15,2) NOT NULL
, L_EXTENDEDPRICE DECIMAL(15,2) NOT NULL
, L_DISCOUNT     DECIMAL(15,2) NOT NULL
, L_TAX           DECIMAL(15,2) NOT NULL
, L_RETURNFLAG    CHAR(1)     NOT NULL
, L_LINESTATUS    CHAR(1)     NOT NULL
, L_SHIPDATE      DATE        NOT NULL
, L_COMMITDATE    DATE        NOT NULL
, L_RECEIPTDATE   DATE        NOT NULL
, L_SHIPINSTRUCT  CHAR(25)    NOT NULL
, L_SHIPMODE      CHAR(10)    NOT NULL
, L_COMMENT       VARCHAR(44) NOT NULL
);

CREATE TABLE ORDERS
(
  O_ORDERKEY      BIGINT      NOT NULL
, O_CUSTKEY       BIGINT      NOT NULL
, O_ORDERSTATUS   CHAR(1)     NOT NULL
, O_TOTALPRICE    DECIMAL(15,2) NOT NULL
, O_ORDERDATE     DATE        NOT NULL
, O_ORDERPRIORITY CHAR(15)    NOT NULL
, O_CLERK         CHAR(15)    NOT NULL
, O_SHIPPRIORITY  BIGINT      NOT NULL
, O_COMMENT       VARCHAR(79) NOT NULL
);
```

The query statements are as follows:

```
explain verbose select
count(*) as numwait
from
lineitem l1,
orders
where
o_orderkey = l1.l_orderkey
and o_orderstatus = 'F'
and l1.l_receiptdate > l1.l_commitdate
and not exists (
select
*
from
lineitem l3
where
l3.l_orderkey = l1.l_orderkey
and l3.l_suppkey <> l1.l_suppkey
and l3.l_receiptdate > l3.l_commitdate
)
order by
numwait desc;
```

If such an issue occurs, you can use the following methods to check whether statistics in tables or columns has been collected using **ANALYZE**.

1. Execute **EXPLAIN VERBOSE** to analyze the execution plan and check the warning information:

```
WARNING:Statistics in some tables or columns(public.lineitem.l_receiptdate,
public.lineitem.l_commitdate, public.lineitem.l_orderkey, public.lineitem.l_suppkey,
public.orders.o_orderstatus, public.orders.o_orderkey) are not collected.
HINT:Do analyze for them in order to generate optimized plan.
```

2. Check whether the following information exists in the log file in the **pg_log** directory. If it does, the poor query performance was caused by the lack of statistics in some tables or columns.
2017-06-14 17:28:30.336 CST 140644024579856 20971684 [BACKEND] LOG:Statistics in some tables or columns(public.lineitem.l_receiptdate, public.lineitem.l_commitdate, public.lineitem.l_orderkey, public.lineitem.l_suppkey, public.orders.o_orderstatus, public.orders.o_orderkey) are not collected.
2017-06-14 17:28:30.336 CST 140644024579856 20971684 [BACKEND] HINT:Do analyze for them in order to generate optimized plan.

By using any of the preceding methods, you can identify tables or columns whose statistics have not been collected using **ANALYZE**. You can execute **ANALYZE** to warnings or tables and columns recorded in logs to resolve the problem that the query becomes slow because statistics are not collected.

6.6.4 Optimizing Operators

Context

A query statement needs to go through multiple operator procedures to generate the final result. Sometimes, the overall query performance deteriorates due to long execution time of certain operators, which are regarded as bottleneck operators. In this case, you need to execute the **EXPLAIN ANALYZE** or **EXPLAIN PERFORMANCE** command to view the bottleneck operators, and then perform optimization.

For example, in the following execution process, the execution time of the **Hashagg** operator accounts for about 66% $[(51016-13535)/56476 \approx 66\%]$ of the total execution time. Therefore, the **Hashagg** operator is the bottleneck operator for this query. Optimize this operator first.

id	operation	A-time	A-rows	E-rows	Peak Memory	E-memory	A-width	E-width	E-costs
1	-> Row Adapter	56476.397	100000000	237060	19KB			20	2093222.75
2	-> Vector Streaming (type: GATHER)	55664.220	100000000	237060	243KB			20	2093222.75
3	-> Vector Hash Aggregate	55124.685, 55132.180	100000000	237060	[29349KB, 29441KB]	10MB	[20,20]	20	20918406.50
4	-> Vector Streaming(type: REDISTRIBUTE)	[55519.784, 57709.779]	339264604	4856184	[12199KB, 32199KB]	3MB		20	10461210.85
5	-> Vector Hash Aggregate	[35675.636, 51016.424]	339264604	4856184	[732850KB, 746894KB]	16MB	[20,20]	20	10457195.65
6	-> Vector Partition Iterator	[9035.202, 13565.884]	970000000	935838097	[9KB, 9KB]	1MB		20	10195891.68
7	-> Partitioned CStore Scan on xujie_gp_day_energy_mv_1	[9015.645, 13535.346]	970000000	935838097	[845KB, 845KB]	1MB		20	10195891.68

Operator Optimization Examples

Example 1: Scan the base table. For queries requiring large volume of data filtering, such as point queries or queries that need range scanning, a full table scan using SeqScan will take a long time. To facilitate scanning, you can create indexes on the condition column and select IndexScan.

```
openGauss=# explain (analyze on, costs off) select * from store_sales where ss_sold_date_sk = 2450944;
```

id	operation	A-time	A-rows	Peak Memory	A-width
1	-> Streaming (type: GATHER)	3666.020	3360	195KB	
2	-> Seq Scan on store_sales	[3594.611, 3594.611]	3360	[34KB, 34KB]	

(2 rows)

Predicate Information (identified by plan id)

```
2 --Seq Scan on store_sales
  Filter: (ss_sold_date_sk = 2450944)
 Rows Removed by Filter: 4968936
```

```
openGauss=# create index idx on store_sales_row(ss_sold_date_sk);
```

```
CREATE INDEX
```

```
openGauss=# explain (analyze on, costs off) select * from store_sales_row where ss_sold_date_sk = 2450944;
```

id	operation	A-time	A-rows	Peak Memory	A-width

```
+-----+-----+-----+-----+
1 | -> Streaming (type: GATHER) | 81.524 | 3360 | 195KB |
2 | -> Index Scan using idx on store_sales_row | [13.352,13.352] | 3360 | [34KB, 34KB] |
(2 rows)
```

In this example, the full table scan filters much data and returns 3360 records. After an index has been created on the **ss_sold_date_sk** column, the scanning efficiency is significantly boosted from 3.6s to 13 ms by using **IndexScan**.

Example 2: If a nested-loop join is used for joining tables with a large number of rows, the join may take a long time. In the following example, the nested loop takes 181s. If **enable_mergejoin** is set to **off** to disable merge join and **enable_nestloop** is set to **off** to disable the nested loop so that the optimizer selects hash join, the join takes more than 200 ms.

```
openGauss=# explain analyze select count(*) from store_sales ss, item i where ss.ss_item_sk = i.i_item_sk;
id | operation | A-time | A-rows | E-rows | Peak Memory | E-memory | A-width | E-width | E-costs
---+-----+-----+-----+-----+-----+-----+-----+-----+
1 | -> Row Adapter | 184300.301 | 1 | 1 | 11KB | | | | 0 | 48629179.77
2 | -> Vector Aggregate | 184300.280 | 1 | 1 | 181KB | | | | 0 | 48629179.77
3 | -> Vector Streaming (type: GATHER) | 184300.186 | 4 | 4 | 188KB | | | | 0 | 48629179.77
4 | -> Vector Aggregate | [165575.384,184252.369] | 4 | 4 | [140KB, 140KB] | 1MB | | | 0 | 48629179.61
5 | -> Vector Hash Join (6,7) | [162915.848,181438.162] | 2880404 | 2880404 | [74KB, 74KB] | 1MB | | | 0 | 48627379.35
6 | -> CStore Scan on store_sales ss | [15.660,16.229] | 2880404 | 2880404 | [490KB, 490KB] | 1MB | | | 4 | 16683.10
7 | -> Vector Materialize | [118314.521,132478.454] | 12968211302 | 18000 | [869KB, 900KB] | 16MB | [8,8] | | 4 | 3890.00
8 | -> CStore Scan on item i | [0.234,0.243] | 18000 | 18000 | [476KB, 476KB] | 1MB | | | 4 | 3867.50
(8 rows)
```

```
openGauss=# set enable_nestloop=off;
SET
openGauss=# set enable_mergejoin=off;
SET
openGauss=# explain analyze select count(*) from store_sales ss, item i where ss.ss_item_sk = i.i_item_sk;
id | operation | A-time | A-rows | E-rows | Peak Memory | E-memory | A-width | E-width | E-costs
---+-----+-----+-----+-----+-----+-----+-----+-----+
1 | -> Row Adapter | 291.066 | 1 | 1 | 11KB | | | | 0 | 32308.66
2 | -> Vector Aggregate | 291.052 | 1 | 1 | 181KB | | | | 0 | 32308.66
3 | -> Vector Streaming (type: GATHER) | 290.973 | 4 | 4 | 188KB | | | | 0 | 32308.66
4 | -> Vector Aggregate | [220.792,234.532] | 4 | 4 | [140KB, 140KB] | 1MB | | | 0 | 32308.50
5 | -> Vector Hash Join (6,7) | [209.987,223.345] | 2880404 | 2880404 | [236KB, 241KB] | 16MB | [8,8] | | 0 | 30508.24
6 | -> CStore Scan on store_sales ss | [13.132,13.717] | 2880404 | 2880404 | [490KB, 490KB] | 1MB | | | 4 | 16683.10
7 | -> CStore Scan on item i | [0.214,0.246] | 18000 | 18000 | [477KB, 477KB] | 1MB | | | 4 | 3867.50
(7 rows)
```

Example 3: Generally, query performance can be improved by selecting HashAgg. If Sort and GroupAgg are used for a large result set, you need to set **enable_sort** to **off**. HashAgg consumes less time than Sort and GroupAgg.

```
openGauss=# explain analyze select count(*) from store_sales group by ss_item_sk;
id | operation | A-time | A-rows | E-rows | Peak Memory | E-memory | A-width | E-width | E-costs
---+-----+-----+-----+-----+-----+-----+-----+-----+
1 | -> Row Adapter | 1977.385 | 18000 | 17644 | 20KB | | | | 4 | 92875.24
2 | -> Vector Streaming (type: GATHER) | 1973.617 | 18000 | 17644 | 1946KB | | | | 4 | 92875.24
3 | -> Vector Sort Aggregate | [1784.800,1883.243] | 18000 | 17644 | [273KB, 273KB] | 1MB | | | 4 | 92136.02
4 | -> Vector Sort | [1752.270,1848.357] | 2880404 | 2880404 | [12846KB, 135135KB] | 16MB | [8,8] | | 4 | 88541.40
5 | -> CStore Scan on store_sales | [12.483,13.548] | 2880404 | 2880404 | [490KB, 490KB] | 1MB | | | 4 | 16683.10
(5 rows)
```

```
openGauss=# set enable_sort=off;
SET
openGauss=# explain analyze select count(*) from store_sales group by ss_item_sk;
id | operation | A-time | A-rows | E-rows | Peak Memory | E-memory | A-width | E-width | E-costs
---+-----+-----+-----+-----+-----+-----+-----+-----+
1 | -> Row Adapter | 838.215 | 18000 | 17644 | 20KB | | | | 4 | 21016.93
2 | -> Vector Streaming (type: GATHER) | 834.264 | 18000 | 17644 | 229KB | | | | 4 | 21016.93
3 | -> Vector Hash Aggregate | [585.017,758.204] | 18000 | 17644 | [262552KB, 262564KB] | 16MB | [8,8] | | 4 | 20327.72
4 | -> CStore Scan on store_sales | [12.540,13.941] | 2880404 | 2880404 | [490KB, 490KB] | 1MB | | | 4 | 16683.10
(4 rows)
```

6.7 Experience in Rewriting SQL Statements

Based on the SQL execution mechanism and a large number of practices, SQL statements can be optimized by following certain rules to enable the database to execute SQL statements more quickly and obtain correct results. You can comply with these rules to improve service query efficiency.

- Replace **UNION** with **UNION ALL**.

UNION eliminates duplicate rows while merging two result sets but **UNION ALL** merges the two result sets without deduplication. Therefore, replace **UNION** with **UNION ALL** if you are sure that the two result sets do not contain duplicate rows based on the service logic.

- Add **NOT NULL** to the join columns.

If there are many NULL values in the **JOIN** columns, you can add the filter criterion **IS NOT NULL** to filter data in advance to improve the **JOIN** efficiency.

- Convert **NOT IN** to **NOT EXISTS**.

nestloop anti join must be used to implement **NOT IN**, and **hash anti join** is required for **NOT EXISTS**. If no NULL value exists in the **JOIN** columns, **NOT IN** is equivalent to **NOT EXISTS**. Therefore, if you are sure that no NULL value exists, you can convert **NOT IN** to **NOT EXISTS** to generate **hash join** and to improve the query performance.

As shown in the following statement, if the **t2.d2** column does not contain null (the **t2.d2** column is not null in the table definition), the query can be modified as follows:

```
SELECT * FROM t1 WHERE NOT EXISTS (SELECT * FROM t2 WHERE t1.c1=t2.d2);
```

The generated execution plan is as follows:

QUERY PLAN

```
-----  
Hash Anti Join  
Hash Cond: (t1.c1 = t2.d2)  
-> Seq Scan on t1  
-> Hash  
-> Seq Scan on t2  
(5 rows)
```

- Use **hashagg**.

If a plan involving groupAgg and SORT operations generated by the **GROUP BY** statement is poor in performance, you can set **work_mem** to a larger value to generate a **hashagg** plan, which does not require sorting and improves the performance.

- Replace functions with **CASE** statements.

The GaussDB performance greatly deteriorates if a large number of functions are called. In this case, you can change the pushdown functions to **CASE** statements.

- Do not use functions or expressions for indexes.

Using functions or expressions for indexes stops indexing. Instead, it enables scanning on the full table.

- Do not use **!=** or **<>** operators, **NULL**, **OR**, or implicit parameter conversion in **WHERE** clauses.

- If the values of **>=** and **<=** are the same in **WHERE** condition, change the condition to **=** because range equivalence class derivation is not supported currently.

For example, change **SELECT * FROM t1 WHERE c1 >= 1 AND c1 <= 1** to **SELECT * FROM t1 WHERE c1 = 1**.

For range queries, the optimizer has a larger error when calculating selectivity than equivalent queries. Therefore, change range queries to equivalent queries as much as possible.

- Split complex SQL statements.

You can split an SQL statement into several ones and save the execution result to a temporary table if the SQL statement is too complex to be tuned using the solutions above, including but not limited to the following scenarios:

- The same subquery is involved in multiple SQL statements of a job and the subquery contains large amounts of data.
- Incorrect plan cost causes a small hash bucket of subquery. For example, the actual number of rows is 10 million, but only 1000 rows are in hash bucket.
- Functions such as **substr** and **to_number** cause incorrect measures for subqueries containing large amounts of data.

6.8 Configuring Key Parameters for SQL Optimization

This section describes key configuration parameters of the primary database node, which affect GaussDB SQL optimization.

Table 6-2 Parameters of the primary database node

Parameter/ Reference Value	Description
enable_nestloop=off	Specifies how the optimizer uses nested-loop join. If this parameter is set to on, the optimizer preferentially uses nested-loop join. If it is set to off, the optimizer preferentially uses other methods, if any. NOTE To temporarily change the value of this parameter for the current database connection (session), execute the following SQL statement: <pre>SET enable_nestloop to off;</pre> Disabling this parameter does not affect the index nested-loop join operator (corresponding to the enable_index_nestloop parameter). During real-world optimization, it is important to assess whether to disable this parameter based on the specific circumstances. Three types of joins are supported: nested loops, merge joins, and hash joins. Specifically, nested loops are best suited for scenarios with small data volumes or indexes, while hash joins are ideal for big data analysis.
enable_index_nestloop=on	Specifies whether to enable the index nested-loop join operator. NOTE To temporarily change the value of this parameter for the current database connection (session), execute the following SQL statement: <pre>SET enable_index_nestloop to off;</pre> Differences between enable_index_nestloop and enable_nestloop: enable_index_nestloop helps manage scenarios where the inner table undergoes index scans, and the outer table supplies index condition parameters for the inner table.

Parameter/ Reference Value	Description
enable_bitmapscan=on	Specifies whether the optimizer uses bitmap scanning. If the value is on, bitmap scanning is used. If the value is off, it is not used. NOTE To temporarily change the value of this parameter for the current database connection (session), execute the following SQL statement: SET enable_bitmapscan to off; The bitmap scanning applies only in the query condition where a > 1 and b > 1 and indexes are created on columns a and b. During optimization, if the query performance is poor and bitmapscan operators are in the execution plan, set this parameter to off and check whether the performance is improved.
enable_hashagg=on	Specifies whether to enable the optimizer's use of Hash-aggregation plan types.
enable_hashjoin=on	Specifies whether to enable the optimizer's use of Hash-join plan types.
enable_mergejoin=on	Specifies whether to enable the optimizer's use of Hash-merge plan types.
enable_indexscan=on	Specifies whether the optimizer uses index scan plans.
enable_indexonlyscan=on	Specifies whether the optimizer uses index-only scan plans.
enable_seqscan=on	Specifies whether the optimizer uses bitmap scanning. It is impossible to suppress sequential scans entirely, but setting this variable to off encourages the optimizer to choose other methods if available.
enable_sort=on	Specifies the optimizer sorts. It is impossible to fully suppress explicit sorts, but setting this variable to off allows the optimizer to preferentially choose other methods if available.
rewrite_rule	Specifies whether the optimizer enables the following rewriting rules: LAZYAGG, MAGICSET, UNIQUECHECK, INTARGETLIST, PREDPUSHNORMAL, PREDPUSHFORCE, PREDPUSH, DISABLE_PULLUP_EXPR_SUBLINK, DISABLE_PULLUP_NOT_IN_SUBLINK, and ENABLE_SUBLINK_PULLUP_ROWNUM.

Parameter/ Reference Value	Description
sql_beta_feature	Specifies whether the optimizer enables the following beta features: SEL_SEMI_POISSON, SEL_EXPR_INSTR, PARAM_PATH_GEN, RAND_COST_OPT, PARAM_PATH_OPT, PAGE_EST_OPT, CANONICAL_PATHKEY, PREDPUSH_SAME_LEVEL, PARTITION_FDW_ON, and DISABLE_BITMAP_COST_WITH_LOSSY_PAGES.

6.9 Hint-based Tuning

6.9.1 Plan Hint Optimization

In plan hints, you can specify a join order, join and scan operations, and the number of rows in a result to tune an execution plan, improving query performance.

GaussDB also provides the SQL patch function. You can create an SQL patch to make hints take effect without modifying service statements.

Description

Plan hints are specified in the following format after keywords such as SELECT, INSERT, UPDATE, DELETE, and MERGE:

```
/*+ <plan hint>*/
```

You can specify multiple hints for a query plan and separate them by spaces. A hint specified for a query plan does not apply to its subquery plans. To specify a hint for a subquery, add the hint following the **SELECT** of this subquery.

Example:

```
select /*+ <plan_hint1> <plan_hint2> */ * from t1, (select /*+ <plan_hint3> */ * from t2) where 1=1;
```

In the preceding command, *<plan_hint1>* and *<plan_hint2>* are the hints of a query, and *<plan_hint3>* is the hint of its subquery.

You can use the EXPLAIN syntax to analyze the plan hint optimization effect. You can use EXPLAIN to view the plan of the target SQL statement after the plan hint is used and check whether the plan meets the requirements to verify the plan hint effect. EXPLAIN has multiple plan display modes, which are controlled by **explain_perf_mode**. In some examples in this section, **explain_perf_mode** is set to **pretty** to display complete plan information. In some examples, **explain_perf_mode** is set to **normal** to simplify the output information.

NOTICE

If a hint is specified in the **CREATE VIEW** statement, the hint will be applied each time this view is used.

If the random plan function is enabled (**plan_mode_seed** is set to a value other than 0), the specified hint will not be used.

Scope

Currently, the following hints are supported:

- Join order hints (**leading**).
- Join operation hints, excluding the **semi join**, **anti join**, and **unique plan** hints.
- Rows hints.
- Scan operation hints, supporting only the **tablescan**, **indexscan**, and **indexonlyscan** hints.
- Sublink name hints.

Precautions

Hints do not support **Agg**, **Sort**, **Setop**, or **Subplan**.

Examples

The following is the original plan and is used for comparing with the optimized ones:

```
create table store
(
  s_store_sk          integer          not null,
  s_store_id          char(16)         not null,
  s_rec_start_date    date              ,
  s_rec_end_date      date              ,
  s_closed_date_sk    integer          ,
  s_store_name        varchar(50)      ,
  s_number_employees  integer          ,
  s_floor_space       integer          ,
  s_hours             char(20)         ,
  s_manager           varchar(40)      ,
  s_market_id         integer          ,
  s_geography_class   varchar(100)     ,
  s_market_desc       varchar(100)     ,
  s_market_manager    varchar(40)     ,
  s_division_id       integer          ,
  s_division_name     varchar(50)      ,
  s_company_id        integer          ,
  s_company_name      varchar(50)      ,
  s_street_number     varchar(10)      ,
  s_street_name       varchar(60)      ,
  s_street_type       char(15)         ,
  s_suite_number      char(10)         ,
  s_city              varchar(60)      ,
  s_county            varchar(30)      ,
  s_state             char(2)          ,
  s_zip              char(10)          ,
  s_country           varchar(20)      ,
  s_gmt_offset        decimal(5,2)     ,
```



```
s_tax_percentage    decimal(5,2)      ,
primary key (s_store_sk)
);
create table store_sales
(
  ss_sold_date_sk    integer          ,
  ss_sold_time_sk    integer          ,
  ss_item_sk         integer          not null,
  ss_customer_sk     integer          ,
  ss_cdemo_sk        integer          ,
  ss_hdemo_sk        integer          ,
  ss_addr_sk         integer          ,
  ss_store_sk        integer          ,
  ss_promo_sk        integer          ,
  ss_ticket_number    integer          not null,
  ss_quantity        integer          ,
  ss_wholesale_cost   decimal(7,2)    ,
  ss_list_price       decimal(7,2)    ,
  ss_sales_price      decimal(7,2)    ,
  ss_ext_discount_amt decimal(7,2)    ,
  ss_ext_sales_price  decimal(7,2)    ,
  ss_ext_wholesale_cost decimal(7,2)  ,
  ss_ext_list_price   decimal(7,2)    ,
  ss_ext_tax          decimal(7,2)    ,
  ss_coupon_amt       decimal(7,2)    ,
  ss_net_paid         decimal(7,2)    ,
  ss_net_paid_inc_tax decimal(7,2)    ,
  ss_net_profit       decimal(7,2)    ,
  primary key (ss_item_sk, ss_ticket_number)
);
create table store_returns
(
  sr_returned_date_sk integer          ,
  sr_return_time_sk   integer          ,
  sr_item_sk          integer          not null,
  sr_customer_sk      integer          ,
  sr_cdemo_sk         integer          ,
  sr_hdemo_sk         integer          ,
  sr_addr_sk          integer          ,
  sr_store_sk         integer          ,
  sr_reason_sk        integer          ,
  sr_ticket_number     integer          not null,
  sr_return_quantity   integer          ,
  sr_return_amt        decimal(7,2)    ,
  sr_return_tax        decimal(7,2)    ,
  sr_return_amt_inc_tax decimal(7,2)    ,
  sr_fee              decimal(7,2)    ,
  sr_return_ship_cost  decimal(7,2)    ,
  sr_refunded_cash     decimal(7,2)    ,
  sr_reversed_charge   decimal(7,2)    ,
  sr_store_credit      decimal(7,2)    ,
  sr_net_loss          decimal(7,2)    ,
  primary key (sr_item_sk, sr_ticket_number)
);
create table customer
(
  c_customer_sk      integer          not null,
  c_customer_id      char(16)         not null,
  c_current_cdemo_sk integer          ,
  c_current_hdemo_sk integer          ,
  c_current_addr_sk  integer          ,
  c_first_shipto_date_sk integer          ,
  c_first_sales_date_sk integer          ,
  c_salutation       char(10)         ,
  c_first_name       char(20)         ,
  c_last_name        char(30)         ,
  c_preferred_cust_flag char(1)       ,
  c_birth_day        integer          ,
  c_birth_month      integer          ,
```

```
c_birth_year      integer          ,
c_birth_country   varchar(20)      ,
c_login           char(13)         ,
c_email_address   char(50)         ,
c_last_review_date char(10)        ,
primary key (c_customer_sk)
);
create table promotion
(
  p_promo_sk       integer          not null,
  p_promo_id       char(16)         not null,
  p_start_date_sk  integer          ,
  p_end_date_sk    integer          ,
  p_item_sk        integer          ,
  p_cost           decimal(15,2)    ,
  p_response_target integer          ,
  p_promo_name     char(50)         ,
  p_channel_dmail  char(1)          ,
  p_channel_email  char(1)          ,
  p_channel_catalog char(1)         ,
  p_channel_tv     char(1)          ,
  p_channel_radio  char(1)          ,
  p_channel_press  char(1)          ,
  p_channel_event  char(1)          ,
  p_channel_demo   char(1)          ,
  p_channel_details varchar(100)    ,
  p_purpose          char(15)         ,
  p_discount_active char(1)         ,
  primary key (p_promo_sk)
);
create table customer_address
(
  ca_address_sk    integer          not null,
  ca_address_id    char(16)         not null,
  ca_street_number char(10)         ,
  ca_street_name   varchar(60)      ,
  ca_street_type   char(15)         ,
  ca_suite_number  char(10)         ,
  ca_city          varchar(60)      ,
  ca_county        varchar(30)      ,
  ca_state         char(2)          ,
  ca_zip           char(10)         ,
  ca_country       varchar(20)      ,
  ca_gmt_offset    decimal(5,2)     ,
  ca_location_type char(20)         ,
  primary key (ca_address_sk)
);
create table item
(
  i_item_sk        integer          not null,
  i_item_id        char(16)         not null,
  i_rec_start_date date              ,
  i_rec_end_date   date              ,
  i_item_desc      varchar(200)     ,
  i_current_price  decimal(7,2)     ,
  i_wholesale_cost decimal(7,2)     ,
  i_brand_id       integer          ,
  i_brand          char(50)         ,
  i_class_id       integer          ,
  i_class          char(50)         ,
  i_category_id    integer          ,
  i_category       char(50)         ,
  i_manufact_id    integer          ,
  i_manufact       char(50)         ,
  i_size           char(20)         ,
  i_formulation    char(20)         ,
  i_color          char(20)         ,
  i_units          char(10)         ,
  i_container      char(10)         ,
```

```

    i_manager_id      integer      ,
    i_product_name    char(50)    ,
    primary key (i_item_sk)
);
explain
select i_product_name product_name
,i_item_sk item_sk
,s_store_name store_name
,s_zip store_zip
,ad2.ca_street_number c_street_number
,ad2.ca_street_name c_street_name
,ad2.ca_city c_city
,ad2.ca_zip c_zip
,count(*) cnt
,sum(ss_wholesale_cost) s1
,sum(ss_list_price) s2
,sum(ss_coupon_amt) s3
FROM   store_sales
,store_returns
,store
,customer
,promotion
,customer_address ad2
,item
WHERE  ss_store_sk = s_store_sk AND
ss_customer_sk = c_customer_sk AND
ss_item_sk = i_item_sk and
ss_item_sk = sr_item_sk and
ss_ticket_number = sr_ticket_number and
c_current_addr_sk = ad2.ca_address_sk and
ss_promo_sk = p_promo_sk and
i_color in ('maroon','burnished','dim','steel','navajo','chocolate') and
i_current_price between 35 and 35 + 10 and
i_current_price between 35 + 1 and 35 + 15
group by i_product_name
,i_item_sk
,s_store_name
,s_zip
,ad2.ca_street_number
,ad2.ca_street_name
,ad2.ca_city
,ad2.ca_zip
;
HashAggregate (cost=23.52..23.53 rows=1 width=880)
  Group By Key: item.i_product_name, item.i_item_sk, store.s_store_name, store.s_zip, ad2.ca_street_number,
ad2.ca_street_name, ad2.ca_city, ad2.ca_zip
  -> Nested Loop (cost=4.27..23.49 rows=1 width=776)
    -> Nested Loop (cost=4.27..22.80 rows=1 width=416)
      -> Nested Loop (cost=4.27..22.39 rows=1 width=420)
        -> Nested Loop (cost=4.27..21.98 rows=1 width=420)
          -> Nested Loop (cost=4.27..21.57 rows=1 width=262)
            Join Filter: (item.i_item_sk = store_sales.ss_item_sk)
              -> Nested Loop (cost=4.27..20.78 rows=2 width=216)
                -> Seq Scan on item (cost=0.00..11.16 rows=1 width=208)
                  Filter: ((i_current_price >= 35::numeric) AND
(i_current_price <= 45::numeric) AND (i_current_price >= 36::numeric) AND
(i_current_price <= 50::numeric) AND (i_color = ANY
('{maroon,burnished,dim,steel,navajo,chocolate}'::bpchar[])))
                -> Bitmap Heap Scan on store_returns (cost=4.27..9.61 rows=2 width=8)
                  Recheck Cond: (sr_item_sk = item.i_item_sk)
                  -> Bitmap Index Scan pn store_returns_pkey (cost=0.00..4.27 rows=2 width=0)
                    Index Cond: (sr_item_sk = item.i_item_sk)
                -> Index Scan using store_sales_pkey on store_sales (cost=0.00..0.38 rows=1 width=62)
                  Index Cond: ((ss_item_sk = store_returns.sr_item_sk) AND (ss_ticket_number =
store_returns.sr_ticket_number))
                -> Index Scan using store_pkey on store (cost=0.00..0.40 rows=1 width=166)
                  Index Cond: (s_store_sk = store_sales.ss_store_sk)
                -> Index Scan using customer_pkey on customer (cost=0.00..0.40 rows=1 width=8)
                  Index Cond: (c_customer_sk = store_sales.ss_customer_sk)
```

```
-> Index Scan using promotion_pkey on promotion (cost=0.00..0.40 rows=1 width=4)
    Index Cond: (p_promo_sk = store_sales.ss_promo_sk)
-> Index Scan using customer_address_pkey on customer_address ad2 (cost=0.00..0.68 rows=1
width=368)
    Index Cond: (ca_address_sk = customer.c_current_addr_sk)
(25 rows)
```

6.9.2 Join Order Hints

Description

Specifies the join order and outer/inner tables.

Syntax

- Specify only the join order.

```
leading(join_table_list)
```

- Specify the join order and outer/inner tables. The outer/inner tables are specified by the outermost parentheses.

```
leading((join_table_list))
```

Parameters

join_table_list specifies the tables to be joined. The values can be table names or table aliases. If a subquery is pulled up, the value can also be the subquery alias. Separate the values with spaces. You can add parentheses to specify the join priorities of tables.

NOTICE

A table name or alias can only be a string without a schema name.

An alias (if any) is used to represent a table.

To prevent semantic errors, tables in the list must meet the following requirements:

- The tables must exist in the query or its subquery to be pulled up.
- The table names must be unique in the query or subquery to be pulled up. If they are not, their aliases must be unique.
- A table appears only once in the list.
- An alias (if any) is used to represent a table.

For example:

leading(t1 t2 t3 t4 t5): **t1**, **t2**, **t3**, **t4**, and **t5** are joined. The join order and outer/inner tables are not specified.

leading((t1 t2 t3 t4 t5)): **t1**, **t2**, **t3**, **t4**, and **t5** are joined in sequence. The table on the right is used as the inner table in each join.

leading(t1 (t2 t3 t4) t5): First, **t2**, **t3**, and **t4** are joined and the outer/inner tables are not specified. Then, the result is joined with **t1** and **t5**, and the outer/inner tables are not specified.

leading((t1 (t2 t3 t4) t5)): First, **t2**, **t3**, and **t4** are joined and the outer/inner tables are not specified. Then, the result is joined with **t1**, and **(t2 t3 t4)** is used as the inner table. Finally, the result is joined with **t5**, and **t5** is used as the inner table.

leading((t1 (t2 t3) t4 t5)) leading((t3 t2)): First, **t2** and **t3** are joined and **t2** is used as the inner table. Then, the result is joined with **t1**, and **(t2 t3)** is used as the inner table. Finally, the result is joined with **t4** and then **t5**, and the table on the right in each join is used as the inner table.

Example

Hint the query plan in [Examples](#) as follows:

explain

```
select /*+ leading((((store_sales store) promotion) item) customer) ad2) store_returns) leading((store_sales))*/ i_product_name product_name ...
```

First, **store_sales** and **store** are joined and **store_sales** is the inner table. Then, the result is joined with **promotion**, **item**, **customer**, **ad2**, and **store_returns** in sequence. The optimized plan is as follows:

```
WARNING: Duplicated or conflict hint: Leading(store_sales store), will be discarded.
QUERY PLAN
-----
HashAggregate (cost=55.24..55.25 rows=1 width=800)
  Group By Key: item.i_product_name, item.i_item_sk, store.s_store_name, store.s_zip, ad2.ca_street_number, ad2.ca_street_name, ad2.ca_city, ad2.ca_zip
  -> Nested Loop (cost=29.92..55.21 rows=1 width=78)
    -> Nested Loop (cost=29.92..54.80 rows=1 width=784)
      -> Nested Loop (cost=29.92..53.70 rows=1 width=424)
        Join Filter: (store_sales.ss_item_sk = item.i_item_sk)
        -> Seq Scan on item (cost=0.00..11.16 rows=1 width=208)
          Filter: ((i_current_price >= 35::numeric) AND (i_current_price <= 45::numeric) AND (i_current_price <= 50::numeric) AND (i_color = ANY ('{maroon,burnished,dim,steel,navajo,chocolate}':bpchar)))
        -> Hash Join (cost=29.92..41.89 rows=44 width=216)
          Hash Cond: (promotion.p_promo_sk = store_sales.ss_promo_sk)
          -> Seq Scan on promotion (cost=0.00..11.16 rows=118 width=4)
          -> Hash (cost=29.00..29.00 rows=74 width=220)
            -> Hash Join (cost=17.61..29.00 rows=74 width=220)
              Hash Cond: (store.s_store_sk = store_sales.ss_store_sk)
              -> Seq Scan on store (cost=0.00..10.44 rows=44 width=166)
              -> Hash (cost=13.38..13.38 rows=338 width=62)
                -> Seq Scan on store_sales (cost=0.00..13.38 rows=338 width=62)
            -> Index Scan using customer_pkey on customer (cost=0.00..0.40 rows=1 width=8)
              Index Cond: (c_customer_sk = store_sales.ss_customer_sk)
          -> Index Scan using customer_address_pkey on customer_address ad2 (cost=0.00..0.68 rows=1 width=368)
              Index Cond: (ca_address_sk = customer.c_current_addr_sk)
          -> Index Only Scan using store_returns_pkey on store_returns (cost=0.00..0.41 rows=1 width=8)
              Index Cond: ((sr_item_sk = store_sales.ss_item_sk) AND (sr_ticket_number = store_sales.ss_ticket_number))
(24 rows)
```

For details about the warning at the top of the plan, see [Hint Errors, Conflicts, and Other Warnings](#).

6.9.3 Join Operation Hints

Description

Specifies the join method, which can be nested-loop join, hash join, or merge join.

Syntax

```
[no] nestloop|hashjoin|mergejoin(table_list)
```

Parameters

- no** indicates that the specified hint will not be used for a join.
- table_list** specifies the tables to be joined. The values are the same as those of [join_table_list](#) but contain no parentheses.

For example:

no nestloop(t1 t2 t3): **nestloop** is not used for joining **t1**, **t2**, and **t3**. The three tables may be joined in either of the two ways: Join **t2** and **t3**, and then **t1**; join **t1**

and **t2**, and then **t3**. This hint takes effect only for the last join. If necessary, you can hint other joins. For example, you can add **no nestloop(t2 t3)** to join **t2** and **t3** first and to forbid the use of **nestloop**.

Example

Hint the query plan in [Examples](#) as follows:

```
explain
select /*+ nestloop(store_sales store_returns item) */ i_product_name product_name ...
```

nestloop is used for the last join between **store_sales**, **store_returns**, and **item**. The optimized plan is as follows:

```
QUERY PLAN
HashAggregate (cost=23.52..23.53 rows=1 width=880)
  Group By Key: item.i_product_name, item.i_item_sk, store.s_store_name, store.s_zip, ad2.ca_street_number, ad2.ca_street_name, ad2.ca_city, ad2.ca_zip
  -> Nested Loop (cost=4.27..22.48 rows=1 width=776)
    -> Nested Loop (cost=4.27..22.80 rows=1 width=416)
      -> Nested Loop (cost=4.27..22.39 rows=1 width=420)
        -> Nested Loop (cost=4.27..21.98 rows=1 width=262)
          Join Filter: (item.i_item_sk = store_sales.ss_item_sk)
          -> Nested Loop (cost=4.27..20.78 rows=2 width=216)
            -> Seq Scan on item (cost=0.00..11.16 rows=1 width=288)
              Filters: ((i_current_price >= 35::numeric) AND (i_current_price <= 45::numeric) AND (i_current_price >= 36::numeric) AND (i_current_price <= 58::numeric) AND (i_color = ANY ({maroon,burnished,dim,steel,navajo,chocolate}::bpchar[])))
            -> Bitmap Heap Scan on store_returns (cost=4.27..9.61 rows=2 width=8)
              Recheck Cond: (sr_item_sk = item.i_item_sk)
              -> Bitmap Index Scan on store_returns_pkey (cost=0.00..4.27 rows=2 width=0)
                Index Cond: (sr_item_sk = item.i_item_sk)
          -> Index Scan using store_sales_pkey on store_sales (cost=0.00..0.38 rows=1 width=62)
            Index Cond: ((ss_item_sk = store_returns.sr_item_sk) AND (ss_ticket_number = store_returns.sr_ticket_number))
        -> Index Scan using store_pkey on store (cost=0.00..0.40 rows=1 width=168)
          Index Cond: (s_store_sk = store_sales.ss_store_sk)
      -> Index Scan using customer_pkey on customer (cost=0.00..0.40 rows=1 width=8)
        Index Cond: (c_customer_sk = store_sales.ss_customer_sk)
    -> Index Only Scan using promotion_pkey on promotion (cost=0.00..0.40 rows=1 width=4)
      Index Cond: (p_promo_sk = store_sales.ss_promo_sk)
  -> Index Scan using customer_address_pkey on customer_address ad2 (cost=0.00..0.68 rows=1 width=368)
    Index Cond: (ca_address_sk = customer.c_current_addr_sk)
```

6.9.4 Rows Hints

Description

Specifies the number of rows in an intermediate result set. Both absolute values and relative values are supported.

Syntax

```
rows(table_list #|+|-|* const)
```

Parameters

- #**, **+**, **-**, and ***** are operators used for hinting the estimation. **#** indicates that the original estimation is used without any calculation. **+**, **-**, and ***** indicate that the original estimation is calculated using these operators. The minimum calculation result is 1. **table_list** specifies the tables to be joined. The values are the same as those of **table_list** in [Join Operation Hints](#).
- const** can be any non-negative number and supports scientific notation.

Example:

rows(t1 #5): The result set of **t1** is five rows.

rows(t1 t2 t3 *1000): Multiply the result set of joined **t1**, **t2**, and **t3** by 1000.

Suggestion

- The hint using ***** for two tables is recommended. This hint will be triggered if the two tables appear on two sides of a join. For example, if the hint is **rows(t1 t2 * 3)**, the join result of **(t1 t3 t4)** and **(t2 t5 t6)** will be multiplied by 3 because **t1** and **t2** appear on both sides of the join.


- **rows** hints can be specified for the result sets of a single table, multiple tables, function tables, and subquery scan tables.

Examples

Hint the query plan in [Examples](#) as follows:

```
explain
select /*+ rows(store_sales store_returns *50) */ i_product_name product_name ...
```

Multiply the result set of joined **store_sales** and **store_returns** by 50. The optimized plan is as follows:



```
QUERY PLAN
HashAggregate (cost=23.52..23.53 rows=1 width=880)
  Group By Key: item_i_product_name, item_i_item_sk, store_s_store_name, store_s_zip, ad2_ca_street_number, ad2_ca_street_name, ad2_ca_city, ad2_ca_zip
  -> Nested Loop (cost=4.27..22.80 rows=1 width=776)
    -> Nested Loop (cost=4.27..22.80 rows=1 width=416)
      -> Nested Loop (cost=4.27..22.39 rows=1 width=420)
        -> Nested Loop (cost=4.27..21.98 rows=1 width=420)
          -> Nested Loop (cost=4.27..21.57 rows=1 width=262)
            Join Filter: (item_i_item_sk = store_sales.ss_item_sk)
            -> Nested Loop (cost=4.27..20.79 rows=2 width=216)
              -> Seq Scan on item (cost=0.00..11.16 rows=1 width=208)
                Filter: ((i_current_price >= 35::numeric) AND (i_current_price <= 45::numeric) AND (i_current_price <= 50::numeric) AND (i_color = ANY ('{maroon,burnished,dim,steel,navajo,chocolate}':bpchar)))
              -> Bitmap Heap Scan on store_returns (cost=4.27..9.61 rows=2 width=8)
                Bitmap Index Scan on store_returns_pkey (cost=0.00..4.27 rows=2 width=0)
                Index Cond: (sr_item_sk = item_i_item_sk)
            -> Index Scan using store_sales_pkey on store_sales (cost=0.00..0.38 rows=1 width=62)
              Index Cond: ((ss_item_sk = store_returns.sr_item_sk) AND (ss_ticket_number = store_returns.sr_ticket_number))
          -> Index Scan using store_pkey on store (cost=0.00..0.40 rows=1 width=168)
            Index Cond: (s_store_sk = store_sales.ss_store_sk)
        -> Index Scan using customer_pkey on customer (cost=0.00..0.40 rows=1 width=8)
          Index Cond: (c_customer_sk = store_sales.ss_customer_sk)
      -> Index Only Scan using promotion_pkey on promotion (cost=0.00..0.40 rows=1 width=4)
        Index Cond: (p_promo_sk = store_sales.ss_promo_sk)
    -> Index Scan using customer_address_pkey on customer_address ad2 (cost=0.00..0.68 rows=1 width=368)
      Index Cond: (ca_address_sk = customer.c_current_addr_sk)
```

6.9.5 Scan Operation Hints

Description

These hints specify a scan operation, which can be **tablescan**, **indexscan**, or **indexonlyscan**.

Syntax

```
[no] tablescan|indexscan|indexonlyscan(table [index])
```

Parameters

- **no** specifies that the specified hint will not be used for scanning.
- **table** specifies the table to be scanned. You can specify only one table. Use a table alias (if any) instead of a table name.
- **index** specifies the index for **indexscan** or **indexonlyscan**. You can specify only one index.

NOTE

indexscan and **indexonlyscan** hints can be used only when the specified index belongs to the table.

Scan operation hints can be used for row-store tables, HDFS tables, and subquery tables. HDFS internal tables include base tables and delta tables. The delta tables are invisible to users. Therefore, scan operation hints are used only for base tables.

Example

To specify an index-based hint for a scan, create an index named **i** on the **i_item_sk** column of the **item** table.

```
create index i on item(i_item_sk);
```

Hint the query plan in [Examples](#) as follows:

```
explain
select /*+ indexscan(item i) */ i_product_name product_name ...
```

item is scanned based on an index. The optimized plan is as follows:

```
QUERY PLAN
-----
HashAggregate (cost=38.79..38.80 rows=1 width=880)
  Group By Key: item.i_product_name, item.i_item_sk, store.s_store_name, store.s_zip, ad2.ca_street_number, ad2.ca_street_name, ad2.ca_city, ad2.ca_zip
  -> Nested Loop (cost=18.45..38.76 rows=1 width=776)
    -> Nested Loop (cost=18.45..38.07 rows=1 width=416)
      -> Nested Loop (cost=18.45..37.66 rows=1 width=420)
        -> Nested Loop (cost=18.45..37.25 rows=1 width=420)
          -> Nested Loop (cost=18.45..36.84 rows=1 width=262)
            Join Filter: (store_sales.ss_item_sk = item.i_item_sk)
            -> Hash Join (cost=18.45..35.67 rows=1 width=42)
              Hash Cond: ((store_returns.sr_item_sk = store_sales.ss_item_sk) AND (store_returns.sr_ticket_number = store_sales.ss_ticket_number))
              -> Seq Scan on store_returns (cost=0.00..14.08 rows=408 width=8)
              -> Hash (cost=13.38..13.38 rows=338 width=42)
                -> Seq Scan on store_sales (cost=0.00..13.38 rows=338 width=42)
            -> Index Scan using i on item (cost=0.00..0.40 rows=1 width=208)
              Index Cond: (i_item_sk = store_returns.sr_item_sk)
          Filter: ((i.current_price >= 35::numeric) AND (i.current_price <= 45::numeric) AND (i.current_price >= 36::numeric) AND (i.current_price <= 50::numeric) AND (i.color = ANY (('{maroon,burnished,dia,steel,navajo,chocolate'}::bpchar))))
        -> Index Scan using store_pkey on store (cost=0.00..0.40 rows=1 width=166)
          Index Cond: (s_store_sk = store_sales.ss_store_sk)
      -> Index Scan using customer_pkey on customer (cost=0.00..0.40 rows=1 width=8)
        Index Cond: (c_customer_sk = store_sales.ss_customer_sk)
    -> Index Only Scan using promotion_pkey on promotion (cost=0.00..0.40 rows=1 width=4)
      Index Cond: (p_promo_sk = store_sales.ss_promo_sk)
  -> Index Scan using customer_address_pkey on customer_address ad2 (cost=0.00..0.68 rows=1 width=368)
    Index Cond: (ca_address_sk = customer.c_current_addr_sk)
(24 rows)
```

6.9.6 Sublink Name Hints

Function

These hints specify the name of a sublink block.

Syntax

```
blockname (table)
```

Parameter Description

- table** specifies the name you have specified for a sublink block.

NOTE

- The **blockname** hint is used by an outer query only when the corresponding sublink is not pulled up. Currently, only the **Agg** equivalent join, **IN**, and **EXISTS** sublinks can be pulled up. This hint is usually used together with the hints described in the previous sections.
- The subquery after the **FROM** keyword is hinted by using the subquery alias. In this case, **blockname** becomes invalid.
- If a sublink contains multiple tables, the tables will be joined with the outer-query tables in a random sequence after the sublink is pulled up. In this case, **blockname** also becomes invalid.

Example

```
explain select /*+nestloop(store_sales tt) */ * from store_sales where ss_item_sk in (select /*
+blockname(tt)*/ i_item_sk from item group by 1);
```

tt indicates the sublink block name. After being pulled up, the sublink is joined with the outer-query table **store_sales** by using **nestloop**. The optimized plan is as follows:

```
QUERY PLAN
-----
Nested Loop (cost=10.53..68.39 rows=169 width=212)
  -> HashAggregate (cost=10.53..10.95 rows=42 width=4)
    Group By Key: item.i_item_sk
    -> Seq Scan on item (cost=0.00..10.42 rows=42 width=4)
  -> Index Scan using store_sales_pkey on store_sales (cost=0.00..1.34 rows=2 width=212)
    Index Cond: (ss_item_sk = item.i_item_sk)
(6 rows)
```


6.9.7 Hint Errors, Conflicts, and Other Warnings

Plan hints change an execution plan. You can run **EXPLAIN** to view the changes.

Hints containing errors are invalid and do not affect statement execution. The errors will be displayed in different ways based on statement types. Hint errors in an **EXPLAIN** statement are displayed as a warning on the interface. Hint errors in other statements will be recorded in debug1-level logs containing the **PLANHINT** keyword.

Hint error types are as follows:

- Syntax errors

An error will be reported if the syntax tree fails to be reduced. The No. of the row generating an error is displayed in the error details.

For example, the hint keyword is incorrect, no table or only one table is specified in the **leading** or **join** hint, or no tables are specified in other hints. The parsing of a hint is terminated immediately after a syntax error is detected. Only the hints that have been parsed successfully are valid.

For example:

```
leading((t1 t2)) nestloop(t1) rows(t1 t2 #10)
```

The syntax of **nestloop(t1)** is wrong and its parsing is terminated. Only **leading(t1 t2)** that has been successfully parsed before **nestloop(t1)** is valid.

- Semantic errors

- An error will be reported if the specified tables do not exist, multiple tables are found based on the hint setting, or a table is used more than once in the **leading** or **join** hint.
- An error will be reported if the index specified in a scan hint does not exist.
- If multiple tables with the same name exist after a subquery is pulled up and some of them need to be hinted, add aliases for them to avoid name duplication.

- Duplicated or conflicted hints

If hint duplication or conflicts occur, only the first hint takes effect. A message will be displayed to describe the situation.

- Hint duplication indicates that a hint is used more than once in the same query, for example, **nestloop(t1 t2) nestloop(t1 t2)**.
- A hint conflict indicates that the functions of two hints with the same table list conflict with each other.

For example, if **nestloop (t1 t2) hashjoin (t1 t2)** is used, **hashjoin (t1 t2)** becomes invalid. **nestloop(t1 t2)** does not conflict with **no mergejoin(t1 t2)**.

NOTICE

The table list in the **leading** hint is disassembled. For example, **leading ((t1 t2 t3))** will be disassembled as **leading((t1 t2)) leading(((t1 t2) t3))**, which will conflict with **leading((t2 t1))** (if any). In this case, the latter **leading(t2 t1)** becomes invalid. If two hints use duplicated table lists and only one of them has the specified outer/inner table, the one without a specified outer/inner table becomes invalid.

- A hint becomes invalid after a sublink is pulled up.
In this case, a message will be displayed. Generally, such invalidation occurs when a sublink contains multiple tables to be joined. After the sublink is pulled up, the tables will not be join members.
- Hints are not used.
 - If a **hashjoin** or **mergejoin** hint is specified for non-equivalent joins, it will not be used.
 - If an **indexscan** or **indexonlyscan** hint is specified for a table that does not have an index, it will not be used.
 - If an **indexscan** or **indexonlyscan** hint is specified for a full-table scan, it will not be used. Generally, index paths are generated only when filtering conditions are used on index columns. Indexes are not used during a full table scan.
 - If an **indexonlyscan** hint is specified when the output or predicate condition column does not contain only indexes, it will not be used.
 - In equivalent joins, only the joins containing equivalence conditions are valid. Therefore, the **leading**, **join**, and **rows** hints specified for the joins without an equivalence condition will not be used. For example, **t1**, **t2**, and **t3** are to be joined, and the join between **t1** and **t3** does not contain an equivalence condition. In this case, **leading(t1 t3)** will not be used.
 - If no sublink is pulled up, the specified **blockname** hint will not be used.

6.9.8 Optimizer GUC Parameter Hints

Description

Sets GUC parameters related to query optimization that take effect during the query execution. For details about the application scenarios of hints, see the description of each GUC parameter.

Syntax

```
set(param value)
```

Parameters

- **param** indicates the parameter name.
- **value** indicates the value of a parameter.
- Currently, the following parameters can be set and take effect by using Hint:
 - Boolean

enable_bitmapscan, enable_hashagg, enable_hashjoin, enable_indexscan, enable_indexonlyscan, enable_material, enable_mergejoin, enable_nestloop, enable_index_nestloop, enable_seqscan, enable_sort, and enable_tidscan

- Integer type

query_dop

- Floating point

cost_weight_index, default_limit_rows, seq_page_cost, random_page_cost, cpu_tuple_cost, cpu_index_tuple_cost, cpu_operator_cost, and effective_cache_size

NOTE

- If you set a parameter that is not in the whitelist and the parameter value is invalid or the hint syntax is incorrect, the query execution is not affected. Run **explain(verbose on)**. An error message is displayed, indicating that hint parsing fails.
- The GUC parameter hint takes effect only in the outermost query. That is, the GUC parameter hint in the subquery does not take effect.
- The GUC parameter hint in the view definition does not take effect.
- In the **CREATE TABLE ... AS ...** statement, the outermost GUC parameter hint takes effect.

6.9.9 Hint for Selecting the Custom Plan or Generic Plan

Description

For query statements and DML statements executed in PBE mode, the optimizer generates a custom plan or generic plan based on factors such as rules, costs, and parameters. You can use the hint of **use_cplan** or **use_gplan** to specify the plan to execute.

Syntax

- To select the custom plan, run the following statement:
`use_cplan`
- To select the generic plan, run the following statement:
`use_gplan`

NOTE

- For SQL statements that are executed in non-PBE mode, setting this hint does not affect the execution mode.
- This hint has a higher priority than cost-based selection and the **plan_cache_mode** parameter. That is, this hint does not take effect for statements for which **plan_cache_mode** cannot be forcibly set to specify an execution mode.

Examples

Forcibly use the custom plan.

```
create table t (a int, b int, c int);
prepare p as select /*+ use_cplan */ * from t where a = $1;
explain execute p(1);
```

In the following plan, the filter condition is the actual value of the input parameter, that is, the plan is a custom plan.

QUERY PLAN

```
Seq Scan on t (cost=0.00..34.31 rows=10 width=12)
  Filter: (a = 1)
(2 rows)
```

Forcibly use the generic plan.

```
deallocate p;
prepare p as select /*+ use_gplan */ * from t where a = $1;
explain execute p(1);
```

In the following plan, the filter condition is the input parameter to be added. That is, the plan is a generic plan.

QUERY PLAN

```
Seq Scan on t (cost=0.00..34.31 rows=10 width=12)
  Filter: (a = $1)
(2 rows)
```

6.9.10 Hint Specifying Not to Expand Subqueries

Function

When the database optimizes the query logic, some subqueries can be promoted to the upper layer to avoid nested execution. However, for some subqueries that have a low selection rate and can use indexes to filter access pages, nested execution does not cause too much performance deterioration, while after the promotion, the query search scope is expanded, which may cause performance deterioration. In this case, you can use the **no_expand** hint for debugging. This hint is not recommended in most cases.

Syntax

```
no_expand
```

Examples

Normal query execution:

```
explain select * from t1 where t1.a in (select t2.a from t2);
```

Plan:

QUERY PLAN

```
Hash Join (cost=38.81..92.58 rows=972 width=12)
  Hash Cond: (t1.a = t2.a)
  -> Seq Scan on t1 (cost=0.00..29.45 rows=1945 width=12)
  -> Hash (cost=36.31..36.31 rows=200 width=4)
      -> HashAggregate (cost=34.31..36.31 rows=200 width=4)
          Group By Key: t2.a
          -> Seq Scan on t2 (cost=0.00..29.45 rows=1945 width=4)
(7 rows)
```

After **no_expand** is added:

```
explain select * from t1 where t1.a in (select /*+ no_expand*/ t2.a from t2);
```

Plan:

```
-----  
QUERY PLAN  
-----  
Seq Scan on t1 (cost=34.31..68.62 rows=972 width=12)  
  Filter: (hashed SubPlan 1)  
    SubPlan 1  
      -> Seq Scan on t2 (cost=0.00..29.45 rows=1945 width=4)  
(4 rows)
```

6.9.11 Hint Specifying Not to Use Global Plan Cache

Description

When global plan cache is enabled, you can use the `no_gpc` hint to force a single query statement not to share the plan cache globally. Only the plan cache within the current session lifecycle is retained.

Syntax

`no_gpc`

NOTE

This parameter takes effect only for statements executed by PBE when `enable_global_plancache` is set to `on`.

Example

```
openGauss=# deallocate all;  
DEALLOCATE ALL  
openGauss=# prepare insert_nogpc as insert /*+ no_gpc */ into t1 select c1, c2 from t2 where c1 = $1;  
PREPARE  
openGauss=# execute insert_nogpc(1);  
INSERT 0 1  
openGauss=# select * from db_perf.global_plancache_status where schema_name = 'schema_hint_iud' order by 1,2;  
nodename | query | refcount | valid | databaseid | schema_name | params_num | func_id  
-----+-----+-----+-----+-----+-----+-----+-----  
(0 rows)
```

No result exists in the `db_perf.global_plancache_status` view, that is, no plan is cached globally.

6.9.12 Hint of Parameterized Paths at the Same Level

Function

The `predpush_same_level` hint is used to specify the generation of parameterized paths between tables or MVs at the same level.

Syntax

```
predpush_same_level(src, dest)  
predpush_same_level(src1 src2 ..., dest)
```

NOTE

This parameter takes effect only when the `predpushforce` option in `rewrite_rule` is enabled.

Examples

Prepare parameters, tables, and indexes.

```
openGauss=# set rewrite_rule = 'predpushforce';
SET
openGauss=# create table t1(a int, b int);
CREATE TABLE
openGauss=# create table t2(a int, b int);
CREATE TABLE
openGauss=# create index idx1 on t1(a);
CREATE INDEX
openGauss=# create index idx2 on t2(a);
CREATE INDEX
```

Run the following statement to view the plan:

```
openGauss=# explain select * from t1, t2 where t1.a = t2.a;
QUERY PLAN
-----
Hash Join (cost=27.50..56.25 rows=1000 width=16)
  Hash Cond: (t1.a = t2.a)
    -> Seq Scan on t1 (cost=0.00..15.00 rows=1000 width=8)
    -> Hash (cost=15.00..15.00 rows=1000 width=8)
        -> Seq Scan on t2 (cost=0.00..15.00 rows=1000 width=8)
(5 rows)
```

The filter condition **t1.a = t2.a** is displayed on **Join**. In this case, **predpush_same_level(t1, t2)** can be used to push the condition down to the scan operator of t2.

```
openGauss=# explain select /*+predpush_same_level(t1, t2)*/ * from t1, t2 where t1.a = t2.a;
QUERY PLAN
-----
Nested Loop (cost=0.00..335.00 rows=1000 width=16)
  -> Seq Scan on t1 (cost=0.00..15.00 rows=1000 width=8)
  -> Index Scan using idx2 on t2 (cost=0.00..0.31 rows=1 width=8)
      Index Cond: (a = t1.a)
(4 rows)
```

NOTICE

- **predpush_same_level** can specify multiple **src** parameters in the same condition.
- If the specified **src** and **dest** conditions do not exist or do not meet the parameterized path requirements, this hint does not take effect.

6.9.13 Hint for Materializing a Sub-plan Result

Description

You can materialize a sub-plan result to temporarily store the query record. This hint is used only in the INSERT statement.

When the INSERT INTO ... SELECT statement is used to insert a large amount of data with multiple duplicate rows, the index needs to be compared for multiple times. As a result, the execution takes a long time. This HINT is used to materialize the results of subplans and temporarily store query records to reduce the number of index comparisons and shorten the statement execution time.

Syntax

```
material_subplan
```

Examples

Create a table and insert data into the table.

```
create table test(a int, b int) with(storage_type = ustore);
create index on test(a);
create table test_src(a int, b int);
insert into test_src values(generate_series(1,10), generate_series(1,100000));
```

Normal INSERT INTO...SELECT statement:

```
insert into test select /*+ nestloop(test_src t1)*/ * from test_src where not exists(select 1 from test t1 where
t1.a = test_src.a);
```

Execute plan:

```
QUERY PLAN
-----
Insert on test
-> Nested Loop Anti Join
   -> Seq Scan on test_src
   -> Index Only Scan using test_a_idx on test t1
       Index Cond: (a = test_src.a)
(5 rows)
```

Use the material_subplan hint operator:

```
insert /*+ material_subplan*/ into test select /*+ nestloop(test_src t1)*/ * from test_src where not
exists(select 1 from test t1 where t1.a = test_src.a);
```

The execution plan is as follows:

```
QUERY PLAN
-----
Insert on test
-> Materialize
   -> Nested Loop Anti Join
       -> Seq Scan on test_src
       -> Index Only Scan using test_a_idx on test t1
           Index Cond: (a = test_src.a)
(6 rows)
```

6.10 Tuning with SQL PATCH

SQL PATCH is designed for database administrators (DBAs), O&M personnel, and other roles who need to optimize SQL statements. If performance problems caused by poor plans of service statements are identified through other O&M views or fault locating methods, you can create an SQL patch to optimize service statements based on hints. Currently, the following hints are supported: number of rows, scanning mode, join mode, join sequence, PBE custom/generic plan selection, statement-level parameter setting, and parameterized path. In addition, in case that services are unavailable due to internal system errors that are triggered by specific statements, you can create SQL patches to rectify single-point failures without changing service statements. In this way, errors can be reported in advance to avoid greater loss.

Constraints

1. Patches can be created only by unique SQL ID. If unique SQL IDs conflict, SQL patches that are used for hint-based optimization may affect performance but do not affect semantic correctness.

2. Only hints that do not change SQL semantics can be used as patches. SQL rewriting is not supported.
3. This tool is not applicable to logical backup and restoration.
4. The patch validity cannot be verified during patch creation. If the patch hint has syntax or semantic errors, the query execution is not affected.
5. Only the initial user, O&M administrator, monitoring administrator, and system administrator have the permission to perform this operation.
6. Patches are not shared between databases. When creating SQL patches, you need to connect to the target database.
7. In the centralized deployment scenario where the standby node is readable, you must specify the primary node to run the SQL PATCH command to create, modify, or delete functions and the standby node to report errors.
8. There is a delay in synchronizing an SQL patch to the standby node. The patch takes effect after the standby node replays related logs.
9. This function does not take effect for SQL statements in stored procedures because no unique SQL ID is generated for statements in stored procedures.
10. It is not recommended that the abort patch be used in the database for a long time. It should be used only as a workaround. If the database service is unavailable due to a kernel fault triggered by a specific statement, you must rectify the service fault or upgrade the kernel as soon as possible. After the upgrade, the method of generating unique SQL IDs may change. Therefore, the workaround may become invalid.
11. Currently, except DML statements, unique SQL IDs of SQL statements (such as CREATE TABLE) are generated by hashing the statement text. Therefore, SQL PATCH is sensitive to uppercase and lowercase letters, spaces, and line breaks. That is, even statements of different texts are semantically relative, you still need to create different SQL patches for them. For DML operations, SQL PATCH can take effect for the same statement with different input parameters, regardless of uppercase letters, lowercase letters, and spaces.

Examples

SQL PATCH is implemented based on the unique SQL ID. Therefore, to use SQL PATCH, related O&M parameters must be enabled for the SQL patch to take effect. The unique SQL ID can be obtained from both the WDR and slow SQL view. You must specify the unique SQL ID when creating an SQL patch. The following provides a simple example.

The following provides two simple examples.

Scenario 1: Use SQL PATCH to optimize specific statements based on hints.

```
openGauss=# create table hint_t1(a int);
CREATE TABLE
openGauss=# create index on hint_t1(a);
CREATE INDEX
openGauss=# set track_stmt_stat_level = 'L1,L1'; -- Enable full SQL statistics.
SET
openGauss=# select * from hint_t1 t1 where t1.a = 1; -- Execute the SQL statement.
a | b | c
-----+----
1 | 1 | 1
(1 row)
openGauss=# select unique_query_id, query, query_plan from db_perf.statement_history where query like
```



```
%hint_t1%'; -- Obtain the query plan and unique SQL ID.
-[ RECORD 1 ]-----+-----
unique_query_id | 2578396627
query           | select * from hint_t1 t1 where t1.a = ?;
query_plan      | Datanode Name: sgnode
                | Bitmap Heap Scan on hint_t1 t1 (cost=4.33..15.70 rows=10 p-time=0 p-rows=0 width=12)
                |   Recheck Cond: (a = '****')
                |   -> Bitmap Index Scan on hint_t1_a_idx (cost=0.00..4.33 rows=10 p-time=0 p-rows=0 width=0)
                |       Index Cond: (a = '****')
                |
openGauss=# select * from db_sql_util.create_hint_sql_patch('patch1', 2578396627, 'indexscan(t1)'); --
Specify a hint patch for the specified unique SQL ID.
-[ RECORD 1 ]-----+---
create_hint_sql_patch | t
openGauss=# explain select * from hint_t1 t1 where t1.a = 1; -- Check whether the hint takes effect.
NOTICE: Plan influenced by SQL hint patch
               QUERY PLAN
-----
[Bypass]
Index Scan using hint_t1_a_idx on hint_t1 t1 (cost=0.00..32.43 rows=10 width=12)
  Index Cond: (a = 1)
(3 rows)
openGauss=# select * from hint_t1 t1 where t1.a = 1; -- Execute the statement again.
 a | b | c
---+---+---
 1 | 1 | 1
(1 row)
openGauss=# select unique_query_id, query, query_plan from db_perf.statement_history where query like
'%hint_t1%'; -- The query plan has been changed.
-[ RECORD 1 ]-----+-----
unique_query_id | 2578396627
query           | select * from hint_t1 t1 where t1.a = ?;
query_plan      | Datanode Name: sgnode
                | Bitmap Heap Scan on hint_t1 t1 (cost=4.33..15.70 rows=10 p-time=0 p-rows=0 width=12)
                |   Recheck Cond: (a = '****')
                |   -> Bitmap Index Scan on hint_t1_a_idx (cost=0.00..4.33 rows=10 p-time=0 p-rows=0 width=0)
                |       Index Cond: (a = '****')
                |
-[ RECORD 2 ]-----+-----
unique_query_id | 2578396627
query           | select * from hint_t1 t1 where t1.a = ?;
query_plan      | Datanode Name: sgnode
                | Index Scan using hint_t1_a_idx on hint_t1 t1 (cost=0.00..8.27 rows=1 p-time=0 p-rows=0
width=12)
                |   Index Cond: (a = '****')
                |
```

Scenario 2: Run the SQL PATCH command to report an error for a specific statement in advance.

```
openGauss=# select * from db_sql_util.drop_sql_patch('patch1'); -- Delete patch 1.
drop_sql_patch
-----
t
(1 row)
openGauss=# select * from db_sql_util.create_abort_sql_patch('patch2', 2578396627); -- Create an abort
patch for the unique SQL ID of the statement.
create_abort_sql_patch
-----
t
(1 row)

openGauss=# select * from hint_t1 t1 where t1.a = 1; -- An error is reported in advance when the statement
is executed again.
ERROR: Statement 2578396627 canceled by abort patch patch2
```

Helpful Links

The following table lists the system catalogs and interface functions related to SQL PATCH.

Table 6-3 System catalogs and interface functions related to SQL PATCH

Name		Description
System catalog	GS_SQL_PATCH	GS_SQL_PATCH records the status information about all SQL patches.
Interface functions DBE_SQL_UTIL Schema	DBE_SQL_UTIL.create_hint_sql_patch	create_hint_sql_patch is an interface function used to create SQL patches for hints. It returns whether the execution is successful.
	DBE_SQL_UTIL.create_abort_sql_patch	create_abort_sql_patch is an interface function used to create abort SQL patches. It returns whether the execution is successful.
	DBE_SQL_UTIL.drop_sql_patch	drop_sql_patch deletes SQL patches from the connected CN and returns whether the execution is successful.
	DBE_SQL_UTIL.enable_sql_patch	enable_sql_patch enables SQL patches on the connected CN and returns whether the execution is successful.
	DBE_SQL_UTIL.disable_sql_patch	disable_sql_patch disables SQL patches on the connected CN and returns whether the execution is successful.
	DBE_SQL_UTIL.show_sql_patch	show_sql_patch displays the SQL patch corresponding to a specified patch name and returns the running result.

6.11 Optimization Cases

6.11.1 Case: Modifying the GUC Parameter `rewrite_rule`

`rewrite_rule` contains multiple query rewriting rules: **magicset**, **partialpush**, **uniquecheck**, **disablerep**, **intargetlist**, and **predpush**. The following describes the application scenarios of some important rules:

Promoting the Subquery in the Target Column Using `intargetlist`

The query performance can be greatly improved by converting the subquery in the target column to JOIN. The following is an example:

```
openGauss=# set rewrite_rule='none';
SET
```

```
openGauss=# create table t1(c1 int,c2 int);
CREATE TABLE
openGauss=# create table t2(c1 int,c2 int);
CREATE TABLE
openGauss=# explain (verbose on, costs off) select c1,(select avg(c2) from t2 where t2.c2=t1.c2) from t1
where t1.c1<100 order by t1.c2;
QUERY PLAN
-----
Sort
  Output: t1.c1, ((SubPlan 1)), t1.c2
  Sort Key: t1.c2
  -> Seq Scan on public.t1
    Output: t1.c1, (SubPlan 1), t1.c2
    Filter: (t1.c1 < 100)
    SubPlan 1
      -> Aggregate
        Output: avg(t2.c2)
        -> Seq Scan on public.t2
          Output: t2.c1, t2.c2
          Filter: (t2.c2 = t1.c2)
(12 rows)
```

Because the subquery (**select avg(c2) from t2 where t2.c2=t1.c2**) in the target column cannot be pulled up, execution of the subquery is triggered each time a row of data of **t1** is scanned, and the query efficiency is low. If the **intargetlist** parameter is enabled, the subquery is converted to JOIN to improve the query performance.

```
openGauss=# set rewrite_rule='intargetlist';
SET
openGauss=# explain (verbose on, costs off) select c1,(select avg(c2) from t2 where t2.c2=t1.c2) from t1
where t1.c1<100 order by t1.c2;
QUERY PLAN
-----
Sort
  Output: t1.c1, (avg(t2.c2)), t1.c2
  Sort Key: t1.c2
  -> Hash Left Join
    Output: t1.c1, (avg(t2.c2)), t1.c2
    Hash Cond: (t1.c2 = t2.c2)
    -> Seq Scan on public.t1
      Output: t1.c1, t1.c2
      Filter: (t1.c1 < 100)
    -> Hash
      Output: (avg(t2.c2)), t2.c2
      -> HashAggregate
        Output: avg(t2.c2), t2.c2
        Group By Key: t2.c2
        -> Seq Scan on public.t2
          Output: t2.c2
(16 rows)
```

Promoting the Subquery Without Aggregate Using uniquecheck

Ensure that each condition has only one line of output. The subqueries with aggregate functions can be automatically pulled up. For subqueries without aggregate functions, the following is an example:

```
select t1.c1 from t1 where t1.c1 = (select t2.c1 from t2 where t1.c1=t2.c2) ;
```

Rewrite as follows:

```
select t1.c1 from t1 join (select t2.c1 from t2 where t2.c1 is not null group by
t2.c1(unique check)) tt(c1) on tt.c1=t1.c1;
```

To ensure semantic equivalence, the subquery **tt** must ensure that each **group by t2.c1** has only one line of output. Enable the **uniquecheck** query rewriting

parameter to ensure that the query can be pulled up and equivalent. If more than one row of data is output at run time, an error is reported.

```
openGauss=# set rewrite_rule='uniquecheck';
SET
openGauss=# explain verbose select t1.c1 from t1 where t1.c1 = (select t2.c1 from t2 where t1.c1=t2.c1);
QUERY PLAN
-----
Hash Join (cost=43.36..104.40 rows=2149 distinct=[200, 200] width=4)
  Output: t1.c1
  Hash Cond: (t1.c1 = subquery."?column?")
  -> Seq Scan on public.t1 (cost=0.00..31.49 rows=2149 width=4)
    Output: t1.c1, t1.c2
  -> Hash (cost=40.86..40.86 rows=200 width=8)
    Output: subquery."?column?", subquery.c1
    -> Subquery Scan on subquery (cost=36.86..40.86 rows=200 width=8)
      Output: subquery."?column?", subquery.c1
      -> HashAggregate (cost=36.86..38.86 rows=200 width=4)
        Output: t2.c1, t2.c1
        Group By Key: t2.c1
        Filter: (t2.c1 IS NOT NULL)
        Unique Check Required
        -> Seq Scan on public.t2 (cost=0.00..31.49 rows=2149 width=4)
          Output: t2.c1
(16 rows)
```

Note: Because **group by t2.c1 unique check** occurs before the filter condition **t1.c1=t2.c1**, an error may be reported after the query that does not report an error is rewritten. An example is as follows:

There are tables **t1** and **t2**. The data in the tables is as follows:

```
openGauss=# select * from t1 order by c2;
 c1 | c2
----+----
  1 |  1
  2 |  2
  3 |  3
(3 rows)
openGauss=# select * from t2 order by c2;
 c1 | c2
----+----
  1 |  1
  2 |  2
  3 |  3
  4 |  4
  4 |  4
  5 |  5
(6 rows)
```

Disable and enable the **uniquecheck** parameter for comparison. After the parameter is enabled, an error is reported.

```
openGauss=# select t1.c1 from t1 where t1.c1 = (select t2.c1 from t2 where t1.c1=t2.c2) ;
 c1
----
  1
  2
  3
(3 rows)
openGauss=# set rewrite_rule='uniquecheck';
SET
openGauss=# select t1.c1 from t1 where t1.c1 = (select t2.c1 from t2 where t1.c1=t2.c2) ;
ERROR: more than one row returned by a subquery used as an expression
```

6.11.2 Case: Creating an Appropriate Index

Symptom

Query the information about all personnel in the sales department.

```
-- Create a table.
CREATE TABLE staffs (staff_id NUMBER(6) NOT NULL, first_name VARCHAR2(20), last_name
VARCHAR2(25), employment_id VARCHAR2(10), section_id NUMBER(4), state_name VARCHAR2(10), city
VARCHAR2(10));
CREATE TABLE sections(section_id NUMBER(4), place_id NUMBER(4), section_name VARCHAR2(20));
CREATE TABLE states(state_id NUMBER(4));
CREATE TABLE places(place_id NUMBER(4), state_id NUMBER(4));
-- Query before optimization.
EXPLAIN SELECT staff_id,first_name,last_name,employment_id,state_name,city
FROM staffs,sections,states,places
WHERE sections.section_name='Sales'
AND staffs.section_id = sections.section_id
AND sections.place_id = places.place_id
AND places.state_id = states.state_id
ORDER BY staff_id;
-- Create indexes.
CREATE INDEX loc_id_pk ON places(place_id);
CREATE INDEX state_c_id_pk ON states(state_id);
-- Query after optimization.
EXPLAIN SELECT staff_id,first_name,last_name,employment_id,state_name,city
FROM staffs,sections,states,places
WHERE sections.section_name='Sales'
AND staffs.section_id = sections.section_id
AND sections.place_id = places.place_id
AND places.state_id = states.state_id
ORDER BY staff_id;
```

Optimization Analysis

The original execution plan is as follows before creating the **places.place_id** and **states.state_id** indexes:

```
QUERY PLAN
-----
Sort (cost=125.25..126.34 rows=438 width=254)
  Sort Key: staffs.staff_id
  -> Hash Join (cost=64.91..106.03 rows=438 width=254)
    Hash Cond: (states.state_id = places.state_id)
    -> Seq Scan on states (cost=0.00..29.45 rows=1945 width=12)
    -> Hash (cost=64.35..64.35 rows=45 width=266)
      -> Hash Join (cost=33.09..64.35 rows=45 width=266)
        Hash Cond: (places.place_id = sections.place_id)
        -> Seq Scan on places (cost=0.00..25.13 rows=1513 width=24)
        -> Hash (cost=33.02..33.02 rows=6 width=266)
          -> Hash Join (cost=19.16..33.02 rows=6 width=266)
            Hash Cond: (staffs.section_id = sections.section_id)
            -> Seq Scan on staffs (cost=0.00..12.76 rows=276 width=266)
            -> Hash (cost=19.11..19.11 rows=4 width=24)
              -> Seq Scan on sections (cost=0.00..19.11 rows=4 width=24)
              Filter: ((section_name)::text = 'Sales'::text)
  (16 rows)
```

The optimized execution plan is as follows (two indexes have been created on the **places.place_id** and **states.state_id** columns in [Symptom](#)):

```
QUERY PLAN
-----
Sort (cost=107.40..108.49 rows=438 width=254)
  Sort Key: staffs.staff_id
  -> Hash Join (cost=35.37..88.18 rows=438 width=254)
```

```
Hash Cond: (sections.section_id = staffs.section_id)
-> Nested Loop (cost=19.16..66.85 rows=292 width=12)
    -> Hash Join (cost=19.16..50.27 rows=30 width=24)
        Hash Cond: (places.place_id = sections.place_id)
        -> Seq Scan on places (cost=0.00..25.13 rows=1513 width=24)
        -> Hash (cost=19.11..19.11 rows=4 width=24)
            -> Seq Scan on sections (cost=0.00..19.11 rows=4 width=24)
                Filter: ((section_name)::text = 'Sales'::text)
            -> Index Only Scan using state_c_id_pk on states (cost=0.00..0.45 rows=10 width=12)
                Index Cond: (state_id = places.state_id)
        -> Hash (cost=12.76..12.76 rows=276 width=266)
            -> Seq Scan on staffs (cost=0.00..12.76 rows=276 width=266)
(15 rows)
```

6.11.3 Case: Adding NOT NULL for the JOIN Columns

```
SELECT * FROM join_a a JOIN join_b b ON a.b = b.b;
```

The execution plan is as follows:

```
QUERY PLAN
-----
Hash Join (cost=58.35..14677.69 rows=1074607 width=16) (actual time=23.374..23.384 rows=10 loops=1)
  Hash Cond: (a.b = b.b)
    -> Seq Scan on join_a a (cost=0.00..2248.10 rows=100010 width=8) (actual time=0.495..12.551 rows=100010 loops=1)
    -> Hash (cost=31.49..31.49 rows=2149 width=8) (actual time=0.614..0.614 rows=1000 loops=1)
        Buckets: 32768 Batches: 1 Memory Usage: 40kB
        -> Seq Scan on join_b b (cost=0.00..31.49 rows=2149 width=8) (actual time=0.009..0.183 rows=1000 loops=1)
  Total runtime: 23.716 ms
(7 rows)
```

Optimization Analysis

1. According to the execution plan, the sequential scan phase is time consuming.
2. Therefore, you are advised to manually add **NOT NULL** for the **JOIN** column in the statement, as shown below:

```
SELECT
*
SELECT * FROM join_a a JOIN join_b b ON a.b = b.b where a.b IS NOT NULL;
```

The execution plan is as follows:

```
QUERY PLAN
-----
Hash Join (cost=58.22..14560.97 rows=1063762 width=16) (actual time=13.237..13.247 rows=10 loops=1)
  Hash Cond: (a.b = b.b)
    -> Seq Scan on join_a a (cost=0.00..2248.10 rows=99510 width=8) (actual time=12.417..12.422 rows=10 loops=1)
        Filter: (b IS NOT NULL)
        Rows Removed by Filter: 100000
    -> Hash (cost=31.49..31.49 rows=2138 width=8) (actual time=0.566..0.566 rows=1000 loops=1)
        Buckets: 32768 Batches: 1 Memory Usage: 40kB
        -> Seq Scan on join_b b (cost=0.00..31.49 rows=2138 width=8) (actual time=0.011..0.229 rows=1000 loops=1)
            Filter: (b IS NOT NULL)
  Total runtime: 13.556 ms
(10 rows)
```

6.11.4 Case: Modifying a Partitioned Table

Symptom

In the following simple SQL statements, the performance bottlenecks exist in the scan operation on the **normal_date** table.

```
-----
QUERY PLAN
-----
Seq Scan on normal_date (cost=0.00..259.00 rows=30 width=12) (actual time=0.100..3.466 rows=30 loops=1)
  Filter: (("time" >= '2022-09-01 00:00:00'::timestamp without time zone) AND ("time" <= '2022-10-01 00:00:00'::timestamp without time zone))
  Rows Removed by Filter: 9970
  Total runtime: 3.587 ms
(4 rows)
```

Optimization Analysis

Obviously, there are date features in the **time** column of table data in the service layer, and this meet the features of a partitioned table. Replan the table definition of the **normal_date** table. Set the **time** column as a partition key, and month as an interval unit. Define the partitioned table **normal_date_part**. The modified result is as follows, and the performance is improved by nearly 10 times.

```
-----
QUERY PLAN
-----
Partition Iterator (cost=0.00..480.00 rows=30 width=12) (actual time=0.038..0.085 rows=30 loops=1)
  Iterations: 2
  -> Partitioned Seq Scan on normal_date_part (cost=0.00..480.00 rows=30 width=12) (actual time=0.049..0.063 rows=30 loops=2)
    Filter: (("time" >= '2022-09-01 00:00:00'::timestamp without time zone) AND ("time" <= '2022-10-01 00:00:00'::timestamp without time zone))
    Rows Removed by Filter: 31
    Selected Partitions: 3..4
    Total runtime: 0.360 ms
(7 rows)
```

6.11.5 Case: Rewriting SQL Statements to Eliminate Subqueries

Symptom

```
select
  1,
  (select count(*) from normal_date n where n.id = a.id) as GZCS
from normal_date a;
```

This SQL performance is poor. SubPlan exists in the execution plan as follows:

```
-----
QUERY PLAN
-----
Seq Scan on normal_date a (cost=0.00..888118.42 rows=5129 width=4) (actual time=2.394..22194.907 rows=10000 loops=1)
  SubPlan 1
    -> Aggregate (cost=173.12..173.12 rows=1 width=8) (actual time=22179.496..22179.942 rows=10000 loops=10000)
      -> Seq Scan on normal_date n (cost=0.00..173.11 rows=1 width=0) (actual time=11279.349..22159.608 rows=10000 loops=10000)
```

```
Filter: (id = a.id)
Rows Removed by Filter: 99990000
Total runtime: 22196.415 ms
(7 rows)
```

Optimization

The core of this optimization is to eliminate subqueries. Based on the service scenario analysis, *a.id* is not null. In terms of SQL syntax, you can rewrite the SQL statement as follows:

```
select
count(*)
from normal_date n, normal_date a
where n.id = a.id
group by a.id;
```

The plan is as follows:

```
QUERY PLAN
-----
HashAggregate (cost=480.86..532.15 rows=5129 width=12) (actual time=21.539..24.356 rows=10000
loops=1)
  Group By Key: a.id
    -> Hash Join (cost=224.40..455.22 rows=5129 width=4) (actual time=6.402..13.484 rows=10000 loops=1)
        Hash Cond: (n.id = a.id)
        -> Seq Scan on normal_date n (cost=0.00..160.29 rows=5129 width=4) (actual time=0.087..1.459
rows=10000 loops=1)
            -> Hash (cost=160.29..160.29 rows=5129 width=4) (actual time=6.065..6.065 rows=10000 loops=1)
                Buckets: 32768 Batches: 1 Memory Usage: 352kB
            -> Seq Scan on normal_date a (cost=0.00..160.29 rows=5129 width=4) (actual time=0.046..2.738
rows=10000 loops=1)
        Total runtime: 26.844 ms
    (9 rows)
```

NOTE

To ensure that the modified statements have the same functions, **NOT NULL** is added to *normal_date.id*.

6.11.6 Case: Rewriting SQL Statements and Deleting in-clause

Symptom

in-clause/any-clause is a common SQL statement constraint. Sometimes, the clause following **in** or **any** is a constant. For example:

```
select
count(1)
from calc_empfyc_c1_result_tmp_t1
where ls_pid_cusr1 in ('20120405', '20130405');
```

Or

```
select
count(1)
from calc_empfyc_c1_result_tmp_t1
where ls_pid_cusr1 in any('20120405', '20130405');
```

Sometimes, the **in** or **any** clause is used as follows:

```
SELECT
*
FROM test1 t1, test2 t2
WHERE t1.a = any(values(t2.a),(t2.b));
```


a and **b** are two columns in **t2**, and "**t1.a = any(values(t2.ba,(t2.b)))**" is equivalent to "**t1.a = t2.a or t1.a = t2.b**".

Therefore, join-condition is essentially an inequality, and nestloop must be used for this unequal join operation. The corresponding execution plan is as follows:

```
-----
QUERY PLAN
-----
NestLoop (cost=0.00..138614.38 rows=2309100 width=16) (actual time=0.152..19225.483 rows=1000 loops=1)
  Join Filter: (SubPlan 1)
  Rows Removed by Join Filter: 999000
  -> Seq Scan on test1 t1 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.021..3.309 rows=1000 loops=1)
  -> Materialize (cost=0.00..42.23 rows=2149 width=8) (actual time=0.331..1265.810 rows=1000000 loops=1000)
    -> Seq Scan on test2 t2 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.013..0.268 rows=1000 loops=1)
    SubPlan 1
      -> Values Scan on "*VALUES*" (cost=0.00..0.03 rows=2 width=4) (actual time=2890.741..7372.739 rows=1999000 loops=1000000)
  Total runtime: 19227.328 ms
(9 rows)
```

Optimization

The test result shows that both result sets are too large. As a result, nestloop is time-consuming with more than one hour to return results. Therefore, the key to performance optimization is to eliminate nestloop, using more efficient hash join. From the perspective of semantic equivalence, the SQL statements can be written as follows:

```
SELECT
*
FROM (
  SELECT * FROM test1 t1, test2 t2 WHERE t1.a = t2.a
  UNION
  SELECT * FROM test1 t1, test2 t2 WHERE t1.a = t2.b
);
```

The optimized SQL query consists of two equivalent join subqueries, and each subquery can be used for hash join in this scenario. The optimized execution plan is as follows:

```
-----
QUERY PLAN
-----
HashAggregate (cost=1634.99..2096.81 rows=46182 width=16) (actual time=6.369..6.772 rows=1000 loops=1)
  Group By Key: t1.a, t1.b, t2.a, t2.b
  -> Append (cost=58.35..1173.17 rows=46182 width=16) (actual time=0.833..3.414 rows=2000 loops=1)
    -> Hash Join (cost=58.35..355.67 rows=23091 width=16) (actual time=0.832..1.590 rows=1000 loops=1)
      Hash Cond: (t1.a = t2.a)
      -> Seq Scan on test1 t1 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.015..0.156 rows=1000 loops=1)
      -> Hash (cost=31.49..31.49 rows=2149 width=8) (actual time=0.531..0.531 rows=1000 loops=1)
        Buckets: 32768 Batches: 1 Memory Usage: 40kB
        -> Seq Scan on test2 t2 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.010..0.199 rows=1000 loops=1)
    -> Hash Join (cost=58.35..355.67 rows=23091 width=16) (actual time=0.694..1.421 rows=1000 loops=1)
      Hash Cond: (t1.a = t2.b)
      -> Seq Scan on test1 t1 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.010..0.160 rows=1000 loops=1)
```

```
-> Hash (cost=31.49..31.49 rows=2149 width=8) (actual time=0.524..0.524 rows=1000 loops=1)
      Buckets: 32768 Batches: 1 Memory Usage: 40kB
      -> Seq Scan on test2 t2 (cost=0.00..31.49 rows=2149 width=8) (actual time=0.008..0.177
rows=1000 loops=1)
Total runtime: 7.759 ms
(16 rows)
```

7 SQL Reference

7.1 GaussDB SQL

What Is SQL?

SQL is a standard computer language used to control the access to databases and manage data in databases.

SQL provides different statements to enable you to:

- Query data.
- Insert, update, and delete rows.
- Create, replace, modify, and delete objects.
- Control the access to a database and its objects.
- Maintain the consistency and integrity of a database.

SQL consists of commands and functions that are used to manage databases and database objects. SQL can also forcibly implement the rules for data types, expressions, and texts. Therefore, [SQL Reference](#) describes data types, expressions, functions, and operators in addition to SQL syntax.

Development of SQL Standards

The development history of SQL standards is as follows:

- 1986: ANSI X3.135-1986, ISO/IEC 9075:1986, SQL-86
- 1989: ANSI X3.135-1989, ISO/IEC 9075:1989, SQL-89
- 1992: ANSI X3.135-1992, ISO/IEC 9075:1992, SQL-92 (SQL2)
- 1999: ISO/IEC 9075:1999, SQL:1999
- 2003: ISO/IEC 9075:2003, SQL:2003
- 2011: ISO/IEC 9075:2011, SQL:2011
- 2016: ISO/IEC 9075:2016, SQL:2016
- 2019: ISO/IEC 9075:2019, SQL:2019

SQL Standards Supported by GaussDB

By default, GaussDB supports most features of SQL:2016.

7.2 Keywords

SQL statements contain reserved keywords and non-reserved keywords. Standards require that reserved keywords not be used as other identifiers. Non-reserved keywords have special meanings only in a specific environment and can be used as identifiers in other environments.

NOTICE

1. Currently, the non-reserved keywords have the following restrictions when being used as the identifier of a database object:
 1. It cannot be directly used as a column alias. That is, usage similar to `SELECT 1 ABORT` may cause errors.
 2. Keywords `ENTITYESCAPING`, `NOENTITYESCAPING`, and `WELLFORMED` cannot be used as identifiers of table names, column names, table aliases, or column aliases, regardless of whether they are enclosed with double quotation marks. In addition, they cannot be used as function names without double quotation marks.
 3. The `RAW` keyword without double quotation marks cannot be used as the identifier of a table name or function name.
 4. The `SET` keyword without double quotation marks cannot be used as an identifier of a table alias. That is, usage similar to `SELECT * FROM T1 SET` may cause errors.
 5. Keywords such as `BEGIN`, `BY`, `CLOSE`, `CURSOR`, `DECLARE`, `DELETE`, `EXECUTE`, `FUNCTION`, `IF`, `IMMEDIATE`, `INSERT`, `LOOP`, `MOVE`, `OF`, `REF`, `RELEASE`, `RETURN`, `SAVEPOINT`, `STRICT`, `TYPE`, and `UPDATE` without double quotation marks cannot be used as variable names.
 6. When the `SYS_REFCURSOR` keyword is used as the identifier of a database object, if double quotation marks are not attached, a database object named **REFCURSOR** is created. If double quotation marks are attached, a database object named **SYS_REFCURSOR** is created.
 2. Similar to the non-reserved keywords, the non-reserved (cannot be a function or type) keywords cannot be directly used as column aliases, either.
 3. The reserved keyword `CURRENT_TIMESTAMP` with double quotation marks cannot be used as a function name.
-

Identifier Naming Conventions

Identifier naming must comply with the following rules:

- An identifier name can only contain letters, digits, underscores (`_`), and dollar signs (`$`).
- An identifier name must start with a letter or an underscore (`_`).

 NOTE

- The naming rules are recommended but not required.
- In special cases, double quotation marks (") can be used to avoid special character errors.

SQL Keywords

Table 7-1 SQL keywords

Keyword	GaussDB	SQL:1999	SQL-92
ABORT	Non-reserved	N/A	N/A
ABS	N/A	Non-reserved	N/A
ABSOLUTE	Non-reserved	Reserved	Reserved
ACCESS	Non-reserved	N/A	N/A
ACCOUNT	Non-reserved	N/A	N/A
ACTION	Non-reserved	Reserved	Reserved
ADA	N/A	Non-reserved	Non-reserved
ADD	Non-reserved	Reserved	Reserved
ADMIN	Non-reserved	Reserved	N/A
AFTER	Non-reserved	Reserved	N/A
AGGREGATE	Non-reserved	Reserved	N/A
ALGORITHM	Non-reserved	N/A	N/A
ALIAS	N/A	Reserved	N/A
ALL	Reserved	Reserved	Reserved
ALLOCATE	N/A	Reserved	Reserved
ALSO	Non-reserved	N/A	N/A
ALTER	Non-reserved	Reserved	Reserved
ALWAYS	Non-reserved	N/A	N/A
ANALYSE	Reserved	N/A	N/A
ANALYZE	Reserved	N/A	N/A
AND	Reserved	Reserved	Reserved
ANY	Reserved	Reserved	Reserved
APP	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
APPEND	Non-reserved	N/A	N/A
ARCHIVE	Non-reserved	N/A	N/A
ARE	N/A	Reserved	Reserved
ARRAY	Reserved	Reserved	N/A
AS	Reserved	Reserved	Reserved
ASC	Reserved	Reserved	Reserved
ASENSITIVE	N/A	Non-reserved	N/A
ASSERTION	Non-reserved	Reserved	Reserved
ASSIGNMENT	Non-reserved	Non-reserved	N/A
ASYMMETRIC	Reserved	Non-reserved	N/A
AT	Non-reserved	Reserved	Reserved
ATOMIC	N/A	Non-reserved	N/A
ATTRIBUTE	Non-reserved	N/A	N/A
AUDIT	Non-reserved	N/A	N/A
AUTHID	Reserved	N/A	N/A
AUTHORIZATION	Reserved (functions and types allowed)	Reserved	Reserved
AUTOEXTEND	Non-reserved	N/A	N/A
AUTOMAPPED	Non-reserved	N/A	N/A
AVG	N/A	Non-reserved	Reserved
BACKWARD	Non-reserved	N/A	N/A
BARRIER	Non-reserved	N/A	N/A
BEFORE	Non-reserved	Reserved	N/A
BEGIN	Non-reserved	Reserved	Reserved
BEGIN_NON_ANOYBLOCK	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
BETWEEN	Non-reserved (excluding functions and types)	Non-reserved	Reserved
BIGINT	Non-reserved (excluding functions and types)	N/A	N/A
BINARY	Reserved (functions and types allowed)	Reserved	N/A
BINARY_DOUBLE	Non-reserved (excluding functions and types)	N/A	N/A
BINARY_INTEGER	Non-reserved (excluding functions and types)	N/A	N/A
BIT	Non-reserved (excluding functions and types)	Reserved	Reserved
BIT_LENGTH	N/A	Non-reserved	Reserved
BITVAR	N/A	Non-reserved	N/A
BLANKS	Non-reserved	N/A	N/A
BLOB	Non-reserved	Reserved	N/A
BLOCKCHAIN	Non-reserved	N/A	N/A
BODY	Non-reserved	N/A	N/A
BOOLEAN	Non-reserved (excluding functions and types)	Reserved	N/A
BOTH	Reserved	Reserved	Reserved
BREADTH	N/A	Reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
BUCKETCNT	Non-reserved (excluding functions and types)	N/A	N/A
BUCKETS	Reserved	N/A	N/A
BY	Non-reserved	Reserved	Reserved
BYTEAWITHOUTORDER	Non-reserved (excluding functions and types)	N/A	N/A
BYTEAWITHOUTORDER- WITHEQUAL	Non-reserved (excluding functions and types)	N/A	N/A
C	N/A	Non- reserved	Non-reserved
CACHE	Non-reserved	N/A	N/A
CALL	Non-reserved	Reserved	N/A
CALLED	Non-reserved	Non- reserved	N/A
CANCELABLE	Non-reserved	N/A	N/A
CARDINALITY	N/A	Non- reserved	N/A
CASCADE	Non-reserved	Reserved	Reserved
CASCADEED	Non-reserved	Reserved	Reserved
CASE	Reserved	Reserved	Reserved
CAST	Reserved	Reserved	Reserved
CATALOG	Non-reserved	Reserved	Reserved
CATALOG_NAME	N/A	Non- reserved	Non-reserved
CHAIN	Non-reserved	Non- reserved	N/A
CHAR	Non-reserved (excluding functions and types)	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
CHAR_LENGTH	N/A	Non-reserved	Reserved
CHARACTER	Non-reserved (excluding functions and types)	Reserved	Reserved
CHARACTER_LENGTH	N/A	Non-reserved	Reserved
CHARACTER_SET_CATALOG	N/A	Non-reserved	Non-reserved
CHARACTER_SET_NAME	N/A	Non-reserved	Non-reserved
CHARACTER_SET_SCHEMA	N/A	Non-reserved	Non-reserved
CHARACTERISTICS	Non-reserved	N/A	N/A
CHARACTERSET	Non-reserved	N/A	N/A
CHECK	Reserved	Reserved	Reserved
CHECKED	N/A	Non-reserved	N/A
CHECKPOINT	Non-reserved	N/A	N/A
CLASS	Non-reserved	Reserved	N/A
CLASS_ORIGIN	N/A	Non-reserved	Non-reserved
CLEAN	Non-reserved	N/A	N/A
CLIENT	Non-reserved	N/A	N/A
CLIENT_MASTER_KEY	Non-reserved	N/A	N/A
CLIENT_MASTER_KEYS	Non-reserved	N/A	N/A
CLOB	Non-reserved	Reserved	N/A
CLOSE	Non-reserved	Reserved	Reserved
CLUSTER	Non-reserved	N/A	N/A
COALESCE	Non-reserved (excluding functions and types)	Non-reserved	Reserved
COBOL	N/A	Non-reserved	Non-reserved

Keyword	GaussDB	SQL:1999	SQL-92
COLLATE	Reserved	Reserved	Reserved
COLLATION	Reserved (functions and types allowed)	Reserved	Reserved
COLLATION_CATALOG	N/A	Non-reserved	Non-reserved
COLLATION_NAME	N/A	Non-reserved	Non-reserved
COLLATION_SCHEMA	N/A	Non-reserved	Non-reserved
COLUMN	Reserved	Reserved	Reserved
COLUMN_ENCRYPTION_KEY	Non-reserved	N/A	N/A
COLUMN_ENCRYPTION_KEYS	Non-reserved	N/A	N/A
COLUMN_NAME	N/A	Non-reserved	Non-reserved
COMMAND_FUNCTION	N/A	Non-reserved	Non-reserved
COMMAND_FUNCTION_CODE	N/A	Non-reserved	N/A
COMMENT	Non-reserved	N/A	N/A
COMMENTS	Non-reserved	N/A	N/A
COMMIT	Non-reserved	Reserved	Reserved
COMMITTED	Non-reserved	Non-reserved	Non-reserved
COMPACT	Reserved (functions and types allowed)	N/A	N/A
COMPATIBLE_ILLEGAL_CHARS	Non-reserved	N/A	N/A
COMPLETE	Non-reserved	N/A	N/A
COMPLETION	N/A	Reserved	N/A
CONCURRENTLY	Reserved (functions and types allowed)	N/A	N/A
CONDITION	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
CONDITION_NUMBER	N/A	Non-reserved	Non-reserved
CONFIGURATION	Non-reserved	N/A	N/A
CONNECT	Non-reserved	Reserved	Reserved
CONNECTION	Non-reserved	Reserved	Reserved
CONNECTION_NAME	N/A	Non-reserved	Non-reserved
CONSTANT	Non-reserved	N/A	N/A
CONSTRAINT	Reserved	Reserved	Reserved
CONSTRAINT_CATALOG	N/A	Non-reserved	Non-reserved
CONSTRAINT_NAME	N/A	Non-reserved	Non-reserved
CONSTRAINT_SCHEMA	N/A	Non-reserved	Non-reserved
CONSTRAINTS	Non-reserved	Reserved	Reserved
CONSTRUCTOR	N/A	Reserved	N/A
CONTAINS	N/A	Non-reserved	N/A
CONTENT	Non-reserved	N/A	N/A
CONTINUE	Non-reserved	Reserved	Reserved
CONVIEW	Non-reserved	N/A	N/A
CONVERSION	Non-reserved	N/A	N/A
CONVERT	N/A	Non-reserved	Reserved
COORDINATOR	Non-reserved	N/A	N/A
COORDINATORS	Non-reserved	N/A	N/A
COPY	Non-reserved	N/A	N/A
CORRESPONDING	N/A	Reserved	Reserved
COST	Non-reserved	N/A	N/A
COUNT	N/A	Non-reserved	Reserved
CREATE	Reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
CROSS	Reserved (functions and types allowed)	Reserved	Reserved
CSN	Reserved (functions and types allowed)	N/A	N/A
CSV	Non-reserved	N/A	N/A
CUBE	Non-reserved	Reserved	N/A
CURRENT	Non-reserved	Reserved	Reserved
CURRENT_CATALOG	Reserved	N/A	N/A
CURRENT_DATE	Reserved	Reserved	Reserved
CURRENT_PATH	N/A	Reserved	N/A
CURRENT_ROLE	Reserved	Reserved	N/A
CURRENT_SCHEMA	Reserved (functions and types allowed)	N/A	N/A
CURRENT_TIME	Reserved	Reserved	Reserved
CURRENT_TIMESTAMP	Reserved	Reserved	Reserved
CURRENT_USER	Reserved	Reserved	Reserved
CURSOR	Non-reserved	Reserved	Reserved
CURSOR_NAME	N/A	Non-reserved	Non-reserved
CYCLE	Non-reserved	Reserved	N/A
DATA	Non-reserved	Reserved	Non-reserved
DATABASE	Non-reserved	N/A	N/A
DATAFILE	Non-reserved	N/A	N/A
DATANODE	Non-reserved	N/A	N/A
DATANODES	Non-reserved	N/A	N/A
DATATYPE_CL	Non-reserved	N/A	N/A
DATE	Non-reserved (excluding functions and types)	Reserved	Reserved
DATE_FORMAT	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
DATETIME_INTERVAL_CODE	N/A	Non-reserved	Non-reserved
DATETIME_INTERVAL_PRECISION	N/A	Non-reserved	Non-reserved
DAY	Non-reserved	Reserved	Reserved
DBCOMPATIBILITY	Non-reserved	N/A	N/A
DEALLOCATE	Non-reserved	Reserved	Reserved
DEC	Non-reserved (excluding functions and types)	Reserved	Reserved
DECIMAL	Non-reserved (excluding functions and types)	Reserved	Reserved
DECLARE	Non-reserved	Reserved	Reserved
DECODE	Non-reserved (excluding functions and types)	N/A	N/A
DEFAULT	Reserved	Reserved	Reserved
DEFAULTS	Non-reserved	N/A	N/A
DEFERRABLE	Reserved	Reserved	Reserved
DEFERRED	Non-reserved	Reserved	Reserved
DEFINED	N/A	Non-reserved	N/A
DEFINER	Non-reserved	Non-reserved	N/A
DELETE	Non-reserved	Reserved	Reserved
DELIMITER	Non-reserved	N/A	N/A
DELIMITERS	Non-reserved	N/A	N/A
DELTA	Non-reserved	N/A	N/A
DELTAMERGE	Reserved (functions and types allowed)	N/A	N/A
DEPTH	N/A	Reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
DEREF	N/A	Reserved	N/A
DESC	Reserved	Reserved	Reserved
DESCRIBE	N/A	Reserved	Reserved
DESCRIPTOR	N/A	Reserved	Reserved
DESTROY	N/A	Reserved	N/A
DESTRUCTOR	N/A	Reserved	N/A
DETERMINISTIC	Non-reserved	Reserved	N/A
DIAGNOSTICS	N/A	Reserved	Reserved
DICTIONARY	Non-reserved	Reserved	N/A
DIRECT	Non-reserved	N/A	N/A
DIRECTORY	Non-reserved	N/A	N/A
DISABLE	Non-reserved	N/A	N/A
DISCARD	Non-reserved	N/A	N/A
DISCONNECT	Non-reserved	Reserved	Reserved
DISPATCH	N/A	Non-reserved	N/A
DISTINCT	Reserved	Reserved	Reserved
DISTRIBUTE	Non-reserved	N/A	N/A
DISTRIBUTION	Non-reserved	N/A	N/A
DO	Reserved	N/A	N/A
DOCUMENT	Non-reserved	N/A	N/A
DOMAIN	Non-reserved	Reserved	Reserved
DOUBLE	Non-reserved	Reserved	Reserved
DROP	Non-reserved	Reserved	Reserved
DUPLICATE	Non-reserved	N/A	N/A
DYNAMIC	N/A	Reserved	N/A
DYNAMIC_FUNCTION	N/A	Non-reserved	Non-reserved
DYNAMIC_FUNCTION_CODE	N/A	Non-reserved	N/A
EACH	Non-reserved	Reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
ELASTIC	Non-reserved	N/A	N/A
ELSE	Reserved	Reserved	Reserved
ENABLE	Non-reserved	N/A	N/A
ENCLOSED	Non-reserved	N/A	N/A
ENCODING	Non-reserved	N/A	N/A
ENCRYPTED	Non-reserved	N/A	N/A
ENCRYPTED_VALUE	Non-reserved	N/A	N/A
ENCRYPTION	Non-reserved	N/A	N/A
ENCRYPTION_TYPE	Non-reserved	N/A	N/A
END	Reserved	Reserved	Reserved
END-EXEC	N/A	Reserved	Reserved
ENFORCED	Non-reserved	N/A	N/A
ENUM	Non-reserved	N/A	N/A
EOL	Non-reserved	N/A	N/A
EQUALS	N/A	Reserved	N/A
ERRORS	Non-reserved	N/A	N/A
ESCAPE	Non-reserved	Reserved	Reserved
ESCAPING	Non-reserved	N/A	N/A
EVERY	Non-reserved	Reserved	N/A
EXCEPT	Reserved	Reserved	Reserved
EXCEPTION	N/A	Reserved	Reserved
EXCHANGE	Non-reserved	N/A	N/A
EXCLUDE	Non-reserved	N/A	N/A
EXCLUDED	Reserved	N/A	N/A
EXCLUDING	Non-reserved	N/A	N/A
EXCLUSIVE	Non-reserved	N/A	N/A
EXEC	N/A	Reserved	Reserved
EXECUTE	Non-reserved	Reserved	Reserved
EXISTING	N/A	Non-reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
EXISTS	Non-reserved (excluding functions and types)	Non- reserved	Reserved
EXPIRED	Non-reserved	N/A	N/A
EXPLAIN	Non-reserved	N/A	N/A
EXTENSION	Non-reserved	N/A	N/A
EXTERNAL	Non-reserved	Reserved	Reserved
EXTRACT	Non-reserved (excluding functions and types)	Non- reserved	Reserved
FALSE	N/A	Reserved	Reserved
FAMILY	Non-reserved	N/A	N/A
FAST	Non-reserved	N/A	N/A
FEATURES	Non-reserved	N/A	N/A
FENCED	Reserved	N/A	N/A
FETCH	Reserved	Reserved	Reserved
FIELDS	Non-reserved	N/A	N/A
FILEHEADER	Non-reserved	N/A	N/A
FILL_MISSING_FIELDS	Non-reserved	N/A	N/A
FILLER	Non-reserved	N/A	N/A
FILTER	Non-reserved	N/A	Reserved
FINAL	N/A	Non- reserved	N/A
FIRST	Non-reserved	Reserved	Reserved
FIXED	Non-reserved	N/A	Reserved
FLOAT	Non-reserved (excluding functions and types)	Reserved	Reserved
FOLLOWING	Non-reserved	N/A	N/A
FOR	Reserved	Reserved	Reserved
FORCE	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
FOREIGN	Reserved	Reserved	Reserved
FORMATTER	Non-reserved	N/A	N/A
FORTRAN	N/A	Non-reserved	Non-reserved
FORWARD	Non-reserved	N/A	N/A
FOUND	N/A	Reserved	Reserved
FREE	N/A	Reserved	N/A
FREEZE	Reserved (functions and types allowed)	N/A	N/A
FROM	Reserved	Reserved	Reserved
FULL	Reserved (functions and types allowed)	Reserved	Reserved
FUNCTION	Non-reserved	Reserved	N/A
FUNCTIONS	Non-reserved	N/A	N/A
G	N/A	Non-reserved	N/A
GENERAL	N/A	Reserved	N/A
GENERATED	Non-reserved	Non-reserved	N/A
GET	N/A	Reserved	Reserved
GLOBAL	Non-reserved	Reserved	Reserved
GO	N/A	Reserved	Reserved
GOTO	N/A	Reserved	Reserved
GRANT	Reserved	Reserved	Reserved
GRANTED	Non-reserved	Non-reserved	N/A
GREATEST	Non-reserved (excluding functions and types)	N/A	N/A
GROUP	Reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
GROUPING	Non-reserved (excluding functions and types)	Reserved	N/A
GROUPPARENT	Reserved	N/A	N/A
HANDLER	Non-reserved	N/A	N/A
HAVING	Reserved	Reserved	Reserved
HDFSDIRECTORY	Reserved (functions and types allowed)	N/A	N/A
HEADER	Non-reserved	N/A	N/A
HIERARCHY	N/A	Non- reserved	N/A
HOLD	Non-reserved	Non- reserved	N/A
HOST	N/A	Reserved	N/A
HOURL	Non-reserved	Reserved	Reserved
IDENTIFIED	Non-reserved	N/A	N/A
IDENTITY	Non-reserved	Reserved	Reserved
IF	Non-reserved	N/A	N/A
IGNORE	N/A	Reserved	N/A
IGNORE_EXTRA_DATA	Non-reserved	N/A	N/A
ILIKE	Reserved (functions and types allowed)	N/A	N/A
IMMEDIATE	Non-reserved	Reserved	Reserved
IMMUTABLE	Non-reserved	N/A	N/A
IMPLEMENTATION	N/A	Non- reserved	N/A
IMPLICIT	Non-reserved	N/A	N/A
IN	Reserved	Reserved	Reserved
INCLUDE	Non-reserved	N/A	N/A
INCLUDING	Non-reserved	N/A	N/A
INCREMENT	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
INCREMENTAL	Non-reserved	N/A	N/A
INDEX	Non-reserved	N/A	N/A
INDEXES	Non-reserved	N/A	N/A
INDICATOR	N/A	Reserved	Reserved
INFILE	Non-reserved	N/A	N/A
INFIX	N/A	Non-reserved	N/A
INHERIT	Non-reserved	N/A	N/A
INHERITS	Non-reserved	N/A	N/A
INITIAL	Non-reserved	N/A	N/A
INITIALIZE	N/A	Reserved	N/A
INITIALLY	Reserved	Reserved	Reserved
INITTRANS	Non-reserved	N/A	N/A
INLINE	Non-reserved	N/A	N/A
INNER	Reserved (functions and types allowed)	Reserved	Reserved
INOUT	Non-reserved (excluding functions and types)	Reserved	N/A
INPUT	Non-reserved	Reserved	Reserved
INSENSITIVE	Non-reserved	Non-reserved	Reserved
INSERT	Non-reserved	Reserved	Reserved
INSTANCE	N/A	Non-reserved	N/A
INSTANTIABLE	N/A	Non-reserved	N/A
INSTEAD	Non-reserved	N/A	N/A
INT	Non-reserved (excluding functions and types)	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
INTEGER	Non-reserved (excluding functions and types)	Reserved	Reserved
INTERNAL	Non-reserved	N/A	N/A
INTERSECT	Reserved	Reserved	Reserved
INTERVAL	Non-reserved (excluding functions and types)	Reserved	Reserved
INTO	Reserved	Reserved	Reserved
INVOKER	Non-reserved	Non-reserved	N/A
IP	Non-reserved	N/A	N/A
IS	Reserved	Reserved	Reserved
ISNULL	Non-reserved	N/A	N/A
ISOLATION	Non-reserved	Reserved	Reserved
ITERATE	N/A	Reserved	N/A
JOIN	Reserved (functions and types allowed)	Reserved	Reserved
K	N/A	Non-reserved	N/A
KEY	Non-reserved	Reserved	Reserved
KEY_MEMBER	N/A	Non-reserved	N/A
KEY_PATH	Non-reserved	N/A	N/A
KEY_STORE	Non-reserved	N/A	N/A
KEY_TYPE	N/A	Non-reserved	N/A
KILL	Non-reserved	N/A	N/A
LABEL	Non-reserved	N/A	N/A
LANGUAGE	Non-reserved	Reserved	Reserved
LARGE	Non-reserved	Reserved	N/A
LAST	Non-reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
LATERAL	N/A	Reserved	N/A
LC_COLLATE	Non-reserved	N/A	N/A
LC_CTYPE	Non-reserved	N/A	N/A
LEADING	Reserved	Reserved	Reserved
LEAKPROOF	Non-reserved	N/A	N/A
LEAST	Non-reserved (excluding functions and types)	N/A	N/A
LEFT	Reserved (functions and types allowed)	Reserved	Reserved
LENGTH	N/A	Non- reserved	Non-reserved
LESS	Reserved	Reserved	N/A
LEVEL	Non-reserved	Reserved	Reserved
LIKE	Reserved (functions and types allowed)	Reserved	Reserved
LIMIT	Reserved	Reserved	N/A
LIST	Non-reserved	N/A	N/A
LISTEN	Non-reserved	N/A	N/A
LOAD	Non-reserved	N/A	N/A
LOCAL	Non-reserved	Reserved	Reserved
LOCALTIME	Reserved	Reserved	N/A
LOCALTIMESTAMP	Reserved	Reserved	N/A
LOCATION	Non-reserved	N/A	N/A
LOCATOR	N/A	Reserved	N/A
LOCK	Non-reserved	N/A	N/A
LOG	Non-reserved	N/A	N/A
LOGGING	Non-reserved	N/A	N/A
LOGIN_ANY	Non-reserved	N/A	N/A
LOGIN_FAILURE	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
LOGIN_SUCCESS	Non-reserved	N/A	N/A
LOGOUT	Non-reserved	N/A	N/A
LOOP	Non-reserved	N/A	N/A
LOWER	N/A	Non-reserved	Reserved
M	N/A	Non-reserved	N/A
MAP	N/A	Reserved	N/A
MAPPING	Non-reserved	N/A	N/A
MASKING	Non-reserved	N/A	N/A
MASTER	Non-reserved	N/A	N/A
MATCH	Non-reserved	Reserved	Reserved
MATCHED	Non-reserved	N/A	N/A
MATERIALIZED	Non-reserved	N/A	N/A
MAX	N/A	Non-reserved	Reserved
MAXEXTENTS	Non-reserved	N/A	N/A
MAXSIZE	Non-reserved	N/A	N/A
MAXTRANS	Non-reserved	N/A	N/A
MAXVALUE	Reserved	N/A	N/A
MERGE	Non-reserved	N/A	N/A
MESSAGE_LENGTH	N/A	Non-reserved	Non-reserved
MESSAGE_OCTET_LENGTH	N/A	Non-reserved	Non-reserved
MESSAGE_TEXT	N/A	Non-reserved	Non-reserved
METHOD	N/A	Non-reserved	N/A
MIN	N/A	Non-reserved	Reserved
MINEXTENTS	Non-reserved	N/A	N/A
MINUS	Reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
MINUTE	Non-reserved	Reserved	Reserved
MINVALUE	Non-reserved	N/A	N/A
MOD	N/A	Non-reserved	N/A
MODE	Non-reserved	N/A	N/A
MODEL	Non-reserved	N/A	N/A
MODIFIES	N/A	Reserved	N/A
MODIFY	Reserved	Reserved	N/A
MODULE	N/A	Reserved	Reserved
MONTH	Non-reserved	Reserved	Reserved
MORE	N/A	Non-reserved	Non-reserved
MOVE	Non-reserved	N/A	N/A
MOVEMENT	Non-reserved	N/A	N/A
MUMPS	N/A	Non-reserved	Non-reserved
NAME	Non-reserved	Non-reserved	Non-reserved
NAMES	Non-reserved	Reserved	Reserved
NATIONAL	Non-reserved (excluding functions and types)	Reserved	Reserved
NATURAL	Reserved (functions and types allowed)	Reserved	Reserved
NCHAR	Non-reserved (excluding functions and types)	Reserved	Reserved
NCLOB	N/A	Reserved	N/A
NEW	N/A	Reserved	N/A
NEXT	Non-reserved	Reserved	Reserved
NO	Non-reserved	Reserved	Reserved
NOCYCLE	Reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
NODE	Non-reserved	N/A	N/A
NOLOGGING	Non-reserved	N/A	N/A
NOMAXVALUE	Non-reserved	N/A	N/A
NOMINVALUE	Non-reserved	N/A	N/A
NONE	Non-reserved (excluding functions and types)	Reserved	N/A
NOT	Reserved	Reserved	Reserved
NOTHING	Non-reserved	N/A	N/A
NOTIFY	Non-reserved	N/A	N/A
NOTNULL	Reserved (functions and types allowed)	N/A	N/A
NOWAIT	Non-reserved	N/A	N/A
NULL	Reserved	Reserved	Reserved
NULLABLE	N/A	Non- reserved	Non-reserved
NULLCOLS	Non-reserved	N/A	N/A
NULLIF	Non-reserved (excluding functions and types)	Non- reserved	Reserved
NULLS	Non-reserved	N/A	N/A
NUMBER	Non-reserved (excluding functions and types)	Non- reserved	Non-reserved
NUMERIC	Non-reserved (excluding functions and types)	Reserved	Reserved
NUMSTR	Non-reserved	N/A	N/A
NVARCHAR2	Non-reserved (excluding functions and types)	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
NVL	Non-reserved (excluding functions and types)	N/A	N/A
OBJECT	Non-reserved	Reserved	N/A
OCTET_LENGTH	N/A	Non- reserved	Reserved
OF	Non-reserved	Reserved	Reserved
OFF	Non-reserved	Reserved	N/A
OFFSET	Reserved	N/A	N/A
OIDS	Non-reserved	N/A	N/A
OLD	N/A	Reserved	N/A
ON	Reserved	Reserved	Reserved
ONLY	Reserved	Reserved	Reserved
OPEN	N/A	Reserved	Reserved
OPERATION	N/A	Reserved	N/A
OPERATOR	Non-reserved	N/A	N/A
OPTIMIZATION	Non-reserved	N/A	N/A
OPTION	Non-reserved	Reserved	Reserved
OPTIONALLY	Non-reserved	N/A	N/A
OPTIONS	Non-reserved	Non- reserved	N/A
OR	Reserved	Reserved	Reserved
ORDER	Reserved	Reserved	Reserved
ORDINALITY	N/A	Reserved	N/A
OUT	Non-reserved (excluding functions and types)	Reserved	N/A
OUTER	Reserved (functions and types allowed)	Reserved	Reserved
OUTPUT	N/A	Reserved	Reserved
OVER	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
OVERLAPS	Reserved (functions and types allowed)	Non-reserved	Reserved
OVERLAY	Non-reserved (excluding functions and types)	Non-reserved	N/A
OVERRIDING	N/A	Non-reserved	N/A
OWNED	Non-reserved	N/A	N/A
OWNER	Non-reserved	N/A	N/A
PACKAGE	Non-reserved	N/A	N/A
PACKAGES	Non-reserved	N/A	N/A
PAD	N/A	Reserved	Reserved
PARAMETER	N/A	Reserved	N/A
PARAMETER_MODE	N/A	Non-reserved	N/A
PARAMETER_NAME	N/A	Non-reserved	N/A
PARAMETER_ORDINAL_POSITION	N/A	Non-reserved	N/A
PARAMETER_SPECIFIC_CATALOG	N/A	Non-reserved	N/A
PARAMETER_SPECIFIC_NAME	N/A	Non-reserved	N/A
PARAMETER_SPECIFIC_SCHEMA	N/A	Non-reserved	N/A
PARAMETERS	N/A	Reserved	N/A
PARSER	Non-reserved	N/A	N/A
PARTIAL	Non-reserved	Reserved	Reserved
PARTITION	Non-reserved	N/A	N/A
PARTITIONS	Non-reserved	N/A	N/A
PASCAL	N/A	Non-reserved	Non-reserved
PASSING	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
PASSWORD	Non-reserved	N/A	N/A
PATH	N/A	Reserved	N/A
PCTFREE	Non-reserved	N/A	N/A
PER	Non-reserved	N/A	N/A
PERCENT	Non-reserved	N/A	N/A
PERFORMANCE	Reserved	N/A	N/A
PERM	Non-reserved	N/A	N/A
PLACING	Reserved	N/A	N/A
PLAN	Non-reserved	N/A	N/A
PLANS	Non-reserved	N/A	N/A
PLI	N/A	Non-reserved	Non-reserved
POLICY	Non-reserved	N/A	N/A
POOL	Non-reserved	N/A	N/A
POSITION	Non-reserved (excluding functions and types)	Non-reserved	Reserved
POSTFIX	N/A	Reserved	N/A
PRECEDING	Non-reserved	N/A	N/A
PRECISION	Non-reserved (excluding functions and types)	Reserved	Reserved
PREDICT	Non-reserved	N/A	N/A
PREFERRED	Non-reserved	N/A	N/A
PREFIX	Non-reserved	Reserved	N/A
PREORDER	N/A	Reserved	N/A
PREPARE	Non-reserved	Reserved	Reserved
PREPARED	Non-reserved	N/A	N/A
PRESERVE	Non-reserved	Reserved	Reserved
PRIMARY	Reserved	Reserved	Reserved
PRIOR	Non-reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
PRIORER	Reserved	N/A	N/A
PRIVATE	Non-reserved	N/A	N/A
PRIVILEGE	Non-reserved	N/A	N/A
PRIVILEGES	Non-reserved	Reserved	Reserved
PROCEDURAL	Non-reserved	N/A	N/A
PROCEDURE	Reserved	Reserved	Reserved
PROFILE	Non-reserved	N/A	N/A
PUBLIC	N/A	Reserved	Reserved
PUBLICATION	Non-reserved	N/A	N/A
PUBLISH	Non-reserved	N/A	N/A
PURGE	Non-reserved	N/A	N/A
QUERY	Non-reserved	N/A	N/A
QUOTE	Non-reserved	N/A	N/A
RANDOMIZED	Non-reserved	N/A	N/A
RANGE	Non-reserved	N/A	N/A
RATIO	Non-reserved	N/A	N/A
RAW	Non-reserved	N/A	N/A
READ	Non-reserved	Reserved	Reserved
READS	N/A	Reserved	N/A
REAL	Non-reserved (excluding functions and types)	Reserved	Reserved
REASSIGN	Non-reserved	N/A	N/A
REBUILD	Non-reserved	N/A	N/A
RECHECK	Non-reserved	N/A	N/A
RECURSIVE	Non-reserved	Reserved	N/A
RECYCLEBIN	Reserved (functions and types allowed)	N/A	N/A
REDISANYVALUE	Non-reserved	N/A	N/A
REF	Non-reserved	Reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
REFERENCES	Reserved	Reserved	Reserved
REFERENCING	N/A	Reserved	N/A
REFRESH	Non-reserved	N/A	N/A
REINDEX	Non-reserved	N/A	N/A
REJECT	Reserved	N/A	N/A
RELATIVE	Non-reserved	Reserved	Reserved
RELEASE	Non-reserved	N/A	N/A
REOPTIONS	Non-reserved	N/A	N/A
REMOTE	Non-reserved	N/A	N/A
REMOVE	Non-reserved	N/A	N/A
RENAME	Non-reserved	N/A	N/A
REPEATABLE	Non-reserved	Non-reserved	Non-reserved
REPLACE	Non-reserved	N/A	N/A
REPLICA	Non-reserved	N/A	N/A
RESET	Non-reserved	N/A	N/A
RESIZE	Non-reserved	N/A	N/A
RESOURCE	Non-reserved	N/A	N/A
RESTART	Non-reserved	N/A	N/A
RESTRICT	Non-reserved	Reserved	Reserved
RESULT	N/A	Reserved	N/A
RETURN	Non-reserved	Reserved	N/A
RETURNED_LENGTH	N/A	Non-reserved	Non-reserved
RETURNED_OCTET_LENGTH	N/A	Non-reserved	Non-reserved
RETURNED_SQLSTATE	N/A	Non-reserved	Non-reserved
RETURNING	Reserved	N/A	N/A
RETURNS	Non-reserved	Reserved	N/A
REUSE	Non-reserved	N/A	N/A
REVOKE	Non-reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
RIGHT	Reserved (functions and types allowed)	Reserved	Reserved
ROLE	Non-reserved	Reserved	N/A
ROLES	Non-reserved	N/A	N/A
ROLLBACK	Non-reserved	Reserved	Reserved
ROLLUP	Non-reserved	Reserved	N/A
ROTATION	Non-reserved	N/A	N/A
ROUTINE	N/A	Reserved	N/A
ROUTINE_CATALOG	N/A	Non-reserved	N/A
ROUTINE_NAME	N/A	Non-reserved	N/A
ROUTINE_SCHEMA	N/A	Non-reserved	N/A
ROW	Non-reserved (excluding functions and types)	Reserved	N/A
ROW_COUNT	N/A	Non-reserved	Non-reserved
ROWNUM	Reserved	N/A	N/A
ROWS	Non-reserved	Reserved	Reserved
ROWTYPE	Non-reserved	N/A	N/A
RULE	Non-reserved	N/A	N/A
SAMPLE	Non-reserved	N/A	N/A
SAVEPOINT	Non-reserved	Reserved	N/A
SCALE	N/A	Non-reserved	Non-reserved
SCHEMA	Non-reserved	Reserved	Reserved
SCHEMA_NAME	N/A	Non-reserved	Non-reserved
SCOPE	N/A	Reserved	N/A
SCROLL	Non-reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
SEARCH	Non-reserved	Reserved	N/A
SECOND	Non-reserved	Reserved	Reserved
SECTION	N/A	Reserved	Reserved
SECURITY	Non-reserved	Non-reserved	N/A
SELECT	Reserved	Reserved	Reserved
SELF	N/A	Non-reserved	N/A
SENSITIVE	N/A	Non-reserved	N/A
SEQUENCE	Non-reserved	Reserved	N/A
SEQUENCES	Non-reserved	N/A	N/A
SERIALIZABLE	Non-reserved	Non-reserved	Non-reserved
SERVER	Non-reserved	N/A	N/A
SERVER_NAME	N/A	Non-reserved	Non-reserved
SESSION	Non-reserved	Reserved	Reserved
SESSION_USER	Reserved	Reserved	Reserved
SET	Non-reserved	Reserved	Reserved
SETOF	Non-reserved (excluding functions and types)	N/A	N/A
SETS	Non-reserved	Reserved	N/A
SHARE	Non-reserved	N/A	N/A
SHIPPABLE	Non-reserved	N/A	N/A
SHOW	Non-reserved	N/A	N/A
SHUTDOWN	Non-reserved	N/A	N/A
SIBLINGS	Non-reserved	N/A	N/A
SIMILAR	Reserved (functions and types allowed)	Non-reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
SIMPLE	Non-reserved	Non-reserved	N/A
SIZE	Non-reserved	Reserved	Reserved
SKIP	Non-reserved	N/A	N/A
SLICE	Non-reserved	N/A	N/A
SMALLDATETIME	Non-reserved (excluding functions and types)	N/A	N/A
SMALLDATETIME_FORMAT	Non-reserved	N/A	N/A
SMALLINT	Non-reserved (excluding functions and types)	Reserved	Reserved
SNAPSHOT	Non-reserved	N/A	N/A
SOME	Reserved	Reserved	Reserved
SOURCE	Non-reserved	Non-reserved	N/A
SPACE	Non-reserved	Reserved	Reserved
SPECIFIC	N/A	Reserved	N/A
SPECIFIC_NAME	N/A	Non-reserved	N/A
SPECIFICTYPE	N/A	Reserved	N/A
SPILL	Non-reserved	N/A	N/A
SPLIT	Non-reserved	N/A	N/A
SQL	N/A	Reserved	Reserved
SQLCODE	N/A	N/A	Reserved
SQLERROR	N/A	N/A	Reserved
SQL EXCEPTION	N/A	Reserved	N/A
SQLSTATE	N/A	Reserved	Reserved
SQLWARNING	N/A	Reserved	N/A
STABLE	Non-reserved	N/A	N/A
STANDALONE	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
START	Non-reserved	Reserved	N/A
STATE	N/A	Reserved	N/A
STATEMENT	Non-reserved	Reserved	N/A
STATEMENT_ID	Non-reserved	N/A	N/A
STATIC	N/A	Reserved	N/A
STATISTICS	Non-reserved	N/A	N/A
STDIN	Non-reserved	N/A	N/A
STDOUT	Non-reserved	N/A	N/A
STORAGE	Non-reserved	N/A	N/A
STORE	Non-reserved	N/A	N/A
STORED	Non-reserved	N/A	N/A
STRATIFY	Non-reserved	N/A	N/A
STREAM	Non-reserved	N/A	N/A
STRICT	Non-reserved	N/A	N/A
STRIP	Non-reserved	N/A	N/A
STRUCTURE	N/A	Reserved	N/A
STYLE	N/A	Non-reserved	N/A
SUBCLASS_ORIGIN	N/A	Non-reserved	Non-reserved
SUBLIST	N/A	Non-reserved	N/A
SUBPARTITION	Non-reserved	N/A	N/A
SUBSCRIPTION	Non-reserved	N/A	N/A
SUBSTRING	Non-reserved (excluding functions and types)	Non-reserved	Reserved
SUM	N/A	Non-reserved	Reserved
SYMMETRIC	Reserved	Non-reserved	N/A
SYNONYM	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
SYS_REFCURSOR	Non-reserved	N/A	N/A
SYSDATE	Reserved	N/A	N/A
SYSID	Non-reserved	N/A	N/A
SYSTEM	Non-reserved	Non-reserved	N/A
SYSTEM_USER	N/A	Reserved	Reserved
TABLE	Reserved	Reserved	Reserved
TABLE_NAME	N/A	Non-reserved	Non-reserved
TABLES	Non-reserved	N/A	N/A
TABLESAMPLE	Reserved (functions and types allowed)	N/A	N/A
TABLESPACE	Non-reserved	N/A	N/A
TARGET	Non-reserved	N/A	N/A
TEMP	Non-reserved	N/A	N/A
TEMPLATE	Non-reserved	N/A	N/A
TEMPORARY	Non-reserved	Reserved	Reserved
TERMINATE	N/A	Reserved	N/A
TERMINATED	Non-reserved	N/A	N/A
TEXT	Non-reserved	N/A	N/A
THAN	Non-reserved	Reserved	N/A
THEN	Reserved	Reserved	Reserved
TIME	Non-reserved (excluding functions and types)	Reserved	Reserved
TIME_FORMAT	Non-reserved	N/A	N/A
TIMECAPSULE	Reserved (functions and types allowed)	N/A	N/A
TIMESTAMP	Non-reserved (excluding functions and types)	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
TIMESTAMP_FORMAT	Non-reserved	N/A	N/A
TIMESTAMPDIFF	Non-reserved (excluding functions and types)	N/A	N/A
TIMEZONE_HOUR	N/A	Reserved	Reserved
TIMEZONE_MINUTE	N/A	Reserved	Reserved
TINYINT	Non-reserved (excluding functions and types)	N/A	N/A
TO	Reserved	Reserved	Reserved
TRAILING	Reserved	Reserved	Reserved
TRANSACTION	Non-reserved	Reserved	Reserved
TRANSACTION_ACTIVE	N/A	Non- reserved	N/A
TRANSACTIONS_COMMITTED	N/A	Non- reserved	N/A
TRANSACTIONS_ROLLED_BACK	N/A	Non- reserved	N/A
TRANSFORM	Non-reserved	Non- reserved	N/A
TRANSFORMS	N/A	Non- reserved	N/A
TRANSLATE	N/A	Non- reserved	Reserved
TRANSLATION	N/A	Reserved	Reserved
TREAT	Non-reserved (excluding functions and types)	Reserved	N/A
TRIGGER	Non-reserved	Reserved	N/A
TRIGGER_CATALOG	N/A	Non- reserved	N/A
TRIGGER_NAME	N/A	Non- reserved	N/A

Keyword	GaussDB	SQL:1999	SQL-92
TRIGGER_SCHEMA	N/A	Non-reserved	N/A
TRIM	Non-reserved (excluding functions and types)	Non-reserved	Reserved
TRUE	N/A	Reserved	Reserved
TRUNCATE	Non-reserved	N/A	N/A
TRUSTED	Non-reserved	N/A	N/A
TSFIELD	Non-reserved	N/A	N/A
TSTAG	Non-reserved	N/A	N/A
TSTIME	Non-reserved	N/A	N/A
TYPE	Non-reserved	Non-reserved	Non-reserved
TYPES	Non-reserved	N/A	N/A
UNBOUNDED	Non-reserved	N/A	N/A
UNCOMMITTED	Non-reserved	Non-reserved	Non-reserved
UNDER	N/A	Reserved	N/A
UNENCRYPTED	Non-reserved	N/A	N/A
UNION	Reserved	Reserved	Reserved
UNIQUE	Reserved	Reserved	Reserved
UNKNOWN	Non-reserved	Reserved	Reserved
UNLIMITED	Non-reserved	N/A	N/A
UNLISTEN	Non-reserved	N/A	N/A
UNLOCK	Non-reserved	N/A	N/A
UNLOGGED	Non-reserved	N/A	N/A
UNNAMED	N/A	Non-reserved	Non-reserved
UNNEST	N/A	Reserved	N/A
UNTIL	Non-reserved	N/A	N/A
UNUSABLE	Non-reserved	N/A	N/A
UPDATE	Non-reserved	Reserved	Reserved

Keyword	GaussDB	SQL:1999	SQL-92
UPPER	N/A	Non-reserved	Reserved
USAGE	N/A	Reserved	Reserved
USEEOF	Non-reserved	N/A	N/A
USER	Reserved	Reserved	Reserved
USER_DEFINED_TYPE_CATALOG	N/A	Non-reserved	N/A
USER_DEFINED_TYPE_NAME	N/A	Non-reserved	N/A
USER_DEFINED_TYPE_SCHEMA	N/A	Non-reserved	N/A
USING	Reserved	Reserved	Reserved
VACUUM	Non-reserved	N/A	N/A
VALID	Non-reserved	N/A	N/A
VALIDATE	Non-reserved	N/A	N/A
VALIDATION	Non-reserved	N/A	N/A
VALIDATOR	Non-reserved	N/A	N/A
VALUE	Non-reserved	Reserved	Reserved
VALUES	Non-reserved (excluding functions and types)	Reserved	Reserved
VARCHAR	Non-reserved (excluding functions and types)	Reserved	Reserved
VARCHAR2	Non-reserved (excluding functions and types)	N/A	N/A
VARIABLE	N/A	Reserved	N/A
VARIABLES	Non-reserved	N/A	N/A
VARIADIC	Reserved	N/A	N/A
VARYING	Non-reserved	Reserved	Reserved
VCGROUP	Non-reserved	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
VERBOSE	Reserved (functions and types allowed)	N/A	N/A
VERIFY	Reserved	N/A	N/A
VERSION	Non-reserved	N/A	N/A
VIEW	Non-reserved	Reserved	Reserved
VOLATILE	Non-reserved	N/A	N/A
WAIT	Non-reserved	N/A	N/A
WEAK	Non-reserved	N/A	N/A
WHEN	Reserved	Reserved	Reserved
WHENEVER	N/A	Reserved	Reserved
WHERE	Reserved	Reserved	Reserved
WHITESPACE	Non-reserved	N/A	N/A
WINDOW	Reserved	N/A	N/A
WITH	Reserved	Reserved	Reserved
WITHIN	Non-reserved	N/A	N/A
WITHOUT	Non-reserved	Reserved	N/A
WORK	Non-reserved	Reserved	Reserved
WORKLOAD	Non-reserved	N/A	N/A
WRAPPER	Non-reserved	N/A	N/A
WRITE	Non-reserved	Reserved	Reserved
XML	Non-reserved	N/A	N/A
XMLATTRIBUTES	Non-reserved (excluding functions and types)	N/A	N/A
XMLCONCAT	Non-reserved (excluding functions and types)	N/A	N/A
XMLELEMENT	Non-reserved (excluding functions and types)	N/A	N/A

Keyword	GaussDB	SQL:1999	SQL-92
XMLEXISTS	Non-reserved (excluding functions and types)	N/A	N/A
XMLFOREST	Non-reserved (excluding functions and types)	N/A	N/A
XMLPARSE	Non-reserved (excluding functions and types)	N/A	N/A
XMLPI	Non-reserved (excluding functions and types)	N/A	N/A
XMLROOT	Non-reserved (excluding functions and types)	N/A	N/A
XMLSERIALIZE	Non-reserved (excluding functions and types)	N/A	N/A
YEAR	Non-reserved	Reserved	Reserved
YES	Non-reserved	N/A	N/A
ZONE	Non-reserved	Reserved	Reserved

Fields listed in the following table cannot be used as column names during table creation.

CTID	XMIN	CMIN	XMAX	CMAX
TABLEOID	XC_NODE_ID	TID	GS_TUPLE_UID	TABLEBUCKETID

7.3 Data Type

Data type is a basic attribute of data that used to distinguish different types of data. Different data types occupy different storage space and support different operations. Data is stored in data tables in the database. Each column of a data

table defines the data type. During storage, data must be stored according to data types.

GaussDB supports implicit conversions between certain data types. For details, see [PG_CAST](#).

7.3.1 Numeric Types

Table 7-2 lists all available types. For digit operators and related built-in functions, see [Mathematical Functions and Operators](#).

Table 7-2 Integer types

Name	Description	Storage Space	Range
TINYINT	Tiny integer, also called INT1	1 byte	0-255
SMALLINT	Small integer, also called INT2	2 bytes	-32,768 to +32,767
INTEGER	Typical choice for integers, also called INT4	4 bytes	-2,147,483,648 to +2,147,483,647
BINARY_INTEGER	INTEGER alias, which is an A-compatible database type	4 bytes	-2,147,483,648 to +2,147,483,647
BIGINT	Big integer, also called INT8	8 bytes	-9,223,372,036,854,775,808 to +9,223,372,036,854,775,807
int16	A 16-byte integer cannot be used to create tables.	16 bytes	-170,141,183,460,469,231,731,687,303,715,884,105,728 to +170,141,183,460,469,231,731,687,303,715,884,105,727

Example:

```
-- Create a table containing TINYINT data.
openGauss=# CREATE TABLE int_type_t1
(
    IT_COL1 TINYINT
);

-- Insert data to the created table.
openGauss=# INSERT INTO int_type_t1 VALUES(10);

-- View data.
openGauss=# SELECT * FROM int_type_t1;
it_col1
-----
10
```



```
(1 row)

-- Delete the table.
openGauss=# DROP TABLE int_type_t1;
-- Create a table containing TINYINT, INTEGER, and BIGINT data.
openGauss=# CREATE TABLE int_type_t2
(
  a TINYINT,
  b TINYINT,
  c INTEGER,
  d BIGINT
);

-- Insert data.
openGauss=# INSERT INTO int_type_t2 VALUES(100, 10, 1000, 10000);

-- View data.
openGauss=# SELECT * FROM int_type_t2;
 a | b | c | d
-----+-----+-----+-----
100 | 10 | 1000 | 10000
(1 row)

-- Delete the table.
openGauss=# DROP TABLE int_type_t2;
```

NOTE

- Only numbers of the TINYINT, SMALLINT, INTEGER, BIGINT, or INT16 type, that is, integers can be stored. Saving a number with a decimal in any of the data types will result in errors.
- The INTEGER type is the common choice, as it offers the best balance between range, storage size, and performance. Generally, use the SMALLINT type only if you are sure that the value range is within the SMALLINT value range. The storage speed of INTEGER is much faster. BIGINT is used only when the range of INTEGER is not large enough.

Table 7-3 Arbitrary precision types

Name	Description	Storage Space	Range
NUMERIC[(p[,s])], DECIMAL[(p[,s])]	The value range of p (precision) is [1,1000], and the value range of s (scale) is [0, <i>p</i>]. NOTE p indicates the total digits, and s indicates the decimal digit.	The precision is specified by users. Every four decimal digits occupy two bytes, and an extra eight-byte overhead is added to the entire data.	Up to 131,072 digits before the decimal point; and up to 16,383 digits after the decimal point when no precision is specified
NUMBER[(p[,s])]	Alias of the NUMERIC type.	The precision is specified by users. Every four decimal digits occupy two bytes, and an extra eight-byte overhead is added to the entire data.	Up to 131,072 digits before the decimal point; and up to 16,383 digits after the decimal point when no precision is specified

Example:

```
-- Create a table.
openGauss=# CREATE TABLE decimal_type_t1
(
    DT_COL1 DECIMAL(10,4)
);

-- Insert data.
openGauss=# INSERT INTO decimal_type_t1 VALUES(123456.122331);

-- Query data in the table.
openGauss=# SELECT * FROM decimal_type_t1;
 dt_col1
-----
123456.1223
(1 row)

-- Delete the table.
openGauss=# DROP TABLE decimal_type_t1;
-- Create a table.
openGauss=# CREATE TABLE numeric_type_t1
(
    NT_COL1 NUMERIC(10,4)
);

-- Insert data.
openGauss=# INSERT INTO numeric_type_t1 VALUES(123456.12354);

-- Query data in the table.
openGauss=# SELECT * FROM numeric_type_t1;
 nt_col1
-----
123456.1235
(1 row)

-- Delete the table.
openGauss=# DROP TABLE numeric_type_t1;
```

** NOTE**

- Compared to the integer types, the arbitrary precision numbers require larger storage space and have lower storage efficiency, operation efficiency, and poorer compression ratio results. The integer type is the common choice when number types are defined. Arbitrary precision numbers are used only when numbers exceed the maximum range indicated by the integers.
- When NUMERIC/DECIMAL is used for defining a column, you are advised to specify the precision (p) and scale (s) for the column.

Table 7-4 Sequence integer

Name	Description	Storage Space	Range
SMALLSERIAL	Two-byte serial integer	2 bytes.	-32,768 to +32,767
SERIAL	Four-byte serial integer	4 bytes.	-2,147,483,648 to +2,147,483,647
BIGSERIAL	Eight-byte serial integer	8 bytes.	-9,223,372,036,854,775,808 to +9,223,372,036,854,775,807

Name	Description	Storage Space	Range
LARGESERIAL	By default, a 16-byte auto-incrementing integer is inserted. The actual value type is the same as that of numeric.	Variable-length type. Every four decimal digits occupy two bytes, and an extra eight-byte overhead is added to the entire data.	There can be a maximum of 131,072 digits before the decimal point and 16,383 digits after the decimal point.

Example:

```
-- Create a table.
openGauss=# CREATE TABLE smallserial_type_tab(a SMALLSERIAL);

-- Insert data.
openGauss=# INSERT INTO smallserial_type_tab VALUES(default);

-- Insert data again.
openGauss=# INSERT INTO smallserial_type_tab VALUES(default);

-- View data.
openGauss=# SELECT * FROM smallserial_type_tab;
a
----
1
2
(2 rows)

-- Create a table.
openGauss=# CREATE TABLE serial_type_tab(b SERIAL);

-- Insert data.
openGauss=# INSERT INTO serial_type_tab VALUES(default);

-- Insert data again.
openGauss=# INSERT INTO serial_type_tab VALUES(default);

-- View data.
openGauss=# SELECT * FROM serial_type_tab;
b
----
1
2
(2 rows)

-- Create a table.
openGauss=# CREATE TABLE bigserial_type_tab(c BIGSERIAL);

-- Insert data.
openGauss=# INSERT INTO bigserial_type_tab VALUES(default);

-- Insert data again.
openGauss=# INSERT INTO bigserial_type_tab VALUES(default);

-- View data.
openGauss=# SELECT * FROM bigserial_type_tab;
c
----
1
2
```

```
(2 rows)

-- Create a table.
openGauss=# CREATE TABLE largeserial_type_tab(c LARGESERIAL);

-- Insert data.
openGauss=# INSERT INTO largeserial_type_tab VALUES(default);

-- Insert data.
openGauss=# INSERT INTO largeserial_type_tab VALUES(default);

-- View data.
openGauss=# SELECT * FROM largeserial_type_tab;
 c
---
 1
 2
(2 rows)

-- Delete the table.
openGauss=# DROP TABLE smallserial_type_tab;

openGauss=# DROP TABLE serial_type_tab;

openGauss=# DROP TABLE bigserial_type_tab;
```

 NOTE

SMALLSERIAL, SERIAL, BIGSERIAL, and LARGESERIAL are not real types. They are concepts used for setting a unique identifier for a table. Therefore, an integer column is created and its default value plans to be read from a sequencer. A NOT NULL constraint is used to ensure NULL is not inserted. In most cases you would also want to attach a **UNIQUE** or **PRIMARY KEY** constraint to prevent duplicate values from being inserted unexpectedly, but this is not automatic. Finally, the sequencer belongs to the column. In this case, when the column or the table is deleted, the sequencer is also deleted. Currently, the **SERIAL** column can be specified only when you create a table. You cannot add the **SERIAL** column in an existing table. In addition, **SERIAL** columns cannot be created in temporary tables. Because SERIAL is not a data type, columns cannot be converted to this type.

Table 7-5 Floating point types

Name	Description	Storage Space	Range
REAL, FLOAT4	Single precision floating points, inexact	4 bytes.	-3.402E+38 to +3.402E+38, 6-bit decimal digits.
DOUBLE PRECISION , FLOAT8	Double precision floating points, inexact	8 bytes.	-1.79E+308 to +1.79E+308, 15-bit decimal digits.
FLOAT[(p)]	Floating points, inexact. The value range of p (precision) is [1,53]. NOTE p is the precision, indicating the total number of binary bits.	4 bytes or 8 bytes.	REAL or DOUBLE PRECISION is selected as an internal identifier based on precision (p). If no precision is specified, DOUBLE PRECISION is used as the internal identifier.

Name	Description	Storage Space	Range
BINARY_DOUBLE	Alias of DOUBLE PRECISION.	8 bytes.	-1.79E+308 to +1.79E+308, 15-bit decimal digits.
DEC[(p[,s])]	The value range of p is [1,1000], and the value range of s is [0,<i>p</i>]. If the scenario in the following "NOTE" is met and the precision and scale are not specified, the precision p is 10 and the scale s is 0 by default. This type maps to NUMERIC. For details about the application scenario, see NUMERIC.	The precision is specified by users. Every four decimal digits occupy two bytes, and an extra eight-byte overhead is added to the entire data.	Maximum 131,072 digits before the decimal point and 16,383 digits after the decimal point when the precision and scale are specified to the maximum.
INTEGER[(p[,s])]	The value range of p is [1,1000], and the value range of s is [0,<i>p</i>]. If the precision and scale are not specified, the precision p is 10 and the scale s is 0 by default. If the precision and scale are not specified, this type is mapped to INTEGER. If the precision and scale are specified, this type is mapped to NUMERIC.	The precision is specified by users. Every four decimal digits occupy two bytes, and an extra eight-byte overhead is added to the entire data.	Maximum 131,072 digits before the decimal point and 16,383 digits after the decimal point when the precision and scale are specified to the maximum. If the precision and scale are not specified, the value ranges from -2,147,483,648 to +2,147,483,647.

 NOTE

The binary floating-point data types REAL, FLOAT4, DOUBLE, DOUBLE PRECISION, FLOAT8, FLOAT[(p)], and BINARY_DOUBLE are imprecise numeric types. Their internal storage is approximate, which may result in slight differences during storage and retrieval. When using binary floating-point data types, pay attention to the following points:

- Precise storage and calculation: If precise storage and calculation are required (such as for monetary values), use exact data types (for example, numeric) instead.
- Complex calculation: When you perform complex calculation with imprecise data types to obtain critical data, it is important to carefully evaluate the results.
- Floating-point number comparison: The result of comparing two floating-point numbers for equality may differ from expectations.
- Underflow error: If a floating-point number is too close to zero, it cannot be accurately represented. As a result, an underflow error occurs.

Example:

```
-- Create a table.
openGauss=# CREATE TABLE float_type_t2
(
    FT_COL1 INTEGER,
    FT_COL2 FLOAT4,
    FT_COL3 FLOAT8,
    FT_COL4 FLOAT(3),
    FT_COL5 BINARY_DOUBLE,
    FT_COL6 DECIMAL(10,4),
    FT_COL7 INTEGER(6,3)
);

-- Insert data.
openGauss=# INSERT INTO float_type_t2 VALUES(10,10.365456,123456.1234,10.3214, 321.321, 123.123654,
123.123654);

-- View data.
openGauss=# SELECT * FROM float_type_t2 ;
 ft_col1 | ft_col2 | ft_col3 | ft_col4 | ft_col5 | ft_col6 | ft_col7
-----+-----+-----+-----+-----+-----+-----
      10 | 10.3655 | 123456.1234 | 10.3214 | 321.321 | 123.1237 | 123.124
(1 row)

-- Delete the table.
openGauss=# DROP TABLE float_type_t2;
```

7.3.2 Monetary Types

The money type stores a currency amount with fixed fractional precision.

The range shown in [Table 7-6](#) assumes there are two fractional digits. Input is accepted in a variety of formats, including integer and floating-point literals, as well as typical currency formatting, such as "\$1,000.00". Output is generally in the last format but depends on the locale.

Table 7-6 Monetary type

Name	Description	Storage Space	Range
money	Currency amount	8 bytes	-92233720368547758.08 to +92233720368547758.07

Values of the numeric, int, and bigint data types can be cast to money. Conversion from the real or double precision data type to a money value can be done by casting to numeric type first, for example:

```
openGauss=# SELECT '12.34'::float8::numeric::money;  
money  
-----  
$12.34  
(1 row)
```

However, this is not recommended. Floating point numbers should not be used to handle money due to the potential for rounding errors.

A money value can be cast to numeric without loss of precision. Conversion to other types could potentially lose precision, and must also be done in two stages:

```
openGauss=# SELECT '52093.89'::money::numeric::float8;  
float8  
-----  
52093.89  
(1 row)
```

When a money value is divided by another money value, the result is of the double precision type (that is, a pure number, not money); the currency units cancel each other out in the division.

7.3.3 Boolean Types

Table 7-7 Boolean types

Name	Description	Storage Space	Value
BOOLEAN	Boolean type	1 byte	• true • false • null (unknown)

- Valid literal values for the "true" state include:
TRUE, **'t'**, **'true'**, **'y'**, **'yes'**, **'1'**, **'TRUE'**, **true**, **on**, and all non-zero integers.
- Valid literal values for the "false" state include:
FALSE, **'f'**, **'false'**, **'n'**, **'no'**, **'0'**, **0**, **'FALSE'**, **false**, and **off**.

TRUE and **FALSE** are standard expressions, compatible with SQL statements.

Examples

Boolean values are displayed using the letters t and f.

```
-- Create a table.  
openGauss=# CREATE TABLE bool_type_t1  
(  
    BT_COL1 BOOLEAN,  
    BT_COL2 TEXT  
);  
  
-- Insert data.  
openGauss=# INSERT INTO bool_type_t1 VALUES (TRUE, 'sic est');
```

```
openGauss=# INSERT INTO bool_type_t1 VALUES (FALSE, 'non est');

-- View data.
openGauss=# SELECT * FROM bool_type_t1;
 bt_col1 | bt_col2 
-----+-----
 t      | sic est
 f      | non est
(2 rows)

openGauss=# SELECT * FROM bool_type_t1 WHERE bt_col1 = 't';
 bt_col1 | bt_col2 
-----+-----
 t      | sic est
(1 row)

-- Delete the table.
openGauss=# DROP TABLE bool_type_t1;
```

7.3.4 Character Types

Table 7-8 lists the character data types supported by GaussDB. For string operators and related built-in functions, see [Character Processing Functions and Operators](#).

Table 7-8 Character types

Name	Description	Storage Space
CHAR(n) CHARACTER(n) NCHAR(n)	Fixed-length character string. Empty characters are filled in with blank spaces. n indicates the byte length. If it is not specified, the default precision 1 is used.	The maximum size is 10 MB.
VARCHAR(n) CHARACTER VARYING(n)	Variable-length string. In PG-compatible mode, n indicates the string length. In other compatibility modes, n indicates the byte length.	The maximum size is 10 MB. If n is not contained, the maximum storage length is 1 GB – 85 – Length of the first <i>n</i> columns. For example, the maximum length (of a int, b varchar) is 1 GB – 85 – 4 = 1,073,741,735.
VARCHAR2(n)	Variable-length string. It is the alias of the VARCHAR(n) type. n indicates the byte length.	The maximum size is 10 MB. If n is not contained, the maximum storage length is 1 GB – 85 – Length of the first <i>n</i> columns. For example, the maximum length (of a int, b varchar) is 1 GB – 85 – 4 = 1,073,741,735.

Name	Description	Storage Space
NVARCHAR2(n)	Variable-length string. n indicates the string length.	The maximum size is 10 MB. If n is not contained, the maximum storage length is 1 GB – 85 – Length of the first <i>n</i> columns. For example, the maximum length (of a int, b varchar) is 1 GB – 85 – 4 = 1,073,741,735.
TEXT	Variable-length string.	The maximum size is 1 GB minus 1 byte. However, the size of the column description header and the size of the tuple (less than 1 GB minus 1 byte) where the column is located must also be considered. Therefore, the maximum size of the TEXT type may be less than 1 GB minus 1 byte.
CLOB	Big text object. It is the alias of the TEXT type.	The maximum size is 32 TB minus 1 byte. However, the size of the column description header and the size of the tuple (less than 32 TB minus 1 byte) where the column is located must also be considered. Therefore, the maximum size of the CLOB type may be less than 32 TB minus 1 byte.

 NOTE

1. In addition to the restriction on the size of each column, the total size of each tuple cannot exceed 1 GB minus 1 byte and is affected by the control header information of the column, the control header information of the tuple, and whether null fields exist in the tuple.
2. NCHAR is the alias of the bpchar type, and VARCHAR2(n) is the alias of the VARCHAR(n) type.
3. Only advanced package db_lob supports CLOBs whose size is greater than 1 GB. System functions do not support CLOBs whose size is greater than 1 GB.

In GaussDB, there are two other fixed-length character types, as shown in [Table 7-9](#). The **name** type exists only for the storage of identifiers in the internal system catalogs and is not intended for use by general users. Its length is currently defined as 64 bytes (63 usable characters plus terminator). The char type uses only one byte of storage. It is internally used in the system catalogs as a simplistic enumeration type.

Table 7-9 Special character types

Name	Description	Storage Space
name	Internal type for object names	64 bytes
"char"	Single-byte internal type	1 byte

Examples

```
-- Create a table.
openGauss=# CREATE TABLE char_type_t1
(
    CT_COL1 CHARACTER(4)
);

-- Insert data.
openGauss=# INSERT INTO char_type_t1 VALUES ('ok');

-- Query data in the table.
openGauss=# SELECT ct_col1, char_length(ct_col1) FROM char_type_t1;
 ct_col1 | char_length
-----+-----
    ok   |          4
(1 row)

-- Delete the table.
openGauss=# DROP TABLE char_type_t1;

-- Create a table.
openGauss=# CREATE TABLE char_type_t2
(
    CT_COL1 VARCHAR(5)
);

-- Insert data.
openGauss=# INSERT INTO char_type_t2 VALUES ('ok');

openGauss=# INSERT INTO char_type_t2 VALUES ('good');

-- Specify the type length. An error is reported if an inserted string exceeds this length.
openGauss=# INSERT INTO char_type_t2 VALUES ('too long');
ERROR: value too long for type character varying(5)
CONTEXT: referenced column: ct_col1

-- Specify the type length. A string exceeding this length is truncated.
openGauss=# INSERT INTO char_type_t2 VALUES ('too long'::varchar(5));

-- Query data.
openGauss=# SELECT ct_col1, char_length(ct_col1) FROM char_type_t2;
 ct_col1 | char_length
-----+-----
    ok   |          2
    good |          4
    too l |          5
(3 rows)

-- Delete data.
openGauss=# DROP TABLE char_type_t2;
```

7.3.5 Binary Types

Table 7-10 lists the binary data types supported by GaussDB.

Table 7-10 Binary data types

Name	Description	Storage Space
BLOB	Binary large object (BLOB). Currently, BLOB only supports the following external access APIs: • DBE_LOB.GET_LENGTH • DBE_LOB.READ • DBE_LOB.WRITE • DBE_LOB.WRITE_APPEND • DBE_LOB.COPY • DBE_LOB.ERASE For details about the APIs, see DBE_LOB .	The maximum size is 32 TB (35,184,372,088,832 bytes).
RAW	Variable-length hexadecimal string.	4 bytes plus the actual hexadecimal string. Its maximum length is 1073733621 bytes (1 GB – 8203 bytes).
BYTEA	Variable-length binary string.	4 bytes plus the actual binary string. Its maximum length is 1073733621 bytes (1 GB – 8203 bytes).

 NOTE

In addition to the size limitation on each column, the total size of each tuple cannot exceed 1,073,733,621 bytes (that is, 1 GB minus 8,203 bytes).

Example:

```
-- Create a table.
openGauss=# CREATE TABLE blob_type_t1
(
  BT_COL1 INTEGER,
  BT_COL2 BLOB,
  BT_COL3 RAW,
  BT_COL4 BYTEA
);

-- Insert data.
openGauss=# INSERT INTO blob_type_t1 VALUES(10,empty_blob(),
HEXTORAW('DEADBEEF'),E'\xDEADBEEF');

-- Query data in the table.
openGauss=# SELECT * FROM blob_type_t1;
 bt_col1 | bt_col2 | bt_col3 | bt_col4
-----+-----+-----+-----
    10 |      | DEADBEEF | \xdeadbeef
(1 row)
```

```
-- Delete the table.  
openGauss=# DROP TABLE blob_type_t1;
```

7.3.6 Date/Time Types

Table 7-11 lists the date/time types supported by GaussDB. For the operators and built-in functions of the types, see [Date and Time Processing Functions and Operators](#).

NOTE

If the time format of another database is different from that of GaussDB, modify the value of the **DateStyle** parameter to keep them consistent.

Table 7-11 Date/Time types

Name	Description	Storage Space
DATE	Date. Minimum value: 4713-01-01BC (4713 B.C.). Maximum value: 5874897-12-31AD (5874897 A.C.) NOTE For A compatibility, the database treats empty strings as NULL and replaces DATE with TIMESTAMP(0) WITHOUT TIME ZONE .	4 bytes (8 bytes in A-compatible schema)
TIME [(p)] [WITHOUT TIME ZONE]	Time within one day (without time zone). p indicates the precision after the decimal point. The value ranges from 0 to 6 . Minimum value: 00:00:00 . Maximum value: 24:00:00 .	8 bytes
TIME [(p)] [WITH TIME ZONE]	Time within one day (with time zone). p indicates the precision after the decimal point. The value ranges from 0 to 6 . Minimum value: 00:00:00+1559 . Maximum value: 24:00:00 .	12 bytes
TIMESTAMP[(p)] [WITHOUT TIME ZONE]	Date and time (without a time zone). p indicates the precision after the decimal point. The value ranges from 0 to 6 . Minimum value: 4713-11-24BC 00:00:00.000000 (4713 B.C.). Maximum value: 294277-01-09AD 00:00:00.000000 (294277 A.D.).	8 bytes

Name	Description	Storage Space
TIMESTAMP[(p)] [WITH TIME ZONE]	Date and time (with time zone). TIMESTAMP is also called TIMESTAMPTZ. p indicates the precision after the decimal point. The value ranges from 0 to 6. Minimum value: 4713-11-24BC 00:00:00.000000 (4713 B.C.). Maximum value: 294277-01-09AD 00:00:00.000000 (294277 A.D.). Time zone update: In some countries or regions, the time zone information is frequently updated due to political, economic, and war factors. Therefore, the database system often needs to modify the time zone file accordingly to ensure that the time is correct. Currently, the GaussDB time zone type involves only timestamp with timezone. When a new time zone file takes effect, the existing data is not changed, and the new data is adjusted based on the time zone file information. The capability of the same type of data in the database is different from that in an A-compatible database.	8 bytes
SMALLDATETIME	Date and time (without time zone). The precision level is minute. 31s to 59s are rounded into 1 minute. Minimum value: 4713-11-24BC 00:00:00.000000 (4713 B.C.). Maximum value: 294277-01-09AD 00:00:00.000000 (294277 A.D.).	8 bytes
INTERVAL DAY (l) TO SECOND (p)	Time interval (X days X hours X minutes X seconds). • l: indicates the precision of days. The value ranges from 0 to 6. For compatibility, the precision functions are not supported. • p: indicates the precision of seconds. The value ranges from 0 to 6. The digit 0 at the end of a decimal number is not displayed.	16 bytes

Name	Description	Storage Space
INTERVAL [FIELDS] [(p)]	Time interval. • FIELDS: YEAR, MONTH, DAY, HOUR, MINUTE, SECOND, DAY TO HOUR, DAY TO MINUTE, DAY TO SECOND, HOUR TO MINUTE, HOUR TO SECOND, or MINUTE TO SECOND. • p: indicates the precision of seconds. The value ranges from 0 to 6. p takes effect only when FIELDS is SECOND, DAY TO SECOND, HOUR TO SECOND, or MINUTE TO SECOND. The digit 0 at the end of a decimal number is not displayed.	12 bytes
reltime	Relative time interval. • The format is as follows: X years X mons X days XX:XX:XX. • The Julian calendar is used. It specifies that a year has 365.25 days and a month has 30 days. The relative time interval needs to be calculated based on the input value.	4 bytes
abstime	Date and time. • Format: YYYY-MM-DD hh:mm:ss +timezone. • The value range is from 1901-12-13 20:45:53 GMT to 2038-01-18 23:59:59 GMT. The precision is second.	4 bytes

Example:

```
-- Create a table.
openGauss=# CREATE TABLE date_type_tab(coll date);

-- Insert data.
openGauss=# INSERT INTO date_type_tab VALUES (date '12-10-2010');

-- View data.
openGauss=# SELECT * FROM date_type_tab;
      coll
-----
2010-12-10 00:00:00
(1 row)

-- Delete the table.
openGauss=# DROP TABLE date_type_tab;
```

```
-- Create a table.
openGauss=# CREATE TABLE time_type_tab (da time without time zone ,dai time with time zone,dfgh
timestamp without time zone,dfga timestamp with time zone, vbg smalldatetime);

-- Insert data.
openGauss=# INSERT INTO time_type_tab VALUES ('21:21:21','21:21:21 pst','2010-12-12','2013-12-11
pst','2003-04-12 04:05:06');

-- View data.
openGauss=# SELECT * FROM time_type_tab;
   da   |   dai   |   dfgh   |   dfga   |   vbg
-----+-----+-----+-----+-----
21:21:21 | 21:21:21-08 | 2010-12-12 00:00:00 | 2013-12-11 16:00:00+08 | 2003-04-12 04:05:00
(1 row)

-- Delete the table.
openGauss=# DROP TABLE time_type_tab;

-- Create a table.
openGauss=# CREATE TABLE day_type_tab (a int,b INTERVAL DAY(3) TO SECOND (4));

-- Insert data.
openGauss=# INSERT INTO day_type_tab VALUES (1, INTERVAL '3' DAY);

-- View data.
openGauss=# SELECT * FROM day_type_tab;
 a | b
---+-----
 1 | 3 days
(1 row)

-- Delete the table.
openGauss=# DROP TABLE day_type_tab;

-- Create a table.
openGauss=# CREATE TABLE year_type_tab(a int, b interval year (6));

-- Insert data.
openGauss=# INSERT INTO year_type_tab VALUES(1,interval '2' year);

-- View data.
openGauss=# SELECT * FROM year_type_tab;
 a | b
---+-----
 1 | 2 years
(1 row)

-- Delete the table.
openGauss=# DROP TABLE year_type_tab;
```

Date Inputs

Date and time input is accepted in almost any reasonable formats, including ISO 8601, SQL-compatible, and traditional Postgres. The system allows you to customize the sequence of day, month, and year in the date input. Set the **DateStyle** parameter to **MDY** to select month-day-year interpretation, **DMY** to select day-month-year interpretation, or **YMD** to select year-month-day interpretation.

Remember that any date or time literal input needs to be enclosed with single quotation marks ('), and the syntax is as follows:

```
type [ ( p ) ] 'value'
```

The **p** that can be selected in the precision statement is an integer, indicating the number of fractional digits in the **seconds** column. [Table 7-12](#) shows the input formats of the date type.

Table 7-12 Date input

Example	Description
1999-01-08	ISO 8601 (recommended format). January 8, 1999 in any format.
January 8, 1999	Unambiguous in any datestyle input mode
1/8/1999	January 8 in MDY mode. August 1 in DMY mode
1/18/1999	January 18 in MDY mode, rejected in other modes
01/02/03	- January 2, 2003 in MDY mode - February 1, 2003 in DMY mode - February 3, 2001 in YMD mode
1999-Jan-08	January 8 in any mode
Jan-08-1999	January 8 in any mode
08-Jan-1999	January 8 in any mode
99-Jan-08	January 8 in YMD mode, else error
08-Jan-99	January 8, except error in YMD mode
Jan-08-99	January 8, except error in YMD mode
19990108	ISO 8601: January 8, 1999 in any mode
990108	ISO 8601: January 8, 1999 in any mode
1999.008	Year and day of year
J2451187	Julian date
January 8, 99 BC	Year 99 BC

Example:

```
-- Create a table.
openGauss=# CREATE TABLE date_type_tab(coll date);

-- Insert data.
openGauss=# INSERT INTO date_type_tab VALUES (date '12-10-2010');

-- View data.
openGauss=# SELECT * FROM date_type_tab;
      coll
-----
2010-12-10(1 row)

-- View the date format.
openGauss=# SHOW datestyle;
```



```
DateStyle
-----
ISO, MDY
(1 row)

-- Set the date format.
openGauss=# SET datestyle='YMD';
SET

-- Insert data.
openGauss=# INSERT INTO date_type_tab VALUES(date '2010-12-11');

-- View data.
openGauss=# SELECT * FROM date_type_tab;
      coll
-----
2010-12-10 00:00:00
2010-12-11 00:00:00
(2 rows)

-- Delete the table.
openGauss=# DROP TABLE date_type_tab;
```

Time

The time-of-day types are **TIME [(p)] [WITHOUT TIME ZONE]** and **TIME [(p)] [WITH TIME ZONE]**. **TIME** alone is equivalent to **TIME WITHOUT TIME ZONE**.

If a time zone is specified in the input for **TIME WITHOUT TIME ZONE**, it is silently ignored.

For details about the time input types, see [Table 7-13](#). For details about time zone input types, see [Table 7-14](#).

Table 7-13 Time input

Example	Description
05:06.8	ISO 8601
4:05:06	ISO 8601
4:05	ISO 8601
040506	ISO 8601
4:05 AM	Same as 04:05. Input hours must be less than or equal to 12.
4:05 PM	Same as 16:05. Input hour must be <= 12
04:05:06.789-8	ISO 8601
04:05:06-08:00	ISO 8601
04:05-08:00	ISO 8601
040506-08	ISO 8601
04:05:06 PST	Time zone specified by abbreviation

Example	Description
2003-04-12 04:05:06 America/ New_York	Time zone specified by full name

Table 7-14 Time zone input

Example	Description
PST	Abbreviation (for Pacific Standard Time)
America/New_York	Full time zone name
-8:00	ISO-8601 offset for PST
-800	ISO-8601 offset for PST
-8	ISO-8601 offset for PST

Example:

```
openGauss=# SELECT time '04:05:06';
time
-----
04:05:06
(1 row)

openGauss=# SELECT time '04:05:06 PST';
time
-----
04:05:06
(1 row)

openGauss=# SELECT time with time zone '04:05:06 PST';
timetz
-----
04:05:06-08
(1 row)
```

Special Values

The special values supported by GaussDB are converted to common date/time values when being read. For details, see [Table 7-15](#).

Table 7-15 Special values

Input String	Applicable Type	Description
epoch	date and timestamp	1970-01-01 00:00:00+00 (Unix system time zero)
infinity	timestamp	Infinity, later than any other timestamp
-infinity	timestamp	Infinity, earlier than any other timestamp

Input String	Applicable Type	Description
now	date, time, and timestamp	Start time of the current transaction
today	date and timestamp	Today at midnight
tomorrow	date and timestamp	Tomorrow at midnight
yesterday	date and timestamp	Yesterday at midnight
allballs	time	00:00:00.00 UTC

Example:

```
-- Create a table.
openGauss=# CREATE TABLE realtime_type_special(col1 varchar(20), col2 date, col3 timestamp, col4 time);

-- Insert data.
openGauss=# INSERT INTO realtime_type_special VALUES('epoch', 'epoch', 'epoch', NULL);
openGauss=# INSERT INTO realtime_type_special VALUES('now', 'now', 'now', 'now');
openGauss=# INSERT INTO realtime_type_special VALUES('today', 'today', 'today', NULL);
openGauss=# INSERT INTO realtime_type_special VALUES('tomorrow', 'tomorrow', 'tomorrow', NULL);
openGauss=# INSERT INTO realtime_type_special VALUES('yesterday', 'yesterday', 'yesterday', NULL);

-- View data.
openGauss=# SELECT * FROM realtime_type_special;
 col1 |   col2   |   col3   |   col4
-----+-----+-----+-----
epoch | 1970-01-01 00:00:00 | 1970-01-01 00:00:00 |
now   | 2023-02-27 11:38:13 | 2023-02-27 11:38:13.032815 | 11:38:13.032815
today | 2023-02-27 00:00:00 | 2023-02-27 00:00:00 |
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
yesterday | 2023-02-26 00:00:00 | 2023-02-26 00:00:00 |
(5 rows)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 < 'infinity';
 col1 |   col2   |   col3   |   col4
-----+-----+-----+-----
epoch | 1970-01-01 00:00:00 | 1970-01-01 00:00:00 |
now   | 2023-02-27 11:38:13 | 2023-02-27 11:38:13.032815 | 11:38:13.032815
today | 2023-02-27 00:00:00 | 2023-02-27 00:00:00 |
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
yesterday | 2023-02-26 00:00:00 | 2023-02-26 00:00:00 |
(5 rows)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 > '-infinity';
 col1 |   col2   |   col3   |   col4
-----+-----+-----+-----
epoch | 1970-01-01 00:00:00 | 1970-01-01 00:00:00 |
now   | 2023-02-27 11:38:13 | 2023-02-27 11:38:13.032815 | 11:38:13.032815
today | 2023-02-27 00:00:00 | 2023-02-27 00:00:00 |
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
yesterday | 2023-02-26 00:00:00 | 2023-02-26 00:00:00 |
(5 rows)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 > 'now';
 col1 |   col2   |   col3   |   col4
-----+-----+-----+-----
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
(1 row)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 = 'today';
```

```
col1 | col2 | col3 | col4
-----+-----+-----+-----
today | 2023-02-27 00:00:00 | 2023-02-27 00:00:00 |
(1 row)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 = 'tomorrow';
col1 | col2 | col3 | col4
-----+-----+-----+-----
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
(1 row)

openGauss=# SELECT * FROM realtime_type_special WHERE col3 > 'yesterday';
col1 | col2 | col3 | col4
-----+-----+-----+-----
now | 2023-02-27 11:38:13 | 2023-02-27 11:38:13.032815 | 11:38:13.032815
today | 2023-02-27 00:00:00 | 2023-02-27 00:00:00 |
tomorrow | 2023-02-28 00:00:00 | 2023-02-28 00:00:00 |
(3 rows)

-- Delete the table.
openGauss=# DROP TABLE realtime_type_special;
```

Interval Input

The input of **reltime** can be any valid interval in text format. It can be a number (negative numbers and decimals are also allowed) or a specific time, which must be in SQL standard format, ISO-8601 format, or POSTGRES format. In addition, the text input needs to be enclosed with single quotation marks (").

For details about interval input, see [Table 7-16](#).

Table 7-16 Interval input types

Input	Output	Description
60	2 mons	Numbers are used to indicate intervals. The default unit is day. Decimals and negative numbers are allowed. Particularly, a negative interval syntactically means how long before.
31.25	1 mons 1 days 06:00:00	
-365	-12 mons -5 days	
1 years 1 mons 8 days 12:00:00	1 years 1 mons 8 days 12:00:00	Intervals are in POSTGRES format. They can contain both positive and negative numbers and are case-insensitive. Output is a simplified POSTGRES interval converted from the input.
-13 months -10 hours	-1 years -25 days -04:00:00	
-2 YEARS +5 MONTHS 10 DAYS	-1 years -6 mons -25 days -06:00:00	
P-1.1Y10M	-3 mons -5 days -06:00:00	Intervals are in ISO-8601 format. They can contain both positive and negative numbers and are case-insensitive. Output is a simplified POSTGRES interval converted from the input.
-12H	-12:00:00	

Example:

```
-- Create a table.
openGauss=# CREATE TABLE reltime_type_tab(col1 character(30), col2 reltime);

-- Insert data.
openGauss=# INSERT INTO reltime_type_tab VALUES ('90', '90');
openGauss=# INSERT INTO reltime_type_tab VALUES ('-366', '-366');
openGauss=# INSERT INTO reltime_type_tab VALUES ('1975.25', '1975.25');
openGauss=# INSERT INTO reltime_type_tab VALUES ('-2 YEARS +5 MONTHS 10 DAYS', '-2 YEARS +5 MONTHS 10 DAYS');
openGauss=# INSERT INTO reltime_type_tab VALUES ('30 DAYS 12:00:00', '30 DAYS 12:00:00');
openGauss=# INSERT INTO reltime_type_tab VALUES ('P-1.1Y10M', 'P-1.1Y10M');

-- View data.
openGauss=# SELECT * FROM reltime_type_tab;
      col1      |      col2
-----+-----
90              | 3 mons
-366            | -1 years -18:00:00
1975.25         | 5 years 4 mons 29 days
-2 YEARS +5 MONTHS 10 DAYS | -1 years -6 mons -25 days -06:00:00
30 DAYS 12:00:00 | 1 mon 12:00:00
P-1.1Y10M       | -3 mons -5 days -06:00:00
(6 rows)

-- Delete the table.
openGauss=# DROP TABLE reltime_type_tab;
```

7.3.7 Geometric

Table 7-17 lists the geometric types that can be used in GaussDB. The most fundamental type, the point, forms the basis for all of the other types.

Table 7-17 Geometric types

Name	Storage Space	Description	Representation
point	16 bytes	Point on a plane	(x,y)
lseg	32 bytes	Finite line segment	((x1,y1),(x2,y2))
box	32 bytes	Rectangle	((x1,y1),(x2,y2))
path	16 + 16 x <i>n</i> bytes	Closed path (similar to polygon)	((x1,y1),...)
path	16 + 16 x <i>n</i> bytes	Open path	[(x1,y1),...]
polygon	40 + 16 x <i>n</i> bytes	Polygon (similar to closed path)	((x1,y1),...)
circle	24 bytes	Circle	<(x,y),r> (center point and radius)

A rich set of functions and operators is available in GaussDB to perform various geometric operations, such as scaling, translation, rotation, and determining intersections. For details, see [Geometric Functions and Operators](#).

Points

Points are the fundamental two-dimensional building block for geometric types. Values of the **point** type are specified using either of the following syntax:

```
( x , y )  
x , y
```

x and **y** are the respective coordinates, as floating-point numbers. The value type of the points is float8.

Points are output using the first syntax.

Example:

```
openGauss=# select point(1.1, 2.2);  
point  
-----  
(1.1,2.2)  
(1 row)
```

Line Segments

Line segments (**lseg**) are represented by pairs of points. Values of the **lseg** type are specified using any of the following syntax:

```
[ ( x1 , y1 ) , ( x2 , y2 ) ]  
( ( x1 , y1 ) , ( x2 , y2 ) )  
( x1 , y1 ) , ( x2 , y2 )  
x1 , y1 , x2 , y2
```

(x1,y1) and **(x2,y2)** are the end points of the line segment. The value type of the points is float8.

Line segments are output using the first syntax.

Example:

```
openGauss=# select lseg(point(1.1, 2.2), point(3.3, 4.4));  
lseg  
-----  
[(1.1,2.2),(3.3,4.4)]  
(1 row)
```

Rectangles

Boxes are represented by pairs of points that are opposite corners of the box. Values of the **box** type are specified using any of the following syntax:

```
( ( x1 , y1 ) , ( x2 , y2 ) )  
( x1 , y1 ) , ( x2 , y2 )  
x1 , y1 , x2 , y2
```

(x1,y1) and **(x2,y2)** are any two opposite corners of the rectangle. The value type of the points is float8.

Rectangles are output using the second syntax.

Any two opposite corners can be supplied on input, but in this order, the values will be reordered as needed to store the upper right and lower left corners.

Example:

```
openGauss=# select box(point(1.1, 2.2), point(3.3, 4.4));
      box
-----
(3.3,4.4),(1.1,2.2)
(1 row)
```

Paths

Paths are represented by lists of connected points. Paths can be open, where the first and last points in the list are considered not connected, or closed, where the first and last points are considered connected.

Values of the **path** type are specified using any of the following syntax:

```
[ ( x1 , y1 ) , ... , ( xn , yn ) ]
( ( x1 , y1 ) , ... , ( xn , yn ) )
( x1 , y1 ) , ... , ( xn , yn )
( x1 , y1 , ... , xn , yn )
x1 , y1 , ... , xn , yn
```

The points are the end points of the line segments comprising the path. The value type of the points is float8. Square brackets ([]) indicate an open path, while parentheses (()) indicate a closed path. When the outermost parentheses are omitted, as in the third through fifth syntax, a closed path is assumed.

Paths are output using the first or second syntax.

Example:

```
openGauss=# select path(polygon '((0,0),(1,1),(2,0))');
      path
-----
((0,0),(1,1),(2,0))
(1 row)
```

Polygons

Polygons are represented by lists of points (the vertexes of the polygon). Polygons are very similar to closed paths, but are stored differently and have their own set of support functions.

Values of the **polygon** type are specified using any of the following syntax:

```
(( x1 , y1 ) , ... , ( xn , yn ) )
( x1 , y1 ) , ... , ( xn , yn )
( x1 , y1 , ... , xn , yn )
x1 , y1 , ... , xn , yn
```

A point indicates the vertex of a polygon. The value type of a point is float8.

Polygons are output using the first syntax.

Example:

```
openGauss=# select polygon(box '((0,0),(1,1))');
      polygon
-----
((0,0),(0,1),(1,1),(1,0))
(1 row)
```

Circles

Circles are represented by a center point and radius. Values of the **circle** type are specified using the following syntaxes:

```
< ( x , y ) , r >  
( ( x , y ) , r )  
( x , y ) , r  
x , y , r
```

(x,y) is the center point and **r** is the radius of the circle. The value type of the points is float8.

Circles are output using the first syntax.

Example:

```
openGauss=# select circle(point(0,0),1);  
circle  
-----  
<(0,0),1>  
(1 row)
```

7.3.8 Network Address Types

GaussDB offers data types to store IPv4 and MAC addresses.

It is better to use these types instead of plain-text types to store network addresses, because these types offer input error checking and specialized operators and functions (see [Network Address Functions and Operators](#)).

Table 7-18 Network address types

Name	Storage Space	Description
cidr	7 bytes	IPv4 networks
inet	7 bytes	IPv4 hosts and networks
macaddr	6 bytes	MAC address

cidr

The **cidr** type (Classless Inter-Domain Routing) holds an IPv4 network address. The format for specifying networks is **address/y** where **address** is the network represented as an IPv4 address, and **y** is the number of bits in the netmask. If **y** is omitted, it is calculated using assumptions from the older classful network numbering system, except it will be at least large enough to include all of the octets written in the input.

Table 7-19 cidr type input examples

cidr Input	cidr Output	abbrev(cidr)
192.168.100.128/25	192.168.100.128/25	192.168.100.128/25
192.168/24	192.168.0.0/24	192.168.0/24

cidr Input	cidr Output	abbrev(cidr)
192.168/25	192.168.0.0/25	192.168.0.0/25
192.168.1	192.168.1.0/24	192.168.1/24
192.168	192.168.0.0/24	192.168.0/24
10.1.2	10.1.2.0/24	10.1.2/24
10.1	10.1.0.0/16	10.1/16
10	10.0.0.0/8	10/8
10.1.2.3/32	10.1.2.3/32	10.1.2.3/32

Example:

```
openGauss=# CREATE TABLE cidr_test(id int, c cidr);
CREATE TABLE
openGauss=# INSERT INTO cidr_test VALUES (1, '192.168.100.128/25');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (2, '192.168/24');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (3, '192.168/25');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (4, '192.168.1');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (5, '192.168');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (6, '10.1.2');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (7, '10.1');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (8, '10');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (9, '2001:4f8:3:ba::/64');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (10, '2001:4f8:3:ba:2e0:81ff:fe22:d1f1/128');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (11, '::ffff:127.0.0.0/120');
INSERT 0 1
openGauss=# INSERT INTO cidr_test VALUES (12, '::ffff:127.0.0.0/128');
INSERT 0 1
openGauss=# SELECT * FROM cidr_test ORDER BY id;
 id |          c          |
-----+-----
  1 | 192.168.100.128/25 |
  2 | 192.168.0.0/24    |
  3 | 192.168.0.0/25    |
  4 | 192.168.1.0/24    |
  5 | 192.168.0.0/24    |
  6 | 10.1.2.0/24       |
  7 | 10.1.0.0/16       |
  8 | 10.0.0.0/8        |
  9 | 2001:4f8:3:ba::/64 |
 10 | 2001:4f8:3:ba:2e0:81ff:fe22:d1f1/128 |
 11 | ::ffff:127.0.0.0/120 |
 12 | ::ffff:127.0.0.0/128 |
(12 rows)

openGauss=# DROP TABLE cidr_test;
DROP TABLE
```

inet

The **inet** type holds an IPv4 host address, and optionally its subnet, all in one field. The subnet is represented by the number of network address bits present in the host address (the "netmask"). If the netmask is 32 and the address is an IPv4 address, then the value does not indicate a subnet, only a single host.

The input format for this type is **address/y** where **address** is an IPv4 address and **y** is the number of bits in the netmask. If **/y** is omitted, the subnet mask is **32** for an IPv4 address, and the value represents just a single host. On display, the **/y** portion is suppressed if the netmask specifies a single host.

The essential difference between the **inet** and **cidr** data types is that **inet** accepts values with nonzero bits to the right of the netmask, whereas **cidr** does not.

Example:

```
openGauss=# CREATE TABLE inet_test(id int, i inet);
CREATE TABLE
openGauss=# INSERT INTO inet_test VALUES (1, '192.168.100.128/25');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (2, '192.168.100.128');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (3, '192.168.1.0/24');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (4, '192.168.1.0/25');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (5, '192.168.1.255/24');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (6, '192.168.1.255/25');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (7, '10.1.2.3/8');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (8, '11.1.2.3/16');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (9, '12.1.2.3/24');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (10, '13.1.2.3/32');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (11, '2001:4f8:3:ba::/64');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (12, '2001:4f8:3:ba:2e0:81ff:fe22:d1f1/128');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (13, '::ffff:127.0.0.0/120');
INSERT 0 1
openGauss=# INSERT INTO inet_test VALUES (14, '::ffff:127.0.0.0/128');
INSERT 0 1
openGauss=# SELECT * FROM inet_test ORDER BY id;
 id |          i
-----+-----
  1 | 192.168.100.128/25
  2 | 192.168.100.128
  3 | 192.168.1.0/24
  4 | 192.168.1.0/25
  5 | 192.168.1.255/24
  6 | 192.168.1.255/25
  7 | 10.1.2.3/8
  8 | 11.1.2.3/16
  9 | 12.1.2.3/24
 10 | 13.1.2.3
 11 | 2001:4f8:3:ba::/64
 12 | 2001:4f8:3:ba:2e0:81ff:fe22:d1f1
 13 | ::ffff:127.0.0.0/120
 14 | ::ffff:127.0.0.0
(14 rows)
```

```
openGauss=# DROP TABLE inet_test;
DROP TABLE
```

macaddr

The **macaddr** type stores MAC addresses, known for example from Ethernet card hardware addresses (although MAC addresses are used for other purposes as well). Input is accepted in the following formats:

```
'08:00:2b:01:02:03'
'08-00-2b-01-02-03'
'08002b:010203'
'08002b-010203'
'0800.2b01.0203'
'08002b010203'
```

These examples would all specify the same address. Upper and lower cases are accepted for the digits a through f. Output is always in the first of the forms shown.

Example:

```
openGauss=# CREATE TABLE macaddr_test(id int, m macaddr);
CREATE TABLE
openGauss=# INSERT INTO macaddr_test VALUES (1, '08:00:2b:01:02:03');
INSERT 0 1
openGauss=# INSERT INTO macaddr_test VALUES (2, '08-00-2b-01-02-03');
INSERT 0 1
openGauss=# INSERT INTO macaddr_test VALUES (3, '08002b:010203');
INSERT 0 1
openGauss=# INSERT INTO macaddr_test VALUES (4, '08002b-010203');
INSERT 0 1
openGauss=# INSERT INTO macaddr_test VALUES (5, '0800.2b01.0203');
INSERT 0 1
openGauss=# INSERT INTO macaddr_test VALUES (6, '08002b010203');
INSERT 0 1
openGauss=# SELECT * FROM macaddr_test ORDER BY id;
 id |      m
-----+-----
  1 | 08:00:2b:01:02:03
  2 | 08:00:2b:01:02:03
  3 | 08:00:2b:01:02:03
  4 | 08:00:2b:01:02:03
  5 | 08:00:2b:01:02:03
  6 | 08:00:2b:01:02:03
(6 rows)

openGauss=# DROP TABLE macaddr_test;
DROP TABLE
```

7.3.9 Bit String Types

Bit strings are strings of 1's and 0's. They can be used to store bit masks.

GaussDB supports two bit string types: **bit(n)** and **bit varying(n)**. Here, *n* is a positive integer. The maximum value of *n* is **83886080**, which is equivalent to 10 MB.

The **bit** type data must match the length *n* exactly. It is an error to attempt to store shorter or longer bit strings. Data of the bit varying type is of variable length up to the maximum length *n*. Longer strings will be rejected. Writing **bit** without a length is equivalent to **bit(1)**, while **bit varying** without a length specification means unlimited length.

 NOTE

- If one explicitly casts a bit-string value to **bit(*n*)**, it will be truncated or zero-padded on the right to be exactly *n* bits, without raising an error.
- Similarly, if one explicitly casts a bit-string value to **bit varying(*n*)**, it will be truncated on the right if it is more than *n* bits.
- When the ADMS platform driver version 8.1.3-200 or earlier is used, use **::bit varying** to convert the bit type. Otherwise, an error may occur.

```
-- Create a table.
openGauss=# CREATE TABLE bit_type_t1
(
    BT_COL1 INTEGER,
    BT_COL2 BIT(3),
    BT_COL3 BIT VARYING(5)
);

-- Insert data.
openGauss=# INSERT INTO bit_type_t1 VALUES(1, B'101', B'00');

-- Specify the type length. An error is reported if an inserted string exceeds this length.
openGauss=# INSERT INTO bit_type_t1 VALUES(2, B'10', B'101');
ERROR: bit string length 2 does not match type bit(3)
CONTEXT: referenced column: bt_col2

-- Specify the type length. Data is converted if it exceeds this length.
openGauss=# INSERT INTO bit_type_t1 VALUES(2, B'10'::bit(3), B'101');

-- View data.
openGauss=# SELECT * FROM bit_type_t1;
 bt_col1 | bt_col2 | bt_col3
-----+-----+-----
      1 | 101    | 00
      2 | 100    | 101
(2 rows)
```

```
-- Delete the table.
openGauss=# DROP TABLE bit_type_t1;
```

7.3.10 UUID

The data type UUID stores universally unique identifiers (UUID) as defined by RFC 4122, ISO/IEF 9834-8:2005, and related standards. This identifier is a 128-bit quantity that is generated by an algorithm chosen to make it very unlikely that the same identifier will be generated by anyone else in the known universe using the same algorithm.

A UUID is written as a sequence of lower-case hexadecimal digits, in several groups separated by hyphens, specifically a group of 8 digits followed by three groups of 4 digits followed by a group of 12 digits, for a total of 32 digits representing the 128 bits. An example of a UUID in this standard form is:

```
a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11
```

GaussDB also accepts the following alternative forms for input: use of upper-case letters and digits, the standard format surrounded by braces, omitting some or all hyphens, adding a hyphen after any group of four digits. An example is provided as follows:

```
A0EEBC99-9C0B-4EF8-BB6D-6BB9BD380A11
{a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11}
a0eebc999c0b4ef8bb6d6bb9bd380a11
a0ee-bc99-9c0b-4ef8-bb6d-6bb9-bd38-0a11
```

Output is always in the standard form.

Examples

```
-- Create a table.
openGauss=# CREATE TABLE uuid_test(id int, test uuid);

-- Insert data in the example format.
openGauss=# INSERT INTO uuid_test VALUES(1, 'A0EEBC99-9C0B-4EF8-BB6D-6BB9BD380A11');
openGauss=# INSERT INTO uuid_test VALUES(2, '{a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11}');
openGauss=# INSERT INTO uuid_test VALUES(3, 'a0eebc999c0b4ef8bb6d6bb9bd380a11');
openGauss=# INSERT INTO uuid_test VALUES(4, 'a0ee-bc99-9c0b-4ef8-bb6d-6bb9-bd38-0a11');

-- View data and output the data in the standard format.
openGauss=# SELECT * FROM uuid_test;
id |          test
----+-----
 1 | a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11
 2 | a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11
 3 | a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11
 4 | a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a11
(4 rows)
```

7.3.11 JSON/JSONB Types

JavaScript Object Notation (JSON) data can be a single scalar, an array, or a key-value pair object. The array and object can be called a container:

- Scalar: a number, Boolean, string, or null
- Array: defined in a pair of square brackets ([]), in which elements can be any type of JSON data, and are not necessarily of the same type.
- Object: defined in a pair of braces ({}), in which objects are stored in the format of key:value. Each key must be a string enclosed by a pair of double quotation marks (""), and its value can be any type of JSON data. In case of duplicate keys, the last key value will be used.

GaussDB offers two types for storing JSON data: JSON and JSONB. The JSON data type stores a complete copy of the input, retaining the entered spaces, duplicate keys, and sequence, while JSONB data is stored in its decomposed binary form, with semantic-irrelevant details and duplicate keys removed and key-values sorted. Therefore, the JSONB data does not need to be parsed.

Basically, both are JSON data types. However, they differ greatly in efficiency. Because JSON data type stores an exact copy of the input text, the data must be parsed on every execution. In contrast, JSONB data is stored in its parsed binary form and can be processed faster, though this makes it slightly slower to input due to the conversion mechanism. In addition, because the JSONB data type is normalized, it supports more operations, for example, comparing sizes according to a specific rule. JSONB also supports indexing, which is a significant advantage.

Input Format

An input must be a JSON-compliant string, which is enclosed in single quotation marks (').

null (null-json): The value can only be **null** in lowercase.

```
openGauss=# SELECT 'null':json; -- suc
json
-----
null
(1 row)
```

```
openGauss=# SELECT 'NULL'::jsonb; -- err
ERROR:  invalid input syntax for type json
```

Number (num-json): The value can be a positive or negative integer, decimal fraction, or 0. The scientific notation is supported.

```
openGauss=# SELECT '1'::json;
json
```

```
-----
1
(1 row)
```

```
openGauss=# SELECT '-1.5'::json;
json
```

```
-----
-1.5
(1 row)
```

```
openGauss=# SELECT '-1.5e-5'::jsonb, '-1.5e+2'::jsonb;
```

```
jsonb | jsonb
-----+-----
-0.000015 | -150
(1 row)
```

```
openGauss=# SELECT '001'::json, '+15'::json, 'NaN'::json; -- Redundant leading zeros, plus signs (+), NaN,
and infinity are not supported.
ERROR:  invalid input syntax for type json
```

Boolean (bool-json): The value can only be **true** or **false** in lowercase.

```
openGauss=# SELECT 'true'::json;
json
```

```
-----
true
(1 row)
```

```
openGauss=# SELECT 'false'::jsonb;
jsonb
```

```
-----
false
(1 row)
```

String (str-json): The value must be a string enclosed in double quotation marks ("").

```
openGauss=# SELECT '"a"'::json;
json
```

```
-----
"a"
(1 row)
```

```
openGauss=# SELECT '"abc"'::jsonb;
jsonb
```

```
-----
"abc"
(1 row)
```

Array (array-json): Arrays are enclosed in square brackets ([]). Elements in the array can be any valid JSON data, and are unnecessarily of the same type.

```
openGauss=# SELECT '[1, 2, "foo", null]'::json;
json
```

```
-----
[1, 2, "foo", null]
(1 row)
```

```
openGauss=# SELECT '[]'::json;
json
```

```
-----
```

```
[]  
(1 row)  
  
openGauss=# SELECT '[1, 2, "foo", null, [], {}]'::jsonb;  
          jsonb  
-----  
[1, 2, "foo", null, [], {}]  
(1 row)
```

Object (object-json): The value is enclosed in braces ({}). The key must be a JSON-compliant string, and the value can be any valid JSON string.

```
openGauss=# SELECT '{}'::json;  
          json  
-----  
{}  
(1 row)  
  
openGauss=# SELECT '{"a": 1, "b": {"a": 2, "b": null}}'::json;  
          json  
-----  
{"a": 1, "b": {"a": 2, "b": null}}  
(1 row)  
  
openGauss=# SELECT '{"foo": [true, "bar"], "tags": {"a": 1, "b": null}}'::jsonb;  
          jsonb  
-----  
{"foo": [true, "bar"], "tags": {"a": 1, "b": null}}  
(1 row)
```

NOTICE

- Note that **'null'::json** and **null::json** are different, which are similar to the strings **str=""** and **str=null**.
- For numbers, when scientific notation is used, JSONB expands them, while JSON stores an exact copy of the input text.

JSONB Advanced Features

- Precautions
 - It cannot be used as a partition key.
 - Foreign tables and MOTs are not supported.

The main difference between JSON and JSONB lies in the storage mode. JSONB stores parsed binary data, which reflects the JSON hierarchy and facilitates direct access. Therefore, JSONB has many advanced features that JSON does not have.

- Format normalization
 - After the input object-json string is parsed into JSONB binary, semantically irrelevant details are naturally discarded, for example, spaces:

```
openGauss=# select ' [1, " a ", {"a" :1  } ] '::jsonb;  
          jsonb  
-----  
[1, " a ", {"a": 1}]  
(1 row)
```
 - For object-json, duplicate key-values are deleted and only the last key-value is retained. For example:

```
openGauss=# select '{"a" : 1, "a" : 2}'::jsonb;  
          jsonb
```

```
-----  
{ "a": 2 }  
(1 row)
```

- For object-jsonb, key-values will be re-sorted. The collation rule is as follows: 1. Longer key-values are sorted last. 2. If the key-values are of the same length, the key-values with a larger ASCII code are sorted after the key-values with a smaller ASCII code:

```
openGauss=# select '{"aa" : 1, "b" : 2, "a" : 3} '::jsonb;  
              jsonb  
-----  
{ "a": 3, "b": 2, "aa": 1 }  
(1 row)
```

- Size comparison

Format normalization ensures that only one form of JSONB data exists in the same semantics. Therefore, sizes may be compared according to a specific rule.

- First, type comparison: **object-jsonb** > **array-jsonb** > **bool-jsonb** > **num-jsonb** > **str-jsonb** > **null-jsonb**
- Content comparison if the data type is the same:
 - str-jsonb: The default text sorting rule of the database is used for comparison. A positive value indicates greater than, a negative value indicates less than, and 0 indicates equal.
 - num-jsonb: numeric comparison
 - bool-jsonb: **true** > **false**
 - array-jsonb: long elements > short elements. If the lengths are the same, compare each element in sequence.
 - object-jsonb: If the length of a key-value pair is longer than that of a short key-value pair, the keys are compared first, and then the values are compared.

NOTICE

For comparison within the object-jsonb type, the final result after format sorting is used for comparison. Therefore, the comparison result may not be intuitive compared with the direct input.

- Creating indexes, primary keys, and foreign keys

- B-tree index

B-tree indexes, primary keys, and foreign keys can be created for the **JSONB** type.

- Inclusion and existence

Querying whether a JSON contains some elements or whether some elements exist in a JSON is an important capability of JSONB.

For details about the operators, see [JSON/JSONB Functions and Operators](#).

- Functions and operators

For details about the functions and operators supported by the JSON/JSONB type, see [JSON/JSONB Functions and Operators](#).

7.3.12 HLL

HyperLoglog (HLL) is an approximation algorithm for efficiently counting the number of distinct values in a dataset. It features faster computing and lower space usage. You only need to store HLL data structures, instead of datasets. When new data is added to a dataset, make hash calculation on the data and insert the result to an HLL. Then, you can obtain the final result based on the HLL.

Table 7-20 compares HLL with other algorithms.

Table 7-20 Comparison between HLL and other algorithms

Item	Sorting Algorithm	Hash Algorithm	HLL
Time complexity	$O(n\log n)$	$O(n)$	$O(n)$
Space complexity	$O(n)$	$O(n)$	$\log(\log n)$
Error rate	0	0	$\approx 0.8\%$
Storage space requirement	Size of original data	Size of original data	The maximum size is 16 KB by default.

HLL has advantages over others in the computing speed and storage space requirement. In terms of time complexity, the sorting algorithm needs $O(n\log n)$ time for sorting, and the hash algorithm and HLL need $O(n)$ time for full table scanning. In terms of storage space requirements, the sorting algorithm and hash algorithm need to store raw data before collecting statistics, whereas the HLL algorithm needs to store only the HLL data structures rather than the raw data, and thereby occupying a fixed space of about 16 KB.

NOTICE

- In the current default specifications, the maximum number of distinct values that can be calculated is about $1.1e + 15$, and the error rate is 0.8%. If the calculation result exceeds the maximum, the error rate of the calculation result will increase, or the calculation will fail and an error will be reported.
- When using this feature for the first time, you need to evaluate the distinct values of the service, properly select configuration parameters, and perform verification to ensure that the accuracy meets requirements.
 - By default, the distinct value is $1.1e + 15$. If the distinct value is NaN, you need to adjust log2m or use another algorithm to calculate the distinct value.
 - The hash algorithm has an extremely low probability of collision. However, you are still advised to select 2 or 3 hash seeds for verification when using the hash algorithm for the first time. If there is only a small difference between the distinct values, you can select any one of the seeds as the hash seed.

Table 7-21 describes main HLL data structures.

Table 7-21 Main HLL data structures

Data Type	Description
hll	The HLL header is a 27-byte field. By default, the data length ranges from 0 KB to 16 KB. The distinct value can be obtained.

When you create an HLL data type, 0 to 4 input parameters are supported. The parameter meanings and specifications are the same as those of the **hll_empty** function. The first parameter is **log2m**, indicating the logarithm of the number of buckets, and its value ranges from 10 to 16. The second parameter is **log2explicit**, indicating the threshold in explicit mode, and its value ranges from 0 to 12. The third parameter is **log2sparse**, indicating the threshold of the Sparse mode, and its value ranges from 0 to 14. The fourth parameter is **duplicatecheck**, indicating whether to enable duplicatecheck, and its value ranges from 0 to 1. When the input parameter is set to **-1**, the default value of the HLL parameter is used. You can run the **\d** or **\d+** command to view the parameters of the HLL type.

NOTE

When the HLL data type is created, the result varies depending on the input parameter behavior:

- When creating an HLL type, do not set the input parameter or set it to **-1**. Use the default value of the corresponding HLL parameter.
- If a valid value is set for the input parameter, the corresponding HLL parameter uses the input value.
- If the input value is invalid, an error is reported when the HLL type is created.

```
-- Create an HLL table without specifying input parameters.
openGauss=# CREATE TABLE t1 (id integer, set hll);
openGauss=# \d t1
      Table "public.t1"
  Column | Type      | Modifiers
-----+-----+-----
 id      | integer   |
 set     | hll       |

-- Create an HLL table, specify the first two input parameters, and use the default values for the last two
input parameters.
openGauss=# CREATE TABLE t2 (id integer, set hll(12,4));
openGauss=# \d t2
      Table "public.t2"
  Column | Type      | Modifiers
-----+-----+-----
 id      | integer   |
 set     | hll(12,4,12,0) |

-- Create an HLL table, specify the third input parameter, and use default values for other parameters.
openGauss=# CREATE TABLE t3(id int, set hll(-1,-1,8,-1));
openGauss=# \d t3
      Table "public.t3"
  Column | Type      | Modifiers
-----+-----+-----
 id      | integer   |
 set     | hll(14,10,8,0) |
```

```
-- When a user creates an HLL table and specifies an invalid input parameter, an error is reported.
openGauss=# CREATE TABLE t4(id int, set hll(5,-1));
ERROR: log2m = 5 is out of range, it should be in range 10 to 16, or set -1 as default

-- Drop the created HLL table.
openGauss=# DROP TABLE t1,t2,t3;
DROP TABLE
```

NOTE

When inserting an HLL object to an HLL table, ensure that the parameters of the HLL type are the same as those of the inserted object. Otherwise, an error is reported.

```
-- Create an HLL table.
openGauss=# CREATE TABLE t1(id integer, set hll(14));

-- Insert an HLL object to a table. The insertion succeeds because parameter types are consistent.
openGauss=# INSERT INTO t1 VALUES (1, hll_empty(14,-1));

-- Insert an HLL object to a table. The insertion fails because parameter types are inconsistent.
openGauss=# INSERT INTO t1(id, set) VALUES (1, hll_empty(14,5));
ERROR: log2explicit does not match: source is 5 and dest is 10

-- Drop the table.
openGauss=# DROP TABLE t1;
```

The following describes HLL application scenarios:

- Scenario 1: "Hello World"

The following example shows how to use the HLL data type:

```
-- Create an HLL table.
openGauss=# CREATE TABLE helloworld (id integer, set hll);

-- Insert an empty HLL to the table.
openGauss=# INSERT INTO helloworld(id, set) VALUES (1, hll_empty());

-- Add a hashed integer to the HLL.
openGauss=# UPDATE helloworld SET set = hll_add(set, hll_hash_integer(12345)) WHERE id = 1;

-- Add a hashed string to the HLL.
openGauss=# UPDATE helloworld SET set = hll_add(set, hll_hash_text('hello world')) WHERE id = 1;

-- Obtain the number of distinct values of the HLL.
openGauss=# SELECT hll_cardinality(set) FROM helloworld WHERE id = 1;
 hll_cardinality
-----
                2
(1 row)

-- Drop the table.
openGauss=# DROP TABLE helloworld;
```

- Scenario 2: Collect statistics about website visitors.

The following example shows how an HLL collects statistics on the number of users visiting a website within a period of time:

```
-- Create a raw data table to show that a user has visited the website at a certain time.
openGauss=# CREATE TABLE facts (
    date        date,
    user_id     integer
);

-- Create a raw data table to show that a user has visited the website at a certain time.
openGauss=# INSERT INTO facts VALUES ('2019-02-20', generate_series(1,100));
openGauss=# INSERT INTO facts VALUES ('2019-02-21', generate_series(1,200));
openGauss=# INSERT INTO facts VALUES ('2019-02-22', generate_series(1,300));
openGauss=# INSERT INTO facts VALUES ('2019-02-23', generate_series(1,400));
openGauss=# INSERT INTO facts VALUES ('2019-02-24', generate_series(1,500));
openGauss=# INSERT INTO facts VALUES ('2019-02-25', generate_series(1,600));
```

```
openGauss=# INSERT INTO facts VALUES ('2019-02-26', generate_series(1,700));
openGauss=# INSERT INTO facts VALUES ('2019-02-27', generate_series(1,800));

-- Create another table and specify an HLL column:
openGauss=# CREATE TABLE daily_uniques (
    date          date UNIQUE,
    users         hll
);

-- Group data by date and insert the data into the HLL.
openGauss=# INSERT INTO daily_uniques(date, users)
    SELECT date, hll_add_agg(hll_hash_integer(user_id))
    FROM facts
    GROUP BY 1;

-- Calculate the numbers of users visiting the website every day.
openGauss=# SELECT date, hll_cardinality(users) FROM daily_uniques ORDER BY date;
    date | hll_cardinality
-----+-----
2019-02-20 |          100
2019-02-21 | 200.217913059312
2019-02-22 | 301.76494508014
2019-02-23 | 400.862858326446
2019-02-24 | 502.626933349694
2019-02-25 | 601.922606454213
2019-02-26 | 696.602316769498
2019-02-27 | 798.111731634412
(8 rows)

-- Calculate the number of users who had visited the website in the week from February 20, 2019 to
February 26, 2019.
openGauss=# SELECT hll_cardinality(hll_union_agg(users)) FROM daily_uniques WHERE date >=
'2019-02-20'::date AND date <= '2019-02-26'::date;
    hll_cardinality
-----
702.941844662509
(1 row)

-- Calculate the number of users who had visited the website yesterday but have not visited the
website today:
openGauss=# SELECT date, (#hll_union_agg(users) OVER two_days) - #users AS lost_uniques FROM
daily_uniques WINDOW two_days AS (ORDER BY date ASC ROWS 1
PRECEDING);
    date | lost_uniques
-----+-----
2019-02-20 |          0
2019-02-21 |          0
2019-02-22 |          0
2019-02-23 |          0
2019-02-24 |          0
2019-02-25 |          0
2019-02-26 |          0
2019-02-27 |          0
(8 rows)

-- Drop the table.
openGauss=# DROP TABLE facts;
openGauss=# DROP TABLE daily_uniques;
```

- Scenario 3: The data to be inserted does not meet the requirements of the HLL data structure.

When inserting data into a column of the HLL type, ensure that the data meets the requirements of the HLL data structure. If the data does not meet the requirements after being parsed, an error will be reported. In the following example, 'E\1234' to be inserted does not meet the requirements of the HLL data structure after being parsed. As a result, an error is reported.

```
openGauss=# CREATE TABLE test(id integer, set hll);
openGauss=# INSERT INTO test VALUES(1, 'E\\1234');
ERROR: not a hll type, size=6 is not enough

openGauss=# DROP TABLE test;
```

7.3.13 Range Types

A range type is a data type that represents the range of a value of an element type (called the subtype of a range). For example, the range of timestamp may be used to express a time range in which a conference room is reserved. In this case, the data type is tsrange (short for timestamp range), and timestamp is the subtype. The subtype must have an overall order so that the element value can be clearly specified within a range, before, or after.

Range types are useful because they can express multiple element values in a single range value and can clearly express concepts such as range overlapping (time and date ranges for time arrangement, price ranges, and instrument ranges).

Built-in Range

The following built-in ranges are available:

- int4range: integer range.
- int8range: bigint range.
- numrange: numeric range.
- tsrange: range of timestamp without the time zone.
- tstzrange: range of timestamp with the time zone
- daterange: date range.

In addition, you can define your own range types. For details, see [CREATE TYPE](#).

Example

```
-- Create a table and insert data into the table.
openGauss=# CREATE TABLE reservation (room int, during tsrange);
openGauss=# INSERT INTO reservation VALUES (1108, '[2010-01-01 14:30, 2010-01-01 15:30)');
-- Inclusion
openGauss=# SELECT int4range(10, 20) @> 3;
?column?
-----
f
(1 row)
-- Overlapping
openGauss=# SELECT numrange(11.1, 22.2) && numrange(20.0, 30.0);
?column?
-----
t
(1 row)
-- Upper bound extraction
openGauss=# SELECT upper(int8range(15, 25));
upper
-----
25
(1 row)
-- Intersection set
openGauss=# SELECT int4range(10, 20) * int4range(15, 25);
?column?
-----
```

```
[15,20)
(1 row)
-- Is the range empty?
openGauss=# SELECT isempty(numrange(1, 5));
isempty
-----
f
(1 row)
-- Drop the table.
openGauss=# DROP TABLE reservation;
```

See the complete list of operators and functions on a range type in [Range Functions and Operators](#).

Including and Excluding Bounds

Each non-empty range has two bounds, a lower bound and an upper bound. All values between the upper and lower bounds are included in the range. An inclusion bound means that the bound value itself is included in the range, while an exclusion bound means that the bound value is not included in the range.

In a textual form of a range, the inclusion lower bound is expressed as "[" and an exclusion lower bound is expressed as "(". Similarly, one containing the upper bound is expressed as "]" and one excluding the upper bound is expressed as ")" (for details, see [Range Input/Output](#)).

The lower_inc and upper_inc functions test the upper and lower bounds of a range value, respectively.

Infinite (Unbounded) Range

When the lower bound of a range is unbounded, it means that all values less than the upper bound are included in the range. Similarly, when the upper bound of a range is unbounded, all values greater than the lower bound are included in the range. When both the upper and lower bounds are unbounded, all values of the element type are considered within the range. The missing bounds are automatically converted to exclusions. You can think of these missing values as positive infinity or negative infinity, but they are special range type values and are considered to be positive and negative infinity values that go beyond any range element type.

Element types with the infinity values can be used as explicit bound values. For example, in the timestamp range, [today, infinity) does not include a special timestamp value infinity.

The lower_inf and upper_inf functions test the infinite upper and lower bounds of a range, respectively.

Range Input/Output

The input of a range value must follow one of the following formats:

```
(lower-bound, upper-bound)
(lower-bound, upper-bound]
[lower-bound, upper-bound)
[lower-bound, upper-bound]
Empty
```

The output of a range value must follow one of the following formats:

[lower-bound, upper-bound)
Empty

Parentheses () or square brackets [] indicate whether the upper and lower bounds are excluded or included. Note that the last format is empty, which represents an empty range (a range that does not contain values).

The value of *lower-bound* can be a valid input character string of the subtype or null, indicating that there is no lower bound. Similarly, *upper-bound* can be a valid input character string of the subtype or null, indicating that there is no upper bound.

Each bound value can be referenced using the quotation marks('') character. This is necessary if the bounds value contains parentheses (), square brackets [], commas (,), quotation marks (''), or backslashes (\). Otherwise, those characters will be considered as the part of the range syntax. To put the quotation mark or backslash in a referenced bound value, put a backslash in front of it (and a pair of double quotation marks in its quoted bound value represents one quotation mark character, which is similar to the single quotation mark rule in SQL character strings). In addition, you can avoid referencing and use backslash escapes to protect all data characters, otherwise they will be used as part of the return syntax. Also, if you want to write a bound value that is an empty string, write '', indicating infinite bounds.

Spaces are allowed before and after a range value, but any space between parentheses() or square brackets[] is used as part of the upper or lower bound value (depending on the element type, the space may or may not represent a value).

Examples:

```
-- 3 is included, 7 is not included, and all values between 3 and 7 are included.
openGauss=# SELECT '[3,7)':int4range;
int4range
-----
[3,7)
(1 row)
-- Neither 3 nor 7 is included, but all values between them are included.
openGauss=# SELECT '(3,7)':int4range;
int4range
-----
[4,7)
(1 row)
-- Only 4 is included.
openGauss=# SELECT '[4,4]':int4range;
int4range
-----
[4,5)
(1 row)
-- Exclude any value (and will be normalized to empty).
openGauss=# SELECT '[4,4)':int4range;
int4range
-----
Empty
(1 row)
```

Constructing Range

Each range type has a constructor function with the same name. Using constructor functions is often more convenient than writing a range literal constant because it avoids extra references to bound values. Constructor functions accept two or three parameters. Two parameters form a range in the standard

form, where the lower bound is included and the upper bound is excluded, and three parameters form a range according to the bound specified by the third parameter. The third parameter must be one of the following character strings: (), (], [), or []. For example:

```
-- The complete format is: lower bound, upper bound, and textual parameters indicating the inclusion/
exclusion of bounds.
openGauss=# SELECT numrange(1.0, 14.0, '()');
numrange
-----
(1.0,14.0]
(1 row)
-- If the third parameter is ignored, it is assumed to be '()'.
openGauss=# SELECT numrange(1.0, 14.0);
numrange
-----
[1.0,14.0)
(1 row)
-- Although '()' is specified here, the value will be converted to the standard format when displayed,
because int8range is a discrete range type (see below).
openGauss=# SELECT int8range(1, 14, '()');
int8range
-----
[2,15)
(1 row)
-- Using NULL for a bound causes the range to be unbounded on that side.
openGauss=# SELECT numrange(NULL, 2.2);
numrange
-----
(,2.2)
(1 row)
```

Discrete Range

A range element type has a well-defined "step" such as integer or date. In these types, if there is no valid value between two elements, they can be said to be adjacent. This is in contrast to a continuous range in which other element values can always be identified between two given values. For example, a range above the numeric type is continuous, and the range of timestamp is also continuous. (Although timestamp has limited precision and can be considered as discrete in theory, it can be considered as continuous because the step is not normally considered.)

Another way to consider discrete range types is to have a clear "next" or "previous" value for each element value. With this idea in mind, you can switch between inclusion and exclusion expressions of a range bound by replacing it with the original given next or previous element value. For example, in an integer range type, [4,8] and (3,9) represent the same set of values, but not for numeric ranges.

A discrete range type should have a *regularization* function that knows the expected step size of the element type. The regularization function can convert the equivalents of the range type to the same expression, in particular consistent with the inclusion or exclusion bounds. If you do not specify a regularization function, ranges with different formats will always be considered as unequal, even if they actually express the same set of values.

The built-in range types int4range, int8range, and daterange use a regularized form that includes the lower bound and excludes the upper bound, that is, []). However, user-defined range types can use other conventions.

Defining New Range

Users can define their own range types. A common reason is to use the range on the subtype that is not provided in the built-in range type. For example, to create the range type subtype float8, run the following command:

```
openGauss=# CREATE TYPE floatrange AS RANGE (
    subtype = float8,
    subtype_diff = float8mi
);
openGauss=# SELECT '[1.234, 5.678]':floatrange;
floatrange
-----
[1.234,5.678]
(1 row)
openGauss=# DROP TYPE floatrange;
```

Because float8 does not have a meaningful "step", the regularization function is not defined in this example.

Defining your own range type also allows you to specify a different subtype B-tree operator class or collection to change the sort order to determine which values fall within the given range.

If the subtype is considered to have a discrete value instead of a continuous value, the **CREATE TYPE** command should specify a canonical function. The regularization function receives an input range value and must return an equivalent range value that may have different bounds and formats. For two ranges, for example, [1, 7] and [1, 8) that represent the same value set, the output must be the same. There is no relationship between choosing which expression to use as the regularization function, as long as two values of equal value in different formats can always be mapped to the same value in the same format. In addition to adjusting the inclusion/exclusion bound format, if the expected compensation is larger than the subtype can store, a regularization function may round the bound value. For example, a range type above a timestamp might be defined as having a one-hour epoch, so the regularization function might need to round off bounds that are not multiples of an hour, or might throw an error directly.

The subtype difference function uses two subtype input values and returns a difference expressed as a float8 value (X minus Y). In the example above, we can use functions under the regular float8 subtraction operator. However, for any other subtype, some type conversion may be required. There may also be a need for innovative ideas on how to express differences as numbers. For maximum extensibility, the subtype_diff function should agree with the sort order of the selected operator class and sort rules. That is, if the first parameter of the sort order is greater than the second parameter, the result should be a positive value.

The following is an example of the subtype_diff function:

```
openGauss=# CREATE FUNCTION time_subtype_diff(x time, y time) RETURNS float8 AS 'SELECT
EXTRACT(EPOCH FROM (x - y))' LANGUAGE sql STRICT IMMUTABLE;
openGauss=# CREATE TYPE timerange AS RANGE (
    subtype = time,
    subtype_diff = time_subtype_diff
);
openGauss=# SELECT '[11:10, 23:00]':timerange;
timerange
-----
[11:10:00,23:00:00]
```

```
(1 row)
openGauss=# DROP TYPE timerange;
openGauss=# DROP FUNCTION time_subtype_diff;
```

For details about how to create a range type, see [CREATE TYPE](#).

Index

B-tree and hash indexes can be created on table columns of the range type. For these index types, basically the only useful range operation is equivalence. Using the corresponding < and > operators, there is a B-tree sort order for range value definitions, but that order is fairly arbitrary and is often less useful in the reality. The B-tree and hash support for range types is primarily designed to allow sorting and hashing within a query, rather than creating an index.

7.3.14 OID Types

OIDs are used internally by GaussDB as primary keys for various system catalogs. OIDs are not added to user-created tables by the system. The OID type represents an object identifier.

The OID type is currently implemented as an unsigned four-byte integer. So, using a user-created table's **OID** column as a primary key is discouraged.

Table 7-22 OID types

Name	Reference	Description	Example
OID	N/A	Numeric object identifier	564182
CID	N/A	Command identifier. This is the data type of the system columns cmin and cmax . Command identifiers are 32-bit quantities.	N/A
XID	N/A	Transaction identifier. This is the data type of the system columns xmin and xmax . Transaction identifiers are also 64-bit quantities.	N/A
TID	N/A	Row identifier. This is the data type of the system column ctid . A row ID is a pair (block number, tuple index within block) that identifies the physical location of the row within its table.	N/A
REGCONFIG	pg_ts_config	Text search configuration	english

Name	Reference	Description	Example
REGDICTIONARY	pg_ts_dict	Text search dictionary	simple
REGOPER	pg_operator	Operator name	N/A
REGOPERATOR	pg_operator	Operator with parameter types	*(integer,integer) or -(NONE,integer)
REGPROC	pg_proc	Function name	sum
REGPROCEDURE	pg_proc	Function with parameter types	sum(int4)
REGCLASS	pg_class	Relation name	pg_type
REGTYPE	pg_type	Data type name	integer

The OID type is used for a column in the database system catalog.

Example:

```
openGauss=# SELECT oid FROM pg_class WHERE relname = 'pg_type';
oid
-----
1247
(1 row)
```

The alias type for **OID** is **REGCLASS** which allows simplified search for **OID** values.

Example:

```
openGauss=# SELECT attrelid,attname,atttypid,attstattarget FROM pg_attribute WHERE attrelid =
'pg_type'::REGCLASS;
attrelid | attname | atttypid | attstattarget
-----+-----+-----+-----
1247 | xc_node_id | 23 | 0
1247 | tableoid | 26 | 0
1247 | cmax | 29 | 0
1247 | xmax | 28 | 0
1247 | cmin | 29 | 0
1247 | xmin | 28 | 0
1247 | oid | 26 | 0
1247 | ctid | 27 | 0
1247 | typname | 19 | -1
1247 | typnamespace | 26 | -1
1247 | typowner | 26 | -1
1247 | typplen | 21 | -1
1247 | typbyval | 16 | -1
1247 | typtype | 18 | -1
1247 | typcategory | 18 | -1
1247 | typispreferred | 16 | -1
1247 | typisdefined | 16 | -1
1247 | typdelim | 18 | -1
1247 | typrelid | 26 | -1
1247 | typelem | 26 | -1
1247 | typarray | 26 | -1
1247 | typinput | 24 | -1
1247 | typoutput | 24 | -1
1247 | typepreceive | 24 | -1
```

1247	typsend		24		-1
1247	typmodin		24		-1
1247	typmodout		24		-1
1247	typanalyze		24		-1
1247	typalign		18		-1
1247	typstorage		18		-1
1247	typnotnull		16		-1
1247	typbasetype		26		-1
1247	typtypmod		23		-1
1247	typndims		23		-1
1247	typcollation		26		-1
1247	typdefaultbin		194		-1
1247	typdefault		25		-1
1247	typacl		1034		-1

(38 rows)

7.3.15 Pseudo-Types

GaussDB type system contains a number of special-purpose entries that are collectively called pseudo-types. A pseudo-type cannot be used as a column data type, but it can be used to declare a function's argument or result type.

Each of the available pseudo-types is useful in situations where a function's behavior does not correspond to simply taking or returning a value of a specific SQL data type. [Table 7-23](#) lists all pseudo-types.

Table 7-23 Pseudo-types

Name	Description
any	Indicates that a function accepts any input data type.
anyelement	Indicates that a function accepts any data type.
anyarray	Indicates that a function accepts any array data type.
anynonarray	Indicates that a function accepts any non-array data type.
anyenum	Indicates that a function accepts any enum data type.
anyrange	Indicates that a function accepts any range data type.
cstring	Indicates that a function accepts or returns a null-terminated C string.
internal	Indicates that a function accepts or returns a server-internal data type.
language_handler	Indicates that a procedural language call handler is declared to return language_handler .
fdw_handler	Indicates that a foreign data wrapper handler is declared to return fdw_handler .
record	Identifies a function returning an unspecified row type.
trigger	Indicates that a trigger function is declared to return trigger .
void	Indicates that a function returns no value.

Name	Description
opaque	Indicates an obsolete type name that formerly served all the above purposes.

Functions coded in C (whether built in or dynamically loaded) can be declared to accept or return any of these pseudo data types. It is up to the function author to ensure that the function will behave safely when a pseudo-type is used as an argument type.

Functions coded in procedural languages can use pseudo-types only as allowed by their implementation languages. At present the procedural languages all forbid use of a pseudo-type as argument type, and allow only **void** and **record** as a result type. Some polymorphic functions support the **anyelement**, **anyarray**, **anynonarray**, **anyenum**, and **anyrange** types.

The internal pseudo-type is used to declare functions that are meant only to be called internally by the database system, and not by direct calling in an SQL query. If a function has at least one internal-type argument, then it cannot be called from SQL. You are advised not to create any function that is declared to return internal unless it has at least one internal argument.

Example:

```
-- Create a table.
openGauss=# create table t1 (a int);

-- Insert two data records.
openGauss=# insert into t1 values(1),(2);

-- Create the showall() function.
openGauss=# CREATE OR REPLACE FUNCTION showall() RETURNS SETOF record
AS $$ SELECT count(*) from t1; $$
LANGUAGE SQL;

-- Call the showall() function.
openGauss=# SELECT showall();
showall
-----
(2)
(1 row)

-- Delete the function.
openGauss=# DROP FUNCTION showall();

-- Drop the table.
openGauss=# drop table t1;
```

7.3.16 ACLItem

The aclitem data type is used to store object permission information. Its internal implementation is of the int type and supports the '*user1* = *privs*/*user2*' format.

The aclitem[] data type is an array consisting of ACLItems. The supported format is {*user1* = *privs1*/*user3*, *user2* = *privs2*/*user3*}.

In the preceding command, *user1*, *user2*, and *user3* indicate the existing users or roles in the database, and *privs* indicates the permissions supported by the database. For details, see [Table 12-34](#).

Example:

```
-- Create users.
openGauss=# CREATE USER user1 WITH PASSWORD 'Aa123456789';
openGauss=# CREATE USER user2 WITH PASSWORD 'Aa123456789';
openGauss=# CREATE USER omm WITH PASSWORD 'Aa123456789';

-- Create a data table table_acl that contains three columns of the int, aclitem, and aclitem[] types.
openGauss=# CREATE TABLE table_acl (id int,priv aclitem,privs aclitem[]);
-- Insert a data record whose content is (1,'user1=arw/omm',{'omm=d/user2,omm=w/omm'}) into the table_acl table.
openGauss=# INSERT INTO table_acl VALUES (1,'user1=arw/omm',{'omm=d/user2,omm=w/omm'});
-- Insert a data record whose content is (2,'user1=aw/omm',{'omm=d/user2'}) into the table_acl table.
openGauss=# INSERT INTO table_acl VALUES (2,'user1=aw/omm',{'omm=d/user2'});
openGauss=# SELECT * FROM table_acl;
id | priv | privs
-----+-----
1 | user1=arw/omm | {omm=d/user2,omm=w/omm}
2 | user1=aw/omm | {omm=d/user2}
(2 rows)

-- Drop the table and users.
openGauss=# DROP USER user1;
openGauss=# DROP USER user2;
openGauss=# DROP USER omm;
openGauss=# DROP TABLE table_acl;
```

7.4 Constant and Macro

Table 7-24 lists the constants and macros that can be used in GaussDB.

Table 7-24 Constant and macro

Parameter	Description	Example
CURRENT_CATALOG	Specifies the current database.	testdb=# SELECT CURRENT_CATALOG; current_database ----- testdb (1 row)
CURRENT_ROLE	Specifies the current user.	openGauss=# SELECT CURRENT_ROLE; current_user ----- omm (1 row)
CURRENT_SCHEMA	Specifies the current database schema.	openGauss=# SELECT CURRENT_SCHEMA; current_schema ----- public (1 row)
CURRENT_USER	Specifies the current user.	openGauss=# SELECT CURRENT_USER; current_user ----- omm (1 row)
LOCALTIMESTAMP	Specifies the current session time (without time zone).	openGauss=# SELECT LOCALTIMESTAMP; timestamp ----- 2015-10-10 15:37:30.968538 (1 row)

Parameter	Description	Example
NULL	This parameter is left blank.	N/A
SESSION_USER	Specifies the current system user.	openGauss=# SELECT SESSION_USER; session_user ----- omm (1 row)
SYSDATE	Specifies the current system date.	openGauss=# SELECT SYSDATE; sysdate ----- 2015-10-10 15:48:53 (1 row)
USER	Specifies the current user, also called CURRENT_USER .	openGauss=# SELECT USER; current_user ----- omm (1 row)

7.5 Functions and Operators

Operators can be used to process one or more operands and can be placed before, after, or between operands. Results are returned after the processing.

Functions encapsulate service logic to implement specific functions. A function may or may not have parameters. After a function is executed, the result is returned.

Users can modify system functions. However, after the modification, the meaning of the functions may change, which results in disorder in system control. Therefore, users are not allowed to manually modify system functions.

7.5.1 Logical Operators

Common logical operators include AND, OR, and NOT. The operation result has three values: **TRUE**, **FALSE**, and **NULL**. NULL indicates unknown. Their priorities are NOT > AND > OR.

Table 7-25 lists the calculation rules, where a and b represent logical expressions.

Table 7-25 Operation rules

a	b	a AND b Result	a OR b Result	NOT a Result
TRUE	TRUE	TRUE	TRUE	FALSE
TRUE	FALSE	FALSE	TRUE	FALSE
TRUE	NULL	NULL	TRUE	FALSE
FALSE	FALSE	FALSE	FALSE	TRUE

a	b	a AND b Result	a OR b Result	NOT a Result
FALSE	NULL	FALSE	NULL	TRUE
NULL	NULL	NULL	NULL	NULL

 NOTE

The operators AND and OR are commutative, that is, you can switch the left and right operand without affecting the result.

7.5.2 Comparison Operators

Comparison operators are available for the most data types and return Boolean values.

All comparison operators are binary operators. Only data types that are the same or can be implicitly converted can be compared using comparison operators.

Table 7-26 describes comparison operators provided by GaussDB.

Table 7-26 Comparison operators

Operator	Description
<	Less than
>	Greater than
<=	Less than or equal to
>=	Greater than or equal to
=	Equal to
<>, !=, or ^=	Not equal to

- Comparison operators are available for all relevant data types. All comparison operators are binary operators that returned values of Boolean type. The calculation priority of the inequality sign is higher than that of the equality sign. If the entered data type is different and cannot be implicitly converted, the comparison fails. For example, an expression such as 1<2<3 is invalid because the less-than sign (<) cannot be used to compare Boolean values and 3.
- Besides, each comparison operator has a corresponding function in the **pg_proc** system catalog. If the value of proleakproof attribute of the corresponding function is **f**, the function is not used to prevent data leakage. If a user only has the permission for a system view, but does not have the permission for the corresponding table, the query plan may not be optimal when the user searches the system view.

7.5.3 Character Processing Functions and Operators

String functions and operators provided by GaussDB are for concatenating strings with each other, concatenating strings with non-strings, and matching the patterns of strings. Note: Except length-related functions, other functions and operators of string processing functions do not support parameters of CLOB whose size is greater than 1 GB.

- `bit_length(string)`

Description: Specifies the number of bits occupied by a string.

Return type: int

Example:

```
openGauss=# SELECT bit_length('world');
 bit_length
-----
         40
(1 row)
```

- `btrim(string text [, characters text])`

Description: Removes the longest string consisting only of characters in **characters** (a space by default) from the start and end of **string**.

Return type: text

Example:

```
openGauss=# SELECT btrim('sring', 'ing');
 btrim
-----
 sr
(1 row)
```

- `char_length(string)` or `character_length(string)`

Description: Specifies the number of characters in a string.

Return type: int

Example:

```
openGauss=# SELECT char_length('hello');
 char_length
-----
          5
(1 row)
```

- `instr(text,text,int,int)`

Description: **instr(string1,string2,int1,int2)** returns the text from **int1** to **int2** in **string1**. The first **int** indicates the start position for matching, and the second **int** indicates the number of matching times.

Return type: int

Example:

```
openGauss=# SELECT instr('abcdabcdabcd', 'bcd', 2, 2);
 instr
-----
      6
(1 row)
```

- `lengthb(text/bpchar)`

Description: Obtains the number of bytes of a specified string.

Return type: int

Example:

```
openGauss=# SELECT lengthb('hello');
lengthb
-----
      5
(1 row)
```

- left(str text, n int)

Description: Returns the first *n* characters in a string. When *n* is negative, all but the last **|n|** characters are returned.

Return type: text

Example:

```
openGauss=# SELECT left('abcde', 2);
left
-----
ab
(1 row)
```

- length(string bytea, encoding name)

Description: Specifies the number of characters in **string** in the given **encoding**. The string must be valid in this encoding.

Return type: int

Example:

```
openGauss=# SELECT length('jose', 'UTF8');
length
-----
      4
(1 row)
```

NOTE

If the length of the bytea type is queried and UTF8 encoding is specified, the maximum length can only be **536870888**.

- lpad(string text, length int [, fill text])

Description: Fills up **string** to **length** by appending the characters **fill** (a space by default). If **string** is already longer than **length**, then it is truncated.

Return type: text

Example:

```
openGauss=# SELECT lpad('hi', 5, 'xyza');
lpad
-----
xyzhi
(1 row)
```

- notlike(x bytea name text, y bytea text)

Description: Compares x and y to check whether they are inconsistent.

Return type: Boolean

Example:

```
openGauss=# SELECT notlike(1,2);
notlike
-----
      t
(1 row)
openGauss=# SELECT notlike(1,1);
notlike
-----
      f
(1 row)
```

- `octet_length(string)`

Description: Specifies the number of bytes in a string.

Return type: int

Example:

```
openGauss=# SELECT octet_length('jose');
octet_length
-----
         4
(1 row)
```

- `overlay(string placing string FROM int [for int])`

Description: Replaces substrings. **FROM int** indicates the start position of the replacement in the first string. **for int** indicates the number of characters replaced in the first string.

Return type: text

Example:

```
openGauss=# SELECT overlay('hello' placing 'world' from 2 for 3 );
overlay
-----
hworldo
(1 row)
```

- `position(substring in string)`

Description: Specifies the position of a substring. Parameters are case-sensitive.

Return type: int. If the character string does not exist, **0** is returned.

Example:

```
openGauss=# SELECT position('ing' in 'string');
position
-----
         4
(1 row)
```

- `pg_client_encoding()`

Description: Specifies the current client-side encoding name.

Return type: name

Example:

```
openGauss=# SELECT pg_client_encoding();
pg_client_encoding
-----
UTF8
(1 row)
```

- `quote_ident(string text)`

Description: Returns the given string suitably quoted to be used as an identifier in an SQL statement string (quotation marks are used as required). Quotation marks are added only if necessary (that is, if the string contains non-identifier characters or would be case-folded). Embedded quotation marks are properly doubled.

Return type: text

Example:

```
openGauss=# SELECT quote_ident('hello world');
quote_ident
-----
"hello world"
(1 row)
```

- `quote_literal(string text)`

Description: Returns the given string suitably quoted to be used as a string literal in an SQL statement string (quotation marks are used as required).

Return type: text

Example:

```
openGauss=# SELECT quote_literal('hello');
quote_literal
-----
'hello'
(1 row)
```

If a command similar to the following exists, the text will be escaped.

```
openGauss=# SELECT quote_literal(E'O\'hello');
quote_literal
-----
'O\'hello'
(1 row)
```

If a command similar to the following exists, the backslash will be properly doubled.

```
openGauss=# SELECT quote_literal('O\'hello');
quote_literal
-----
'E'O\\hello'
(1 row)
```

If the parameter is null, **NULL** is returned. If the parameter may be null, you are advised to use **quote_nullable**.

```
openGauss=# SELECT quote_literal(NULL);
quote_literal
-----
(1 row)
```

- `quote_literal(value anyelement)`

Description: Converts the given value to text and then quotes it as a literal.

Return type: text

Example:

```
openGauss=# SELECT quote_literal(42.5);
quote_literal
-----
'42.5'
(1 row)
```

If a command similar to the following exists, the given value will be escaped.

```
openGauss=# SELECT quote_literal(E'O\'42.5');
quote_literal
-----
'O\'42.5'
(1 row)
```

If a command similar to the following exists, the backslash will be properly doubled.

```
openGauss=# SELECT quote_literal('O\'42.5');
quote_literal
-----
'E'O\\42.5'
(1 row)
```

- `quote_nullable(string text)`

Description: Returns the given string suitably quoted to be used as a string literal in an SQL statement string (quotation marks are used as required).

Return type: text

Example:

```
openGauss=# SELECT quote_nullable('hello');
quote_nullable
-----
'hello'
(1 row)
```

If a command similar to the following exists, the text will be escaped.

```
openGauss=# SELECT quote_nullable(E'O\'hello');
quote_nullable
-----
'O\'hello'
(1 row)
```

If a command similar to the following exists, the backslash will be properly doubled.

```
openGauss=# SELECT quote_nullable('O\'hello');
quote_nullable
-----
E'O\\hello'
(1 row)
```

If the parameter is null, **NULL** is returned.

```
openGauss=# SELECT quote_nullable(NULL);
quote_nullable
-----
NULL
(1 row)
```

- `quote_nullable(value anyelement)`

Description: Converts the given value to text and then quotes it as a literal.

Return type: text

Example:

```
openGauss=# SELECT quote_nullable(42.5);
quote_nullable
-----
'42.5'
(1 row)
```

If a command similar to the following exists, the given value will be escaped.

```
openGauss=# SELECT quote_nullable(E'O\'42.5');
quote_nullable
-----
'O\'42.5'
(1 row)
```

If a command similar to the following exists, the backslash will be properly doubled.

```
openGauss=# SELECT quote_nullable('O\'42.5');
quote_nullable
-----
E'O\\42.5'
(1 row)
```

If the parameter is null, **NULL** is returned.

```
openGauss=# SELECT quote_nullable(NULL);
quote_nullable
-----
NULL
(1 row)
```

- `substring_inner(string [from int] [for int])`

Description: Extracts a substring. **from int** indicates the start position of the truncation. **for int** indicates the number of bytes truncated.

Return type: text

Example:

```
openGauss=# SELECT substring_inner('adcde', 2,3);
substring_inner
-----
dcd
(1 row)
```

- `substring(string [from int] [for int])`

Description: Extracts a substring. **from int** indicates the start position of the truncation. **for int** indicates the number of bytes truncated.

Return type: text

Example:

```
openGauss=# SELECT substring('Thomas' from 2 for 3);
substring
-----
hom
(1 row)
```

- `rawcat(raw,raw)`

Description: Indicates the string concatenation function.

Return type: raw

Example:

```
openGauss=# SELECT rawcat('ab','cd');
rawcat
-----
ABCD
(1 row)
```

- `regexp_count(string text, pattern text [, position int [, flags text]])`

Description: obtains the number of substrings used for matching.

Parameter description:

- **string**: source character string used for matching.
- **pattern**: regular expression pattern string used for matching.
- **position**: sequence number of the character to be matched from the source character string. This parameter is optional. The default value is **1**.
- **flags**: contains zero or multiple single-letter flags that change the matching behavior of the function. This parameter is optional.

Return type: int

Example:

```
openGauss=# SELECT regexp_count('foobarbaz','b(..)', 5) AS RESULT;
result
-----
1
(1 row)
```

- `regexp_instr(string text, pattern text [, position int [, occurrence int [, return_opt int [, flags text]]]])`

Description: obtains the position (starting from 1) of the substring that meets the matching condition. If no substring is matched, **0** is returned.

Parameter description:

- **string**: source character string used for matching.
- **pattern**: regular expression pattern string used for matching.

- **position**: start character of the source string used for matching. This parameter is optional. The default value is **1**.
- **occurrence**: sequence number of the matched substring to be replaced. This parameter is optional. The default value is **1**.
- **return_opt**: specifies whether to return the position of the first or last character of the matched substring. This parameter is optional. If the value is **0**, the position of the first character (starting from 1) of the matched substring is returned. If the value is greater than 0, the position of the next character of the end character of the matched substring is returned. The default value is **0**.
- **flags**: contains zero or multiple single-letter flags that change the matching behavior of the function. This parameter is optional.

Return type: int

Example:

```
openGauss=# SELECT regexp_instr('foobabaz','b(..)', 1, 1, 0) AS RESULT;
result
-----
4
(1 row)

openGauss=# SELECT regexp_instr('foobabaz','b(..)', 1, 2, 0) AS RESULT;
result
-----
7
(1 row)
```

- repeat(string text, number int)

Description: Repeats **string** the specified number of times.

Return type: text

Example:

```
openGauss=# SELECT repeat('Pg', 4);
repeat
-----
PgPgPgPg
(1 row)
```

NOTE

The maximum size of memory allocated at a time cannot exceed 1 GB due to the memory allocation mechanism of the database. Therefore, the maximum value of **number** cannot exceed $(1 \text{ GB} - x) / \text{lengthb}(\text{string}) - 1$. **x** indicates the length of the header information, which is usually greater than 4 bytes. The value varies among different scenarios.

- replace(string text, from text, to text)

Description: Replaces all occurrences in **string** of substring **from** with substring **to**.

Return type: text

Example:

```
openGauss=# SELECT replace('abcdefabcdef', 'cd', 'XXX');
replace
-----
abXXXefabXXXef
(1 row)
```

- replace(string, substring)

Description: Deletes all substrings in a string.

String type: text

Substring type: text

Return type: text

Example:

```
openGauss=# SELECT replace('abcdefabcdef', 'cd');
replace
-----
abefabef
(1 row)
```

- **reverse(str)**

Description: Returns the reversed string.

Return type: text

Example:

```
openGauss=# SELECT reverse('abcde');
reverse
-----
edcba
(1 row)
```

- **right(str text, n int)**

Description: Returns the last *n* characters in a string. When *n* is negative, all but the first **|n|** characters are returned.

Return type: text

Example:

```
openGauss=# SELECT right('abcde', 2);
right
-----
de
(1 row)

openGauss=# SELECT right('abcde', -2);
right
-----
cde
(1 row)
```

- **rpad(string text, length int [, fill text])**

Description: Fills up **string** to **length** by appending the characters **fill** (a space by default). If **string** is already longer than **length**, then it is truncated.

Return type: text

Example:

```
openGauss=# SELECT rpad('hi', 5, 'xy');
rpad
-----
hixyx
(1 row)
```

- **rtrim(string text [, characters text])**

Description: Removes the longest string containing only characters from characters (a space by default) from the end of string.

Return type: text

Example:

```
openGauss=# SELECT rtrim('trimxxxx', 'x');
rtrim
-----
```



```
trim
(1 row)
```

- substrb(text,int,int)

Description: Extracts a substring. The first **int** indicates the start position of the subtraction. The second **int** indicates the number of characters subtracted.

Return type: text

Example:

```
openGauss=# SELECT substrb('string',2,3);
substrb
-----
tri
(1 row)
```

- substrb(text,int)

Description: Extracts a substring. **int** indicates the start position of the extraction.

Return type: text

Example:

```
openGauss=# SELECT substrb('string',2);
substrb
-----
tring
(1 row)
```

- substr(bytea,from,count)

Description: Extracts a substring from **bytea**. **from** specifies the position where the extraction starts. **count** specifies the length of the extracted substring.

Return type: text

Example:

```
openGauss=# SELECT substr('string',2,3);
substr
-----
tri
(1 row)
```

- string || string

Description: Concatenates strings.

Return type: text

Example:

```
openGauss=# SELECT 'MPP' || 'DB' AS RESULT;
result
-----
MPPDB
(1 row)
```

- string || non-string or non-string || string

Description: Concatenates strings and non-strings.

Return type: text

Example:

```
openGauss=# SELECT 'Value: ' || 42 AS RESULT;
result
-----
Value: 42
(1 row)
```

- `split_part(string text, delimiter text, field int)`

Description: Splits **string** on **delimiter** and returns the **fieldth** column (counting from text of the first appeared delimiter).

Return type: text

Example:

```
openGauss=# SELECT split_part('abc~@~def~@~ghi', '~@~', 2);
split_part
-----
def
(1 row)
```

- `strpos(string, substring)`

Description: Specifies the position of a substring. It is the same as **position(substring in string)**. However, the parameter sequences of them are reversed.

Return type: int

Example:

```
openGauss=# SELECT strpos('source', 'rc');
strpos
-----
4
(1 row)
```

- `to_hex(number int or bigint)`

Description: Converts a number to a hexadecimal expression.

Return type: text

Example:

```
openGauss=# SELECT to_hex(2147483647);
to_hex
-----
7fffffff
(1 row)
```

- `translate(string text, from text, to text)`

Description: Any character in **string** that matches a character in **from** is replaced by the corresponding character in **to**. If **from** is longer than **to**, extra characters occurred in **from** are removed.

Return type: text

Example:

```
openGauss=# SELECT translate('12345', '143', 'ax');
translate
-----
a2x5
(1 row)
```

- `length(string)`

Description: Obtains the number of characters in a string.

Return type: integer

Example:

```
openGauss=# SELECT length('abcd');
length
-----
4
(1 row)
```

- `lengthb(string)`

Description: Obtains the number of bytes in a string. The value depends on character sets (GBK and UTF8).

Return type: integer

Example:

```
openGauss=# SELECT lengthb('Chinese');
lengthb
-----
      7
(1 row)
```

- substr(string,from)

Description:

Extracts substrings from a string.

from indicates the start position of the extraction.

- If **from** starts at 0, the value **1** is used.
- If the value of **from** is positive, all characters from **from** to the end are extracted.
- If the value of **from** is negative, the last *n* characters in the string are extracted, and *n* indicates the absolute value of **from**.

Return type: text

Example:

If the value of **from** is positive:

```
openGauss=# SELECT substr('ABCDEF',2);
substr
-----
BCDEF
(1 row)
```

If the value of **from** is negative:

```
openGauss=# SELECT substr('ABCDEF',-2);
substr
-----
EF
(1 row)
```

- substr(string,from,count)

Description:

Extracts substrings from a string.

from indicates the start position of the extraction.

count indicates the length of the extracted substring.

- If **from** starts at 0, the value **1** is used.
- If the value of **from** is positive, extract **count** characters starting from **from**.
- If the value of **from** is negative, extract the last **n count** characters in the string, in which **n** indicates the absolute value of **from**.
- If the value of **count** is smaller than **1**, **null** is returned.

Return type: text

Example:

If the value of **from** is positive:

```
openGauss=# SELECT substr('ABCDEF',2,2);
substr
```

```
-----  
BC  
(1 row)
```

If the value of **from** is negative:

```
openGauss=# SELECT substr('ABCDEF',-3,2);  
substr  
-----  
DE  
(1 row)
```

- **substrb(string,from)**

Description: The functionality of this function is the same as that of SUBSTR(string,from). However, the calculation unit is byte.

Return type: text

Example:

```
openGauss=# SELECT substrb('ABCDEF',-2);  
substrb  
-----  
EF  
(1 row)
```

- **substrb(string,from,count)**

Description: The functionality of this function is the same as that of SUBSTR(string,from,count). However, the calculation unit is byte.

Return type: text

Example:

```
openGauss=# SELECT substrb('ABCDEF',2,2);  
substrb  
-----  
BC  
(1 row)
```

- **trim([leading |trailing |both] [characters] from string)**

Description: Removes the longest string containing only the characters (a space by default) from the start/end/both ends of the string.

Return type: text

Example:

```
openGauss=# SELECT trim(BOTH 'x' FROM 'xTomxx');  
btrim  
-----  
Tom  
(1 row)  
openGauss=# SELECT trim(LEADING 'x' FROM 'xTomxx');  
ltrim  
-----  
Tomxx  
(1 row)  
openGauss=# SELECT trim(TRAILING 'x' FROM 'xTomxx');  
rtrim  
-----  
xTom  
(1 row)
```

- **rtrim(string [, characters])**

Description: Removes the longest string containing only characters from characters (a space by default) from the end of string.

Return type: text

Example:

```
openGauss=# SELECT rtrim('TRIMxxxx','x');
rtrim
-----
TRIM
(1 row)
```

- `ltrim(string [, characters])`

Description: Removes the longest string containing only characters from characters (a space by default) from the start of string.

Return type: text

Example:

```
openGauss=# SELECT ltrim('xxxxTRIM','x');
ltrim
-----
TRIM
(1 row)
```

- `upper(string)`

Description: Converts the string into the uppercase.

Return type: text

Example:

```
openGauss=# SELECT upper('tom');
upper
-----
TOM
(1 row)
```

- `lower(string)`

Description: Converts the string into the lowercase.

Return type: text

Example:

```
openGauss=# SELECT lower('TOM');
lower
-----
tom
(1 row)
```

- `rpadd(string varchar, length int [, fill varchar])`

Description: Fills up **string** to **length** by appending the characters **fill** (a space by default). If **string** is already longer than **length**, then it is truncated.

length in GaussDB indicates the character length. One Chinese character is counted as one character.

Return type: text

Example:

```
openGauss=# SELECT rpadd('hi',5,'xyza');
rpadd
-----
hixyz
(1 row)
openGauss=# SELECT rpadd('hi',5,'abcdefg');
rpadd
-----
hiabc
(1 row)
```

- `instr(string,substring[,position,occurrence])`

Description: Queries and returns the value of the substring position that occurs the **occurrence** (1 by default) times from the **position** (1 by default) in the string.

- If the value of **position** is **0**, **0** is returned.
- If the value of **position** is negative, the search is performed backwards from the last n th character in the string, in which n indicates the absolute value of **position**.

In this function, the calculation unit is character. One Chinese character is one character.

Return type: integer

Example:

```
openGauss=# SELECT instr('corporate floor','or', 3);
instr
-----
      5
(1 row)
openGauss=# SELECT instr('corporate floor','or',-3,2);
instr
-----
      2
(1 row)
```

- **initcap(string)**

Description: Capitalizes the first letter of each word in a string.

Return type: text

Example:

```
openGauss=# SELECT initcap('hi THOMAS');
initcap
-----
Hi Thomas
(1 row)
```

- **ascii(string)**

Description: Indicates the ASCII code of the first character in the string.

Return type: integer

Example:

```
openGauss=# SELECT ascii('xyz');
ascii
-----
    120
(1 row)
```

- **replace(string varchar, search_string varchar, replacement_string varchar)**

Description: Replaces all **search_string** in the string with **replacement_string**.

Return type: text

Example:

```
openGauss=# SELECT replace('jack and jue','j','bl');
replace
-----
black and blue
(1 row)
```

- **lpad(string varchar, length int[, repeat_string varchar])**

Description: Adds a series of **repeat_string** (a space by default) on the left of the string to generate a new string with the total length of n .

If the length of the string is longer than the specified length, the function truncates the string and returns the substrings with the specified length.

Return type: varchar

Example:

```
openGauss=# SELECT lpad('PAGE 1',15,'*');
          lpad
-----
*****PAGE 1
(1 row)
openGauss=# SELECT lpad('hello world',5,'abcd');
          lpad
-----
hello
(1 row)
```

- **concat(str1,str2)**

Description: Concatenates **str1** and **str2** and returns the concatenated string. If **str1** or **str2** is set to **NULL**, **NULL** is returned. Note: **concat** calls the output function of the data type and the return value is uncertain. As a result, the optimizer cannot calculate the result in advance when generating a plan. If there are performance requirements, you are advised to use the operator **||**.

 NOTE

- If **sql_compatibility** is set to 'B' and **str1** or **str2** is set to **NULL**, the returned result is **NULL**.
- The return value of the concat function is of the variable-length type. When the concat function is compared with table data, the character string length is lost in the combination result. As a result, the comparison results are different.

Return type: text

Example:

```
openGauss=# SELECT concat('Hello', ' World!');
          concat
-----
Hello World!
(1 row)
openGauss=# SELECT concat('Hello', NULL);
          concat
-----
Hello
(1 row)
```

- **chr(integer)**

Description: Specifies the character of the ASCII code.

Return type: varchar

Example:

```
openGauss=# SELECT chr(65);
          chr
-----
A
(1 row)
```

- **concat_ws(sep text, str"any" [, str"any" [, ...]])**

Description: Uses the first parameter as the separator, which is associated with all following parameters. The **NULL** parameter is ignored.

NOTICE

- If the first parameter value is **NULL**, the returned result is **NULL**.
- If the first parameter is provided but the parameter value is an empty string ('') and the SQL compatibility mode of the database is set to **A**, the returned result is **NULL**. This is because the A-compatible mode treats the empty string ('') as **NULL**. To resolve this problem, you can change the SQL compatibility mode of the database to **B**, **C**, or **PG**.

Return type: text

Example:

```
openGauss=# SELECT concat_ws(',', 'ABCDE', 2, NULL, 22);
concat_ws
-----
ABCDE,2,22
(1 row)
```

- **nlssort**(string text, sort_method text)

Description: Returns the encoding value of a string in the sorting mode specified by **sort_method**. The encoding value can be used for sorting and determines the sequence of the string in the sorting mode. Currently, **sort_method** can be set to **nls_sort=schinese_pinyin_m** or **nls_sort=generic_m_ci**. **nls_sort=generic_m_ci** supports only the case-insensitive order for English characters.

String type: text

sort_method type: text

Return type: text

Example:

```
openGauss=# CREATE TABLE test(a text);

openGauss=# INSERT into test(a) values ('abC');

openGauss=# insert INTO test(a) VALUES ('abC');

openGauss=# INSERT INTO test(a) VALUES ('abc');

openGauss=# SELECT * FROM test ORDER BY nlssort(a,'nls_sort=schinese_pinyin_m');
a
-----
abc
abC
abC
(3 rows)

openGauss=# SELECT * FROM test ORDER BY nlssort(a, 'nls_sort=generic_m_ci');
a
-----
abC
abc
abC
(3 rows)

openGauss=# DROP TABLE test;
```

- **convert**(string bytea, src_encoding name, dest_encoding name)

Description: Converts the string to **dest_encoding**. **src_encoding** specifies the source code encoding. The string must be valid in this encoding.

Return type: bytea

Example:

```
openGauss=# SELECT convert('text_in_utf8', 'UTF8', 'GBK');
          convert
-----
\x746578745f696e5f75746638
(1 row)
```

 NOTE

If the rule for converting between source to target encoding (for example, GBK and LATIN1) does not exist, the string is returned without conversion. See the `pg_conversion` system catalog for details.

Example:

```
openGauss=# SHOW server_encoding;
 server_encoding
-----
LATIN1
(1 row)

openGauss=# SELECT convert_from('some text', 'GBK');
          convert_from
-----
some text
(1 row)

db_latin1=# SELECT convert_to('some text', 'GBK');
          convert_to
-----
\x736f6d652074657874
(1 row)

db_latin1=# SELECT convert('some text', 'GBK', 'LATIN1');
          convert
-----
\x736f6d652074657874
(1 row)
```

- `convert_from(string bytea, src_encoding name)`

Description: Converts a string using the coding mode of the database.

src_encoding specifies the source code encoding. The string must be valid in this encoding.

Return type: text

Example:

```
openGauss=# SELECT convert_from('text_in_utf8', 'UTF8');
          convert_from
-----
text_in_utf8
(1 row)
```

- `convert_to(string text, dest_encoding name)`

Description: Converts a string to **dest_encoding**.

Return type: bytea

Example:

```
openGauss=# SELECT convert_to('some text', 'UTF8');
          convert_to
-----
encode \x736f6d652074657874
(1 row)
```

- `format(formatstr text [, str"any" [, ...] ])`

Description: Formats a string.

Return type: text

Example:

```
openGauss=# SELECT format('Hello %s, %1$s', 'World');
format
-----
Hello World, World
(1 row)
```

- **md5(string)**

Description: Encrypts a string in MD5 mode and returns a value in hexadecimal form.

 NOTE

The MD5 encryption algorithm is not recommended because it has lower security and poses security risks.

Return type: text

Example:

```
openGauss=# SELECT md5('ABC');
md5
-----
902fbdd2b1df0c4f70b4a5d23525e932
(1 row)
```

- **decode(string text, format text)**

Description: Decodes binary data from textual representation.

Return type: bytea

Example:

```
openGauss=# SELECT decode('MTIzAAE=', 'base64');
decode
-----
\x3132330001
(1 row)
```

- **encode(data bytea, format text)**

Description: Encodes binary data into a textual representation.

Return type: text

Example:

```
openGauss=# SELECT encode(E'123\\000\\001', 'base64');
encode
-----
MTIzAAE=
(1 row)
```

 NOTE

- For a string containing newline characters, for example, a string consisting of a newline character and a space, the values of **length** and **lengthb** in GaussDB are 2.
- In GaussDB, *n* in the CHAR(*n*) type indicates the number of characters. Therefore, for multiple-octet coded character sets, the length returned by the LENGTHB function may be longer than *n*.
- GaussDB supports multiple types of databases, and the value can be **A**, **B**, **C**, or **PG**. If the database type is not specified, **A** is used by default. In this case, the lexical analyzer is different from that of the other three databases, which considers an empty character string as **NULL**. Therefore, when **A** is specified, if an empty string is used as a parameter in the preceding character operation function, no output is displayed. For example:
openGauss=# SELECT translate('12345','123','');
translate

```
-----  
(1 row)
```

This is because the kernel checks whether the input parameter contains **NULL** before calling the corresponding function. If the input parameter contains **NULL**, the kernel does not call the corresponding function. As a result, no output is displayed. In PG-compatible mode, the processing of character strings is the same as that of PG. Therefore, the preceding problem does not occur.

7.5.4 Binary String Functions and Operators

String Operators

SQL defines some string functions that use keywords, rather than commas, to separate arguments.

- **octet_length(string)**

Description: Specifies the number of bytes in a binary string.

Return type: int

Example:

```
openGauss=# SELECT octet_length(E'jo\000se'::bytea) AS RESULT;  
result  
-----  
5  
(1 row)
```

- **overlay(string placing string from int [for int])**

Description: Replaces substrings.

Return type: bytea

Example:

```
openGauss=# SELECT overlay(E'Th\000omas'::bytea placing E'\002\003'::bytea from 2 for 3) AS  
RESULT;  
result  
-----  
\x5402036d6173  
(1 row)
```

- **position(substring in string)**

Description: Specifies the location of a specified substring.

Return type: int

Example:

```
openGauss=# SELECT position(E'\000om'::bytea in E'Th\000omas'::bytea) AS RESULT;  
result
```

```
-----  
3  
(1 row)
```

- `substring(string [from int] [for int])`

Description: Truncates a substring.

Return type: `bytea`

Example:

```
openGauss=# SELECT substring(E'Th\000omas'::bytea from 2 for 3) AS RESULT;  
result  
-----  
\x68006f  
(1 row)
```

- `substr(string, from int [, for int])`

Description: Truncates a substring.

Return type: `bytea`

Example:

```
openGauss=# SELECT substr(E'Th\000omas'::bytea,2, 3) AS RESULT;  
result  
-----  
\x68006f  
(1 row)
```

- `trim([both] bytes from string)`

Description: Removes the longest string containing only bytes from **bytes** from the start and end of **string**.

Return type: `bytea`

Example:

```
openGauss=# SELECT trim(E'\000'::bytea from E'\000Tom\000'::bytea) AS RESULT;  
result  
-----  
\x546f6d  
(1 row)
```

Other Binary String Functions

GaussDB provides common syntax used for calling functions.

- `btrim(string bytea,bytes bytea)`

Description: Removes the longest string containing only bytes from **bytes** from the start and end of **string**.

Return type: `bytea`

Example:

```
openGauss=# SELECT btrim(E'\000trim\000'::bytea, E'\000'::bytea) AS RESULT;  
result  
-----  
\x7472696d  
(1 row)
```

- `get_bit(string, offset)`

Description: Extracts bits from a string.

Return type: `int`

Example:

```
openGauss=# SELECT get_bit(E'Th\000omas'::bytea, 45) AS RESULT;  
result
```

```
-----  
1  
(1 row)
```

- `get_byte(string, offset)`

Description: Extracts bytes from a string.

Return type: int

Example:

```
openGauss=# SELECT get_byte(E'Th\000omas'::bytea, 4) AS RESULT;  
result  
-----  
109  
(1 row)
```

- `rawcmp`

Description: Specifies the raw data type comparison function.

Parameter: raw, raw

Return type: integer

- `raweq`

Description: Specifies the raw data type comparison function.

Parameter: raw, raw

Return type: Boolean

- `rawge`

Description: Specifies the raw data type comparison function.

Parameter: raw, raw

Return type: Boolean

- `rawgt`

Description: Specifies the raw data type comparison function.

Parameter: raw, raw

Return type: Boolean

- `rawin`

Description: Specifies the raw data type parsing function.

Parameter:cstring

Return type: bytea

- `rawle`

Description: Specifies the raw data type parsing function.

Parameter: raw, raw

Return type: Boolean

- `rawlike`

Description: Specifies the raw data type parsing function.

Parameter: raw, raw

Return type: Boolean

- `rawlt`

Description: Specifies the raw data type parsing function.

Parameter: raw, raw

Return type: Boolean

- **rawne**
Description: Compares whether the raw types are the same.
Parameter: raw, raw
Return type: Boolean
- **rawnlike**
Description: Checks whether the raw type matches the mode.
Parameter: raw, raw
Return type: Boolean
- **rawout**
Description: Specifies the RAW output API.
Parameter: bytea
Return type: cstring
- **rawsend**
Description: Converts the bytea type to the binary type.
Parameter: raw
Return type: bytea
- **rawtohex**
Description: Converts the raw format to the hexadecimal format.
Parameter: text
Return type: text
- **set_bit(string,offset, newvalue)**
Description: Sets bits in a string.
Return type: bytea

Example:

```
openGauss=# SELECT set_bit(E'Th\\000omas'::bytea, 45, 0) AS RESULT;  
result  
-----  
\x5468006f6d4173  
(1 row)
```

- **set_byte(string,offset, newvalue)**
Description: Sets bytes in a string.
Return type: bytea

Example:

```
openGauss=# SELECT set_byte(E'Th\\000omas'::bytea, 4, 64) AS RESULT;  
result  
-----  
\x5468006f406173  
(1 row)
```

7.5.5 Bit String Functions and Operators

Bit String Operators

Aside from the usual comparison operators, the following operators can be used. Bit string operands of **&**, **|**, and **#** must be of equal length. In case of bit shifting, the original length of the string is preserved by zero padding (if necessary).

- ||

Description: Connects bit strings.

Example:

```
openGauss=# SELECT B'10001' || B'011' AS RESULT;
result
-----
10001011
(1 row)
```

 NOTE

- It is recommended that a column have no more than 180 consecutive internal joins. A column with over 180 joins will be split into joined consecutive strings.
Example: **str1||str2||str3||str4** is split into **(str1||str2)||(str3||str4)**.
- In A-compatible mode, if bit strings contain a null string, the null string is ignored and other strings are joined. In other compatibility modes, the null string is returned.
Take **str1||NULL||str2** as an example. **str1str2** is returned in A-compatible mode and **NULL** is returned in other compatibility modes.

- &

Description: Specifies the AND operation between bit strings.

Example:

```
openGauss=# SELECT B'10001' & B'01101' AS RESULT;
result
-----
00001
(1 row)
```

- |

Description: Specifies the OR operation between bit strings.

Example:

```
openGauss=# SELECT B'10001' | B'01101' AS RESULT;
result
-----
11101
(1 row)
```

- #

Description: Specifies the OR operation between bit strings if they are inconsistent. If the same positions in the two bit strings are both 1 or 0, the position returns 0.

Example:

```
openGauss=# SELECT B'10001' # B'01101' AS RESULT;
result
-----
11100
(1 row)
```

- ~

Description: Specifies the NOT operation between bit strings.

Example:

```
openGauss=# SELECT ~B'10001' AS RESULT;
result
-----
01110
(1 row)
```

- <<

Description: Shifts left in a bit string.

Example:

```
openGauss=# SELECT B'10001' << 3 AS RESULT;
result
-----
01000
(1 row)
```

- >>

Description: Shifts right in a bit string.

Example:

```
openGauss=# SELECT B'10001' >> 2 AS RESULT;
result
-----
00100
(1 row)
```

The following SQL-standard functions work on bit strings as well as strings: **length**, **bit_length**, **octet_length**, **position**, **substring**, and **overlay**.

The following functions work on bit strings as well as binary strings: **get_bit** and **set_bit**. When working with a bit string, these functions number the first (leftmost) bit of the string as bit 0.

In addition, it is possible to convert between integral values and type **bit**. Example:

```
openGauss=# SELECT 44::bit(10) AS RESULT;
result
-----
0000101100
(1 row)

openGauss=# SELECT 44::bit(3) AS RESULT;
result
-----
100
(1 row)

openGauss=# SELECT cast(-44 as bit(12)) AS RESULT;
result
-----
111111010100
(1 row)

openGauss=# SELECT '1110'::bit(4)::integer AS RESULT;
result
-----
14
(1 row)

openGauss=# select substring('10101111'::bit(8), 2);
substring
-----
0101111
(1 row)
```

NOTE

Casting to just "bit" means casting to bit(1), and so will deliver only the least significant bit of the integer.

7.5.6 Pattern Matching Operators

The database provides three independent methods for implementing pattern matching: SQL LIKE operator, SIMILAR TO operator, and POSIX-style regular expressions. Besides these basic operators, functions can be used to extract or replace matching substrings and to split a string at matching locations.

- LIKE

Description: Specifies whether the string matches the mode string following **LIKE**. The LIKE expression returns true if the string matches the supplied pattern. (As expected, the NOT LIKE expression returns false if LIKE returns true, and vice versa.)

Matching rules:

- This operator can succeed only when its pattern matches the entire string. If you want to match a sequence in any position within the string, the pattern must begin and end with a percent sign.
- The underscore (_) represents (matching) any single character. Percentage (%) indicates the wildcard character of any string.
- To match a literal underscore or percent sign without matching other characters, the respective character in **pattern** must be preceded by the escape character. The default escape character is the backslash but a different one can be selected by using the **ESCAPE** clause.
- To match with escape characters, enter two escape characters. For example, to write a **pattern** constant containing a backslash (\), you need to enter two backslashes in SQL statements.

 NOTE

When **standard_conforming_strings** is set to **off**, any backslashes you write in literal string constants will need to be doubled. So, writing a pattern that matches a single backslash actually involves writing four backslashes in the statement (you can avoid this by selecting a different escape character with **ESCAPE** so that the backslash is no longer a special character of **LIKE**. But the backslash is still the special character of the character text analyzer, so you still need two backslashes.)

In MySQL-compatible schema, it is also possible to select no escape character by writing **ESCAPE ''**. This effectively disables the escape mechanism, which makes it impossible to turn off the special meaning of underscore and percent signs in the schema.

- The keyword **ILIKE** can be used instead of **LIKE** to make the match case-insensitive.
- Operator **~~** is equivalent to **LIKE**, and operator **~~*** corresponds to **ILIKE**.

Example:

```
openGauss=# SELECT 'abc' LIKE 'abc' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' LIKE 'a%' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' LIKE '_b_' AS RESULT;
result
```

```
-----  
t  
(1 row)  
openGauss=# SELECT 'abc' LIKE 'c' AS RESULT;  
result  
-----  
f  
(1 row)
```

- **SIMILAR TO**

Description: Returns **true** or **false** depending on whether the pattern matches the given string. It is similar to LIKE, but differs in that it uses the regular expression understanding pattern defined by the SQL standard.

Matching rules:

- Similar to LIKE, this operator succeeds only when its pattern matches the entire string. If you want to match a sequence in any position within the string, the pattern must begin and end with a percent sign.
- The underscore (_) represents (matching) any single character. Percentage (%) indicates the wildcard character of any string.
- SIMILAR TO supports these pattern-matching metacharacters borrowed from POSIX regular expressions:

Metacharacter	Description
	Specifies alternation (either of two alternatives).
*	Specifies repetition of the previous item zero or more times.
+	Specifies repetition of the previous item one or more times.
?	Specifies repetition of the previous item zero or one time.
{m}	Specifies repetition of the previous item exactly <i>m</i> times.
{m,}	Specifies repetition of the previous item <i>m</i> or more times.
{m,n}	Specifies repetition of the previous item at least <i>m</i> times and does not exceed <i>n</i> times.
()	Specifies that parentheses () can be used to group items into a single logical item.
[...]	Specifies a character class, just as in POSIX regular expressions.

- A preamble escape character disables the special meaning of any of these metacharacters. The rules for using escape characters are the same as those for LIKE.

Regular expressions:

The `substring(string from pattern for escape)` function extracts a substring that matches an SQL regular expression pattern.

Example:

```
openGauss=# SELECT 'abc' SIMILAR TO 'abc' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' SIMILAR TO 'a' AS RESULT;
result
-----
f
(1 row)
openGauss=# SELECT 'abc' SIMILAR TO '%(b|d)%' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' SIMILAR TO '(b|c)%' AS RESULT;
result
-----
f
(1 row)
```

- POSIX regular expressions

Description: A regular expression is a collation that is an abbreviated definition of a set of strings (a regular set). If a string is a member of a regular set described by a regular expression, the string matches the regular expression. POSIX regular expressions provide a more powerful means for pattern matching than the LIKE and SIMILAR TO operators. [Table 7-27](#) lists all available operators for pattern matching using POSIX regular expressions.

Table 7-27 Regular expression match operators

Operator	Description	Example
~	Matches a regular expression, which is case-sensitive.	'thomas' ~ '.*thomas.*'
~*	Matches a regular expression, which is case-insensitive.	'thomas' ~* '.*Thomas.*'
!~	Does not match a regular expression, which is case-sensitive.	'thomas' !~ '.*Thomas.*'
!~*	Does not match a regular expression, which is case-insensitive.	'thomas' !~* '.*vadim.*'

Matching rules:

- a. Unlike LIKE patterns, a regular expression is allowed to match anywhere within a string, unless the regular expression is explicitly anchored to the beginning or end of the string.
- b. Besides the metacharacters mentioned above, POSIX regular expressions also support the following pattern matching metacharacters:

Metacharacter	Description
^	Specifies the match starting with a string.
\$	Specifies the match at the end of a string.
.	Matches any single character.

Regular expressions:

POSIX-style regular expressions support the following functions:

- The `substring(string from pattern)` function provides a method for extracting a substring that matches the POSIX-style regular expression pattern.
- The `regexp_count(stringtext,patterntext[,positionint[,flagstext]])` function counts the number of substrings that match the POSIX regular expression pattern.
- The `regexp_instr(stringtext,patterntext[,positionint[,occurrenceint[,return_optint[,flagstext]]]])` function obtains the position of a substring that matches the POSIX regular expression pattern.
- The `regexp_substr(string text, pattern text [, position int [, occurrence int [, flags text]])` function provides a method to extract a substring that matches a POSIX-style regular expression pattern.
- The `regexp_replace(string, pattern, replacement [,flags ])` function replaces the substring that matches the POSIX-style regular expression pattern with the new text.
- The `regexp_matches(string text, pattern text [, flags text])` function returns a text array consisting of all captured substrings that match a POSIX-style regular expression pattern.
- The `regexp_split_to_table(string text, pattern text [, flags text])` function splits a string using a POSIX-style regular expression pattern as a delimiter.
- The `regexp_split_to_array(string text, pattern text [, flags text ])` function behaves the same as `regexp_split_to_table`, except that `regexp_split_to_array` returns its result as an array of text.

NOTE

The regular expression split functions ignore zero-length matches, which occur at the beginning or end of a string or after the previous match. This is contrary to the strict definition of regular expression matching. The latter is implemented by `regexp_matches`, but the former is usually the most commonly used behavior in practice.

Example:

```
openGauss=# SELECT 'abc' ~ 'Abc' AS RESULT;
result
-----
f
(1 row)
openGauss=# SELECT 'abc' ~* 'Abc' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' !~ 'Abc' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' !~* 'Abc' AS RESULT;
result
-----
f
(1 row)
openGauss=# SELECT 'abc' ~ '^a' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' ~ '(b|d)' AS RESULT;
result
-----
t
(1 row)
openGauss=# SELECT 'abc' ~ '^ (b|c)' AS RESULT;
result
-----
f
(1 row)
```

Although most regular expression searches can be performed quickly, regular expressions can still be artificially processed to require any length of time and any amount of memory. It is not recommended that you accept the regular expression search pattern from the non-security pattern source. If you must do this, you are advised to add the statement timeout limit. The search with the SIMILAR TO mode has the same security risks as the SIMILAR TO provides many capabilities that are the same as those of the POSIX- style regular expression. The LIKE search is much simpler than the other two options. Therefore, it is more secure to accept the non-secure mode source search.

7.5.7 Mathematical Functions and Operators

Numeric Operators

- +

Description: Addition

Example:

```
openGauss=# SELECT 2+3 AS RESULT;
result
-----
5
(1 row)
```

- -

Description: Subtraction

Example:

```
openGauss=# SELECT 2-3 AS RESULT;
result
-----
-1
(1 row)
```

- *

Description: Multiplication

Example:

```
openGauss=# SELECT 2*3 AS RESULT;
result
-----
6
(1 row)
```

- /

Description: Division (The result is not rounded.)

Example:

```
openGauss=# SELECT 4/2 AS RESULT;
result
-----
2
(1 row)
openGauss=# SELECT 4/3 AS RESULT;
result
-----
1.3333333333333333
(1 row)
```

- +/-

Description: Positive/Negative

Example:

```
openGauss=# SELECT -2 AS RESULT;
result
-----
-2
(1 row)
```

- %

Description: Modulo (to obtain the remainder)

Example:

```
openGauss=# SELECT 5%4 AS RESULT;
result
-----
1
(1 row)
```

- @

Description: Absolute value

Example:

```
openGauss=# SELECT @ -5.0 AS RESULT;
result
-----
5.0
(1 row)
```

- ^

Description: Power (exponent calculation)

Example:

```
openGauss=# SELECT 2.0^3.0 AS RESULT;
      result
-----
8.0000000000000000
(1 row)
```

- $\sqrt{\quad}$

Description: Square root

Example:

```
openGauss=# SELECT  $\sqrt{25.0}$  AS RESULT;
      result
-----
5
(1 row)
```

- $\sqrt[3]{\quad}$

Description: Cubic root

Example:

```
openGauss=# SELECT  $\sqrt[3]{27.0}$  AS RESULT;
      result
-----
3
(1 row)
```

- $!$

Description: Factorial

Example:

```
openGauss=# SELECT 5! AS RESULT;
      result
-----
120
(1 row)
```

- $!!$

Description: Factorial (prefix operator)

Example:

```
openGauss=# SELECT !!5 AS RESULT;
      result
-----
120
(1 row)
```

- $\&$

Description: Binary AND

Example:

```
openGauss=# SELECT 91&15 AS RESULT;
      result
-----
11
(1 row)
```

- $|$

Description: Binary OR

Example:

```
openGauss=# SELECT 32|3 AS RESULT;
      result
-----
35
(1 row)
```

- #

Description: Binary XOR

Example:

```
openGauss=# SELECT 17#5 AS RESULT;
result
-----
      20
(1 row)
```

- ~

Description: Binary NOT

Example:

```
openGauss=# SELECT ~1 AS RESULT;
result
-----
      -2
(1 row)
```

- <<

Description: Binary shift left

Example:

```
openGauss=# SELECT 1<<4 AS RESULT;
result
-----
      16
(1 row)
```

- >>

Description: Binary shift right

Example:

```
openGauss=# SELECT 8>>2 AS RESULT;
result
-----
       2
(1 row)
```

Numeric Operation Functions

- abs(x)

Description: Absolute value

Return type: same as the input

Example:

```
openGauss=# SELECT abs(-17.4);
abs
-----
 17.4
(1 row)
```

- acos(x)

Description: Arc cosine

Return type: double precision

Example:

```
openGauss=# SELECT acos(-1);
acos
-----
3.14159265358979
(1 row)
```


- **asin(x)**
Description: Arc sine
Return type: double precision
Example:

```
openGauss=# SELECT asin(0.5);
      asin
-----
.523598775598299
(1 row)
```
- **atan(x)**
Description: Arc tangent
Return type: double precision
Example:

```
openGauss=# SELECT atan(1);
      atan
-----
.785398163397448
(1 row)
```
- **atan2(y, x)**
Description: Arc tangent of y/x
Return type: double precision
Example:

```
openGauss=# SELECT atan2(2, 1);
      atan2
-----
1.10714871779409
(1 row)
```
- **bitand(integer, integer)**
Description: Performs an AND (&) operation on two integers.
Return type: bigint
Example:

```
openGauss=# SELECT bitand(127, 63);
      bitand
-----
        63
(1 row)
```
- **cbrt(dp)**
Description: Cubic root
Return type: double precision
Example:

```
openGauss=# SELECT cbrt(27.0);
      cbrt
-----
        3
(1 row)
```
- **ceil(x)**
Description: Minimum integer greater than or equal to the parameter
Return type: integer
Example:

```
openGauss=# SELECT ceil(-42.8);
      ceil
-----
```

```
-----  
-42  
(1 row)
```

- **ceiling(dp or numeric)**

Description: Minimum integer (alias of ceil) greater than or equal to the parameter

Return type: dp or numeric. If implicit type conversion is not considered, the return type is the same as the input type.

Example:

```
openGauss=# SELECT ceiling(-95.3);  
ceiling  
-----  
-95  
(1 row)
```

- **cos(x)**

Description: Cosine

Return type: double precision

Example:

```
openGauss=# SELECT cos(-3.1415927);  
cos  
-----  
-.9999999999999999  
(1 row)
```

- **cot(x)**

Description: Cotangent

Return type: double precision

Example:

```
openGauss=# SELECT cot(1);  
cot  
-----  
.642092615934331  
(1 row)
```

- **degrees(dp)**

Description: Converts radians to angles.

Return type: double precision

Example:

```
openGauss=# SELECT degrees(0.5);  
degrees  
-----  
28.6478897565412  
(1 row)
```

- **div(y numeric, x numeric)**

Description: Integer part of y/x

Return type: numeric

Example:

```
openGauss=# SELECT div(9,4);  
div  
-----  
2  
(1 row)
```

- **exp(x)**

Description: Natural exponent

Return type: dp or numeric. If implicit type conversion is not considered, the return type is the same as the input type.

Example:

```
openGauss=# SELECT exp(1.0);
exp
-----
2.7182818284590452
(1 row)
```

- floor(x)

Description: Maximum integer not larger than the parameter

Return type: same as the input

Example:

```
openGauss=# SELECT floor(-42.8);
floor
-----
-43
(1 row)
```

- int1(in)

Description: Converts the input text parameter to a value of the int1 type and returns the value.

Return type: int1

Example:

```
openGauss=# SELECT int1('123');
int1
-----
123
(1 row)
openGauss=# SELECT int1('a');
int1
-----
0
(1 row)
```

- int2(in)

Description: Converts the input parameter to a value of the int2 type and returns the value.

The supported input parameter types include float4, float8, int16, numeric, and text.

Return type: int2

Example:

```
openGauss=# SELECT int2('1234');
int2
-----
1234
(1 row)
openGauss=# SELECT int2(25.3);
int2
-----
25
(1 row)
```

- int4(in)

Description: Converts the input parameter to a value of the int4 type and returns the value.

The supported input parameter types include bit, Boolean, char, double precision, int16, numeric, real, smallint, and text.

Return type: int4

Example:

```
openGauss=# SELECT int4('789');
int4
-----
789
(1 row)

openGauss=# SELECT int4(99.9);
int4
-----
99
(1 row)
```

- float4(in)

Description: Converts the input parameter to a value of the float4 type and returns the value. The supported input parameter types include bigint, double precision, int16, integer, numeric, smallint, and text.

Return type: float4

Example:

```
openGauss=# SELECT float4('789');
float4
-----
789
(1 row)

openGauss=# SELECT float4(99.9);
float4
-----
99.9
(1 row)
```

- float8(in)

Description: Converts the input parameter to a value of the float8 type and returns the value. The supported input parameter types include bigint, int16, integer, numeric, real, smallint, and text.

Return type: float8

Example:

```
openGauss=# SELECT float8('789');
float8
-----
789
(1 row)

openGauss=# SELECT float8(99.9);
float8
-----
99.9
(1 row)
```

- int16(in)

Description: Converts the input parameter to a value of the int16 type and returns the value. The supported input parameter types include bigint, Boolean, double precision, integer, numeric, oid, real, smallint, and tinyint.

Return type: int16

Example:

```
openGauss=# SELECT int16('789');
int16
-----
789
(1 row)
```

```
openGauss=# SELECT int16(99.9);
int16
-----
  99
(1 row)
```

- **numeric(in)**

Description: Converts the input parameter to a value of the numeric type and returns the value. The supported input parameter types include bigint, Boolean, double precision, int16, integer, money, real, and smallint.

Return type: numeric

Example:

```
openGauss=# SELECT "numeric"('789');
numeric
-----
   789
(1 row)
```

```
openGauss=# SELECT "numeric"(99.9);
numeric
-----
   99.9
(1 row)
```

- **oid(in)**

Description: Converts the input parameter to a value of the oid type and returns the value. The supported input parameter types include bigint and int16.

Return type: oid

- **radians(dp)**

Description: Converts angles to radians.

Return type: double precision

Example:

```
openGauss=# SELECT radians(45.0);
radians
-----
.785398163397448
(1 row)
```

- **random()**

Description: Random number between 0.0 and 1.0

Return type: double precision

Example:

```
openGauss=# SELECT random();
random
-----
.824823560658842
(1 row)
```

- **multiply(x double precision or text, y double precision or text)**

Description: Product of x and y.

Return type: double precision

Example:

```
openGauss=# SELECT multiply(9.0, '3.0');
multiply
-----
```

```
      27
(1 row)
openGauss=# SELECT multiply('9.0', 3.0);
      multiply
-----
      27
(1 row)
```

- **ln(x)**

Description: Natural logarithm

Return type: dp or numeric. If implicit type conversion is not considered, the return type is the same as the input type.

Example:

```
openGauss=# SELECT ln(2.0);
      ln
-----
.6931471805599453
(1 row)
```

- **log(x)**

Description: Logarithm with 10 as the base

Return type: same as the input

Example:

```
openGauss=# SELECT log(100.0);
      log
-----
2.0000000000000000
(1 row)
```

- **log(b numeric, x numeric)**

Description: Logarithm with b as the base

Return type: numeric

Example:

```
openGauss=# SELECT log(2.0, 64.0);
      log
-----
6.0000000000000000
(1 row)
```

- **mod(x,y)**

Description: Remainder of x/y (model) If x equals to 0, 0 is returned.

Return type: same as the parameter type

Example:

```
openGauss=# SELECT mod(9,4);
      mod
-----
      1
(1 row)
openGauss=# SELECT mod(9,0);
      mod
-----
      9
(1 row)
```

- **pi()**

Description: π constant value

Return type: double precision

Example:

```
openGauss=# SELECT pi();
      pi
-----
3.14159265358979
(1 row)
```

- **power(a double precision, b double precision)**

Description: b power of a

Return type: double precision

Example:

```
openGauss=# SELECT power(9.0, 3.0);
      power
-----
729.0000000000000000
(1 row)
```

- **round(x)**

Description: Integer closest to the input parameter

Return type: same as the input

Example:

```
openGauss=# SELECT round(42.4);
      round
-----
42
(1 row)

openGauss=# SELECT round(42.6);
      round
-----
43
(1 row)
```

- **round(v numeric, s int)**

Description: s digits are kept after the decimal point.

Return type: numeric

Example:

```
openGauss=# SELECT round(42.4382, 2);
      round
-----
42.44
(1 row)
```

- **setseed(dp)**

Description: Sets seed for the following random() calling (between -1.0 and 1.0, inclusive).

Return type: void

Example:

```
openGauss=# SELECT setseed(0.54823);
      setseed
-----
(1 row)
```

- **sign(x)**

Description: Returns symbols of this parameter.

Return type: -1 indicates minus. 0 indicates 0, and 1 indicates positive numbers.

Example:

```
openGauss=# SELECT sign(-8.4);
sign
-----
-1
(1 row)
```

- **sin(x)**

Description: Sine

Return type: double precision

Example:

```
openGauss=# SELECT sin(1.57079);
sin
-----
.999999999979986
(1 row)
```

- **sqrt(x)**

Description: Square root

Return type: dp or numeric. If implicit type conversion is not considered, the return type is the same as the input type.

Example:

```
openGauss=# SELECT sqrt(2.0);
sqrt
-----
1.414213562373095
(1 row)
```

- **tan(x)**

Description: Tangent

Return type: double precision

Example:

```
openGauss=# SELECT tan(20);
tan
-----
2.23716094422474
(1 row)
```

- **trunc(x)**

Description: Truncates (the integral part).

Return type: same as the input

Example:

```
openGauss=# SELECT trunc(42.8);
trunc
-----
42
(1 row)
```

- **trunc(v numeric, s int)**

Description: Truncates a number with **s** digits after the decimal point.

Return type: numeric

Example:

```
openGauss=# SELECT trunc(42.4382, 2);
trunc
-----
42.43
(1 row)
```

- **smgrne(a smgr, b smgr)**

Description: Compares two integers of the smgr type to check whether they are different.

Return type: Boolean

- smgreq(a smgr, b smgr)

Description: Compares two integers of the smgr type to check whether they are equivalent.

Return type: Boolean

- int1abs

Description: Returns the absolute value of data of the uint8 type.

Parameter: tinyint

Return type: tinyint

- int1and

Description: Returns the bitwise AND result of two data records of the uint8 type.

Parameter: tinyint, tinyint

Return type: tinyint

- int1cmp

Description: Returns the comparison result of two data records of the uint8 type. If the value of the first parameter is greater, **1** is returned. If the value of the second parameter is greater, **-1** is returned. If they are the same, **0** is returned.

Parameter: tinyint, tinyint

Return type: integer

- int1div

Description: Returns the result of dividing two data records of the uint8 type. The result is of the float8 type.

Parameter: tinyint, tinyint

Return type: tinyint

- int1eq

Description: Compares two pieces of data of the uint8 type to check whether they are the same.

Parameter: tinyint, tinyint

Return type: Boolean

- int1ge

Description: Determines whether the value of the first parameter is greater than or equal to the value of the second parameter in two data records of the uint8 type.

Parameter: tinyint, tinyint

Return type: Boolean

- int1gt

Description: Performs a greater-than operation on an unsigned 1-byte integer.

Parameter: tinyint, tinyint

Return type: Boolean

- **int1larger**
Description: Returns the maximum value of an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1le**
Description: Performs a less-than or an equal-to operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: Boolean
- **int1lt**
Description: Performs a less-than operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: Boolean
- **int1smaller**
Description: Calculates the minimum value of an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1inc**
Description: Unsigned 1-byte integer plus 1.
Parameter: tinyint
Return type: tinyint
- **int1mi**
Description: Performs a minus operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1mod**
Description: Performs a MOD operation on unsigned 1-byte integers.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1mul**
Description: Performs a multiplication operation on unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1ne**
Description: Performs a not-equal-to operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: Boolean
- **int1pl**
Description: Performs an addition operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint

- **int1um**
Description: Returns an unsigned 2-byte integer after subtracting the opposite number from the unsigned 1-byte integer.
Parameter: tinyint
Return type: smallint
- **int1xor**
Description: Performs an exclusive OR operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **cash_div_int1**
Description: Performs a division operation on the money type.
Parameter: money, tinyint
Return type: money
- **cash_mul_int1**
Description: Performs a multiplication operation on the money type.
Parameter: money, tinyint
Return type: money
- **int1not**
Description: Reverts binary bits of an unsigned 1-byte integer.
Parameter: tinyint
Return type: tinyint
- **int1or**
Description: Performs an OR operation on an unsigned 1-byte integer.
Parameter: tinyint, tinyint
Return type: tinyint
- **int1shl**
Description: Shifts an unsigned 1-byte integer leftwards by a specified number of bits.
Parameter: tinyint, integer
Return type: tinyint
- **int1shr**
Description: Shifts an unsigned 1-byte integer rightwards by a specified number of bits.
Parameter: tinyint, integer
Return type: tinyint
- **width_bucket(op numeric, b1 numeric, b2 numeric, count int)**
Description: Returns a bucket to which the operand will be assigned in an equi-depth histogram with **count** buckets, ranging from **b1** to **b2**.
Return type: int
Example:

```
openGauss=# SELECT width_bucket(5.35, 0.024, 10.06, 5);
width_bucket
```

```
-----
      3
(1 row)
```

- width_bucket(op dp, b1 dp, b2 dp, count int)
Description: Returns a bucket to which the operand will be assigned in an equi-depth histogram with **count** buckets, ranging from **b1** to **b2**.
Return type: int
Example:
openGauss=# SELECT width_bucket(5.35, 0.024, 10.06, 5);
width_bucket

 3
(1 row)

7.5.8 Date and Time Processing Functions and Operators

Date and Time Operators



WARNING

When the user uses date/time operators, explicit type prefixes are modified for corresponding operands to ensure that the operands parsed by the database are consistent with what the user expects, and no unexpected results occur.

For example, abnormal mistakes will occur in the following example without an explicit data type.

```
openGauss=# SELECT date '2001-10-01' - '7' AS RESULT;
ERROR:
GAUSS-10416: invalid input syntax for type timestamp: "7"
SQLSTATE: 22007
LINE 1: SELECT date '2001-10-01' - '7' AS RESULT;
           ^
CONTEXT: referenced column: result
```

Table 7-28 Time and date operators

Ope rato r	Example
+	<pre>openGauss=# SELECT date '2001-9-28' + integer '7' AS RESULT; result ----- 2001-10-05 (1 row)</pre>
	<pre>openGauss=# SELECT date '2001-09-28' + interval '1 hour' AS RESULT; result ----- 2001-09-28 01:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT date '2001-09-28' + time '03:00' AS RESULT; result ----- 2001-09-28 03:00:00 (1 row)</pre>

Operator	Example
	<pre>openGauss=# SELECT interval '1 day' + interval '1 hour' AS RESULT; result ----- 1 day 01:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT timestamp '2001-09-28 01:00' + interval '23 hours' AS RESULT; result ----- 2001-09-29 00:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT time '01:00' + interval '3 hours' AS RESULT; result ----- 04:00:00 (1 row)</pre>
-	<pre>openGauss=# SELECT date '2001-10-01' - date '2001-09-28' AS RESULT; result ----- 3days (1 row)</pre>
	<pre>openGauss=# SELECT date '2001-10-01' - integer '7' AS RESULT; result ----- 2001-09-24 00:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT date '2001-09-28' - interval '1 hour' AS RESULT; result ----- 2001-09-27 23:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT time '05:00' - time '03:00' AS RESULT; result ----- 02:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT time '05:00' - interval '2 hours' AS RESULT; result ----- 03:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT timestamp '2001-09-28 23:00' - interval '23 hours' AS RESULT; result ----- 2001-09-28 00:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT interval '1 day' - interval '1 hour' AS RESULT; result ----- 23:00:00 (1 row)</pre>
	<pre>openGauss=# SELECT timestamp '2001-09-29 03:00' - timestamp '2001-09-27 12:00' AS RESULT; result ----- 1 day 15:00:00 (1 row)</pre>

Operator	Example
*	<pre>openGauss=# SELECT 900 * interval '1 second' AS RESULT; result ----- 00:15:00 (1 row)</pre>
	<pre>openGauss=# SELECT 21 * interval '1 day' AS RESULT; result ----- 21 days (1 row)</pre>
	<pre>openGauss=# SELECT double precision '3.5' * interval '1 hour' AS RESULT; result ----- 03:30:00 (1 row)</pre>
/	<pre>openGauss=# SELECT interval '1 hour' / double precision '1.5' AS RESULT; result ----- 00:40:00 (1 row)</pre>

Time/Date Functions

- age(timestamp, timestamp)

Description: Subtracts parameters, producing a result in YYYY-MM-DD format. If the result is negative, the returned result is also negative. The input parameters can contain timezone or not.

Return type: interval

Example:

```
openGauss=# SELECT age(timestamp '2001-04-10', timestamp '1957-06-13');
age
-----
43 years 9 mons 27 days
(1 row)
```

- age(timestamp)

Description: Minuses the current time with the parameter. The input parameter can contain timezone or not.

Return type: interval

Example:

```
openGauss=# SELECT age(timestamp '1957-06-13');
age
-----
60 years 2 mons 18 days
(1 row)
```

- clock_timestamp()

Description: Specifies the current timestamp of the real-time clock.

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT clock_timestamp();
               clock_timestamp
-----
2017-09-01 16:57:36.636205+08
(1 row)
```

- **current_date**

Description: Specifies the current date.

Return type: date

Example:

```
openGauss=# SELECT current_date;
               date
-----
2017-09-01
(1 row)
```

- **current_time**

Description: Specifies the current time.

Return type: time with time zone

Example:

```
openGauss=# SELECT current_time;
               timetz
-----
16:58:07.086215+08
(1 row)
```

- **current_timestamp**

Description: Specifies the current date and time.

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT current_timestamp;
               pg_timestamp
-----
2017-09-01 16:58:19.22173+08
(1 row)
```

- **pg_timestamp()**

Description: Current date and time (start of the current statement).

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT pg_timestamp();
               pg_timestamp
-----
2015-10-14 11:21:28.317367+08
(1 row)
```

- **date_part(text, timestamp)**

Description: Obtains the value of a subdomain in date or time, for example, the year or hour. It is equivalent to **extract(field from timestamp)**.

Timestamp types: abstime, date, interval, reltime, time with time zone, time without time zone, timestamp with time zone, timestamp without time zone

Return type: double precision

Example:

```
openGauss=# SELECT date_part('hour', timestamp '2001-02-16 20:38:40');
               date_part
-----
```

```
20
(1 row)
```

- `timestamp_diff(text, timestamp, timestamp)`

Description: Calculates the difference between two timestamps and truncates the difference to the precision specified by text. In B-compatible database, this function is the same as **`TIMESTAMPDIFF(unit , timestamp_expr1, timestamp_expr2)`**.

Return type: int64

Example:

```
openGauss=# SELECT timestamp_diff('year','2018-01-01','2020-04-01');
timestamp_diff
-----
2
(1 row)
openGauss=# SELECT timestamp_diff('month','2018-01-01','2020-04-01');
timestamp_diff
-----
27
(1 row)
openGauss=# SELECT timestamp_diff('quarter','2018-01-01','2020-04-01');
timestamp_diff
-----
9
(1 row)
openGauss=# SELECT timestamp_diff('week','2018-01-01','2020-04-01');
timestamp_diff
-----
117
(1 row)
openGauss=# SELECT timestamp_diff('day','2018-01-01','2020-04-01');
timestamp_diff
-----
821
(1 row)
openGauss=# SELECT timestamp_diff('hour','2018-01-01 10:10:10','2018-01-01 12:12:12');
timestamp_diff
-----
2
(1 row)
openGauss=# SELECT timestamp_diff('minute','2018-01-01 10:10:10','2018-01-01 12:12:12');
timestamp_diff
-----
122
(1 row)
openGauss=# SELECT timestamp_diff('second','2018-01-01 10:10:10','2018-01-01 10:12:12');
timestamp_diff
-----
122
(1 row)
openGauss=# SELECT timestamp_diff('microsecond','2018-01-01 10:10:10','2018-01-01 10:12:12');
timestamp_diff
-----
122000000
(1 row)
```

- `date_part(text, interval)`

Description: Obtains the subdomain value of the date/time value. When obtaining the month value, if the value is greater than 12, obtain the remainder after it is divided by 12. It is equivalent to **`extract(field from timestamp)`**.

Return type: double precision

Example:


```
openGauss=# SELECT date_part('month', interval '2 years 3 months');
date_part
-----
3
(1 row)
```

- `date_trunc(text, timestamp)`

Description: Truncates to the precision specified by **text**.

Return type: interval, timestamp with time zone, timestamp without time zone

Example:

```
openGauss=# SELECT date_trunc('hour', timestamp '2001-02-16 20:38:40');
date_trunc
-----
2001-02-16 20:00:00
(1 row)
```

- `trunc(timestamp)`

Description: Truncates to day by default.

Example:

```
openGauss=# SELECT trunc(timestamp '2001-02-16
20:38:40');
trunc
-----
2001-02-16 00:00:00
(1 row)
```

- `trunc(arg1, arg2)`

Description: Truncates to the precision specified by **arg2**.

Type of **arg1**: interval, timestamp with time zone, timestamp without time zone

Type of **arg2**: text

Return type: interval, timestamp with time zone, timestamp without time zone

Example:

```
openGauss=# SELECT trunc(timestamp '2001-02-16 20:38:40',
'hour');
trunc
-----
2001-02-16 20:00:00
(1 row)
```

- `daterange(arg1, arg2)`

Description: Obtains time boundary information. The type of **arg1** and **arg2** is **date**.

Return type: daterange

Example:

```
openGauss=# SELECT daterange('2000-05-06','2000-08-08');
daterange
-----
[2000-05-06,2000-08-08)
(1 row)
```

- `daterange(arg1, arg2, text)`

Description: Obtains time boundary information. The type of **arg1** and **arg2** is **date**, and the type of **text** is **text**.

Return type: daterange

Example:

```
openGauss=# SELECT daterange('2000-05-06','2000-08-08','[]');
daterange
-----
[2000-05-06,2000-08-09)
(1 row)
```

- `extract(field from timestamp)`

Description: Obtains the hour.

Return type: double precision

Example:

```
openGauss=# SELECT extract(hour from timestamp '2001-02-16 20:38:40');
date_part
-----
20
(1 row)
```

- `extract(field from interval)`

Description: Obtains the month. If the value is greater than 12, obtain the remainder after it is divided by 12.

Return type: double precision

Example:

```
openGauss=# SELECT extract(month from interval '2 years 3 months');
date_part
-----
3
(1 row)
```

- `isfinite(date)`

Description: Tests for a valid date.

Return type: Boolean

Example:

```
openGauss=# SELECT isfinite(date '2001-02-16');
isfinite
-----
t
(1 row)
```

- `isfinite(timestamp)`

Description: Tests for a valid timestamp.

Return type: Boolean

Example:

```
openGauss=# SELECT isfinite(timestamp '2001-02-16 21:28:30');
isfinite
-----
t
(1 row)
```

- `isfinite(interval)`

Description: Tests for a valid interval.

Return type: Boolean

Example:

```
openGauss=# SELECT isfinite(interval '4 hours');
isfinite
-----
t
(1 row)
```

- **justify_days(interval)**
Description: Adjusts intervals to 30-day time periods, which are represented as months.
Return type: interval
Example:

```
openGauss=# SELECT justify_days(interval '35 days');
justify_days
-----
1 mon 5 days
(1 row)
```
- **justify_hours(interval)**
Description: Sets the time interval in days (24 hours is one day).
Return type: interval
Example:

```
openGauss=# SELECT JUSTIFY_HOURS(INTERVAL '27 HOURS');
justify_hours
-----
1 day 03:00:00
(1 row)
```
- **justify_interval(interval)**
Description: Adjusts **interval** using **justify_days** and **justify_hours**.
Return type: interval
Example:

```
openGauss=# SELECT JUSTIFY_INTERVAL(INTERVAL '1 MON -1 HOUR');
justify_interval
-----
29 days 23:00:00
(1 row)
```
- **localtime**
Description: Specifies the current time.
Return type: time
Example:

```
openGauss=# SELECT localtime AS RESULT;
result
-----
16:05:55.664681
(1 row)
```
- **localtimestamp**
Description: Specifies the current date and time.
Return type: timestamp
Example:

```
openGauss=# SELECT localtimestamp;
timestamp
-----
2017-09-01 17:03:30.781902
(1 row)
```
- **now()**
Description: Specifies the current date and time.
Return type: timestamp with time zone
Example:

```
openGauss=# SELECT now();
          now
-----
2017-09-01 17:03:42.549426+08
(1 row)
```

- `timenow()`

Description: Specifies the current date and time.

Return type: timestamp with time zone

Example:

```
openGauss=# select timenow();
          timenow
-----
2020-06-23 20:36:56+08
(1 row)
```

- `numtodsinterval(num, interval_unit)`

Description: Converts a number to the interval type. **num** is a numeric-typed number. **interval_unit** is a string in the following format: 'DAY' | 'HOUR' | 'MINUTE' | 'SECOND'

You can set the GUC parameter **IntervalStyle** to **a** to be compatible with the interval output format of the function.

Example:

```
openGauss=# SELECT numtodsinterval(100, 'HOUR');
 numtodsinterval
-----
100:00:00
(1 row)

openGauss=# SET intervalstyle = a;
SET
openGauss=# SELECT numtodsinterval(100, 'HOUR');
 numtodsinterval
-----
+0000000004 04:00:00.000000000
(1 row)
```

- `pg_sleep(seconds)`

Description: Specifies the delay time of the server thread in unit of second.

Return type: void

Example:

```
openGauss=# SELECT pg_sleep(10);
 pg_sleep
-----
(1 row)
```

- `statement_timestamp()`

Description: Current date and time (start of the current statement).

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT statement_timestamp();
 statement_timestamp
-----
2017-09-01 17:04:39.119267+08
(1 row)
```

- `sysdate`

Description: Specifies the current date and time.

Return type: timestamp

Example:

```
openGauss=# SELECT sysdate;
           sysdate
-----
2017-09-01 17:04:49
(1 row)
```

- **timeofday()**

Description: Specifies the current date and time (like **clock_timestamp**, but returned as a **text** string)

Return type: text

Example:

```
openGauss=# SELECT timeofday();
           timeofday
-----
Fri Sep 01 17:05:01.167506 2017 CST
(1 row)
```

- **transaction_timestamp()**

Description: Specifies the current date and time (equivalent to **current_timestamp**)

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT transaction_timestamp();
           transaction_timestamp
-----
2017-09-01 17:05:13.534454+08
(1 row)
```

- **add_months(d,n)**

Description: Returns the time point *d* plus *n* months.

d: indicates the value of the timestamp type and the value that can be implicitly converted to the timestamp type.

n: indicates the value of the INTEGER type and the value that can be implicitly converted to the INTEGER type.

Return type: timestamp

Example:

```
openGauss=# SELECT add_months(to_date('2017-5-29', 'yyyy-mm-dd'), 11) FROM sys_dummy;
           add_months
-----
2018-04-29 00:00:00
(1 row)
```

- **last_day(d)**

Description: Returns the date of the last day of the month that contains time point *d*.

Return type: timestamp

Example:

```
openGauss=# SELECT last_day(to_date('2017-01-01', 'YYYY-MM-DD')) AS cal_result;
           cal_result
-----
2017-01-31 00:00:00
(1 row)
```

- **next_day(x,y)**

Description: Calculates the time of the next week y started from x.

Return type: timestamp

Example:

```
openGauss=# SELECT next_day(timestamp '2017-05-25 00:00:00','Sunday')AS cal_result;
cal_result
-----
2017-05-28 00:00:00
(1 row)
```

- **interval(abstime, abstime)**

Description: Creates a time interval with two pieces of absolute time.

Return type: interval

Example:

```
openGauss=# CALL interval(abstime 'May 10, 1947 23:59:12', abstime 'Mon May 1 00:30:30 1995');
interval
-----
["1947-05-10 23:59:12+09" "1995-05-01 00:30:30+08"]
(1 row)
```

- **intervalend(interval)**

Description: Returns the end time of interval.

Return type: abstime

Example:

```
openGauss=# SELECT intervalend(['"Sep 4, 1983 23:59:12" "Oct4, 1983 23:59:12"']);
intervalend
-----
1983-10-04 23:59:12+08
(1 row)
```

- **intervalrel(interval)**

Description: Calculates and returns the relative time of **interval**.

Return type: reltime

Example:

```
openGauss=# SELECT intervalrel(['"Sep 4, 1983 23:59:12" "Oct4, 1983 23:59:12"']);
intervalrel
-----
1 mon
(1 row)
```

- **smalldatetime_ge**

Description: Checks whether the value of the first parameter is greater than or equal to that of the second parameter.

Parameter: smalldatetime, smalldatetime

Return type: Boolean

- **smalldatetime_cmp**

Description: Compares two smalldatetime values to check whether they are the same.

Parameter: smalldatetime, smalldatetime

Return type: integer

- **smalldatetime_eq**

Description: Compares two smalldatetime values to check whether they are the same.

Parameter: smalldatetime, smalldatetime

- Return type: Boolean
- `smalldatetime_gt`
Description: Determines whether the first parameter is greater than the second.
Parameter: `smalldatetime`, `smalldatetime`
Return type: Boolean
 - `smalldatetime_hash`
Description: Calculates the hash value corresponding to a timestamp.
Parameter: `smalldatetime`
Return type: integer
 - `smalldatetime_in`
Description: Inputs a timestamp.
Parameter: `cstring`, `oid`, `integer`
Return type: `smalldatetime`
 - `smalldatetime_larger`
Description: Returns a larger timestamp.
Parameter: `smalldatetime`, `smalldatetime`
Return type: `smalldatetime`
 - `smalldatetime_le`
Description: Checks whether the value of the first parameter is less than or equal to that of the second parameter.
Parameter: `smalldatetime`, `smalldatetime`
Return type: Boolean
 - `smalldatetime_lt`
Description: Determines whether the first parameter is less than the second parameter.
Parameter: `smalldatetime`, `smalldatetime`
Return type: Boolean
 - `smalldatetime_ne`
Description: Compares two timestamps to check whether they are different.
Parameter: `smalldatetime`, `smalldatetime`
Return type: Boolean
 - `smalldatetime_out`
Description: Converts a timestamp into the external form.
Parameter: `smalldatetime`
Return type: `cstring`
 - `smalldatetime_send`
Description: Converts a timestamp to the binary format.
Parameter: `smalldatetime`
Return type: `bytea`

- `smalldatetime_smaller`
Description: Returns a smaller smalldatetime.
Parameter: smalldatetime, smalldatetime
Return type: smalldatetime
- `smalldatetime_to_abstime`
Description: Converts smalldatetime to abstime.
Parameter: smalldatetime
Return type: abstime
- `smalldatetime_to_time`
Description: Converts smalldatetime to time.
Parameter: smalldatetime
Return type: time without time zone
- `smalldatetime_to_timestamp`
Description: Converts smalldatetime to timestamp.
Parameter: smalldatetime
Return type: timestamp without time zone
- `smalldatetime_to_timestamptz`
Description: Converts smalldatetime to timestamptz.
Parameter: smalldatetime
Return type: timestamp with time zone
- `smalldatetime_to_varchar2`
Description: Converts smalldatetime to varchar2.
Parameter: smalldatetime
Return type: character varying

 NOTE

There are multiple methods for obtaining the current time. Select an appropriate API based on the actual service scenario.

1. The following APIs return values based on the start time of the current transaction:

```
CURRENT_DATE  
CURRENT_TIME  
CURRENT_TIME(precision)  
CURRENT_TIMESTAMP(precision)  
LOCALTIME  
LOCALTIMESTAMP  
LOCALTIME(precision)  
LOCALTIMESTAMP(precision)  
transaction_timestamp()  
now()
```

The values transferred by **CURRENT_TIME** and **CURRENT_TIMESTAMP(precision)** contain time zone information. The values transferred by **LOCALTIME** and **LOCALTIMESTAMP** do not contain time zone information. **CURRENT_TIME**, **LOCALTIME**, and **LOCALTIMESTAMP** can specify a precision parameter, which rounds the seconds field of the result to the decimal place. If there is no precision parameter, the result is given the full precision that can be obtained.

Because these functions all return results by the start time of the current transaction, their values do not change throughout the transaction. This can be considered as a feature with the purpose to allow a transaction to have a consistent concept at the "current" time, so that multiple modifications in the same transaction can maintain the same timestamp.

transaction_timestamp() is equivalent to **CURRENT_TIMESTAMP(precision)**, indicating the start time of the transaction where the current statement is located. **now()** is equivalent to **transaction_timestamp()**.

2. The following APIs return the start time of the current statement:
`statement_timestamp()`

`statement_timestamp()` returns the start time of the current statement (more accurately, the time when the last instruction is received from the client). The return values of **statement_timestamp()** and **transaction_timestamp()** are the same during the execution of the first instruction of a transaction, but may be different in subsequent instructions.

3. The following APIs return the actual current time when the function is called:
`clock_timestamp()`
`timeofday()`

clock_timestamp() returns the actual current time, and its value changes even in the same SQL statement. Similar to **clock_timestamp()**, **timeofday()** also returns the actual current time. However, the result of **timeofday()** is a formatted text string instead of a timestamp with time zone information.

Table 7-29 shows the templates for truncating date/time values.

Table 7-29 Truncating date/time values

Item	Format	Description
Microsecond	MICROSECON	Truncates date/time values, accurate to the microsecond (000000–999999).
	US	
	USEC	
	USECOND	

Item	Format	Description
Millisecond	MILLISECON	Truncates date/time values, accurate to the millisecond (000–999).
	MS	
	MSEC	
	MSECOND	
Second	S	Truncates date/time values, accurate to the second (00–59).
	SEC	
	SECOND	
Minute	M	Truncates date/time values, accurate to the minute (00–59).
	MI	
	MIN	
	MINUTE	
Hour	H	Truncates date/time values, accurate to the hour (00–23).
	HH	
	HOUR	
	HR	
Day	D	Truncates date/time values, accurate to the day (01-01 to 12–31)
	DAY	
	DD	
	DDD	
	J	
Week	W	Truncates date/time values, accurate to the week (the first day of the current week).
	WEEK	
Month	MM	Truncates date/time values, accurate to the month (the first day of the current month).
	MON	
	MONTH	
Quarter	Q	Truncates date/time values, accurate to the quarter (the first day of the current quarter).
	QTR	
	QUARTER	
Year	Y	Truncates date/time values, accurate to the year (the first day of the current year).
	YEAR	

Item	Format	Description
	YR	
	YYYY	
Decade	DEC	Truncates date/time values, accurate to the decade (the first day of the current decade).
	DECADE	
Century	C	Truncates date/time values, accurate to the century (the first day of the current century).
	CC	
	CENT	
	CENTURY	
Millennium	MIL	Truncates date/time values, accurate to the millennium (the first day of the current millennium).
	MILLENNIA	
	MILLENNIUM	

TIMESTAMPDIFF

- **TIMESTAMPDIFF**(*unit*, *timestamp_expr1*, *timestamp_expr2*)

The **timestampdiff** function returns the result of **timestamp_expr2** – **timestamp_expr1** in the specified unit. **timestamp_expr1** and **timestamp_expr2** must be value expressions of the timestamp, timestampz, or date type. **unit** specifies the unit of the difference between two dates.

NOTE

This function is valid only when GaussDB is compatible with the MY type (that is, dbcompatibility = 'B'). It is equivalent to **timestamp_diff(text, timestamp, timestamp)**.

- **year**

Year.

```
openGauss=# SELECT TIMESTAMPDIFF(YEAR, '2018-01-01', '2020-01-01');
timestamp_diff
-----
2
(1 row)
```

- **quarter**

Quarter.

```
openGauss=# SELECT TIMESTAMPDIFF(QUARTER, '2018-01-01', '2020-01-01');
timestamp_diff
-----
8
(1 row)
```

- **month**

Month.

```
openGauss=# SELECT TIMESTAMPDIFF(MONTH, '2018-01-01', '2020-01-01');
timestamp_diff
```

- ```

 24
(1 row)
```
- week  
Week.  

```
openGauss=# SELECT TIMESTAMPDIFF(WEEK, '2018-01-01', '2020-01-01');
timestamp_diff

 104
(1 row)
```
- day  
Day.  

```
openGauss=# SELECT TIMESTAMPDIFF(DAY, '2018-01-01', '2020-01-01');
timestamp_diff

 730
(1 row)
```
- hour  
Hour.  

```
openGauss=# SELECT TIMESTAMPDIFF(HOUR, '2020-01-01 10:10:10', '2020-01-01 11:11:11');
timestamp_diff

 1
(1 row)
```
- minute  
Minute.  

```
openGauss=# SELECT TIMESTAMPDIFF(MINUTE, '2020-01-01 10:10:10', '2020-01-01 11:11:11');
timestamp_diff

 61
(1 row)
```
- second  
Second.  

```
openGauss=# SELECT TIMESTAMPDIFF(SECOND, '2020-01-01 10:10:10', '2020-01-01 11:11:11');
timestamp_diff

 3661
(1 row)
```
- microseconds  
The seconds column, including fractional parts, is multiplied by 1,000,000.  

```
openGauss=# SELECT TIMESTAMPDIFF(MICROSECOND, '2020-01-01 10:10:10.000000', '2020-01-01
10:10:10.111111');
timestamp_diff

 111111
(1 row)
```
- timestamp\_expr with the time zone  

```
openGauss=# SELECT TIMESTAMPDIFF(HOUR,'2020-05-01 10:10:10-01','2020-05-01 10:10:10-03');
timestamp_diff

 2
(1 row)
```

## EXTRACT

- **EXTRACT(*field* FROM *source*)**

The **extract** function retrieves subcolumns such as year or hour from date/time values. **source** must be a value expression of type **timestamp**, **time**, or **interval**. (Expressions of type **date** are cast to **timestamp** and can therefore be used as well.) **field** is an identifier or string that selects what column to extract from the source value. The **extract** function returns values of type **double precision**. The following are valid **field** names:

- century

The first century starts at 0001-01-01 00:00:00 AD. This definition applies to all Gregorian calendar countries. There is no century number 0. You go from -1 century to 1 century.

Example:

```
openGauss=# SELECT EXTRACT(CENTURY FROM TIMESTAMP '2000-12-16 12:21:13');
date_part

 20
(1 row)
```

- day

- Specifies the day (1–31) of the month for **timestamp**.

```
openGauss=# SELECT EXTRACT(DAY FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 16
(1 row)
```

- Specifies the number of days for **interval**.

```
openGauss=# SELECT EXTRACT(DAY FROM INTERVAL '40 days 1 minute');
date_part

 40
(1 row)
```

- decade

Specifies the number of decades.

```
openGauss=# SELECT EXTRACT(DECADE FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 200
(1 row)
```

- dow

Specifies the day of the week. The value range from 0 (Sunday) to 6 (Saturday).

```
openGauss=# SELECT EXTRACT(DOW FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 5
(1 row)
```

- doy

Specifies the day of the year (1–365 or 1–366).

```
openGauss=# SELECT EXTRACT(DOY FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 47
(1 row)
```

- epoch

- Specifies the number of seconds since 1970-01-01 00:00:00-00 UTC (can be negative) for **timestamp with time zone** values.

Specifies the number of seconds since 1970-01-01 00:00:00-00 local time for **date** and **timestamp** values.

Specifies the total number of seconds in the interval for **interval** values.

```
openGauss=# SELECT EXTRACT(EPOCH FROM TIMESTAMP WITH TIME ZONE '2001-02-16
20:38:40.12-08');
date_part

982384720.12
(1 row)
openGauss=# SELECT EXTRACT(EPOCH FROM INTERVAL '5 days 3 hours');
date_part

442800
(1 row)
```

- The following converts an epoch value back to a timestamp.

```
openGauss=# SELECT TIMESTAMP WITH TIME ZONE 'epoch' + 982384720.12 * INTERVAL '1
second' AS RESULT;
result

2001-02-17 12:38:40.12+08
(1 row)
```

- hour

Specifies the hour column (0–23).

```
openGauss=# SELECT EXTRACT(HOUR FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

20
(1 row)
```

- isodow

Specifies the day of the week (1–7).

**1** indicates Monday and **7** indicates Sunday.

#### NOTE

This is identical to **dow** except for Sunday.

```
openGauss=# SELECT EXTRACT(ISODOW FROM TIMESTAMP '2001-02-18 20:38:40');
date_part

7
(1 row)
```

- isoyear

The ISO 8601 year that the date falls in (not applicable to intervals).

Each ISO year begins with the Monday of the week containing January 4, so in early January or late December the ISO year may be different from the Gregorian year. See the **week** column for more information.

```
openGauss=# SELECT EXTRACT(ISOYEAR FROM DATE '2006-01-01');
date_part

2005
(1 row)
openGauss=# SELECT EXTRACT(ISOYEAR FROM DATE '2006-01-02');
date_part

2006
(1 row)
```

- microseconds

The seconds column, including fractional parts, is multiplied by 1,000,000.

```
openGauss=# SELECT EXTRACT(MICROSECONDS FROM TIME '17:12:28.5');
date_part

28500000
(1 row)
```

- millennium

Years in the 1900s are in the second millennium. The third millennium started from 0:00 of January 1, 2001.

```
openGauss=# SELECT EXTRACT(MILLENNIUM FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

3
(1 row)
```

- milliseconds

Seconds column, including fractional parts, is multiplied by 1000. Note that this includes full seconds.

```
openGauss=# SELECT EXTRACT(MILLISECONDS FROM TIME '17:12:28.5');
date_part

28500
(1 row)
```

- minute

Specifies the minute column (0–59).

```
openGauss=# SELECT EXTRACT(MINUTE FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

38
(1 row)
```

- month

Specifies the month column (1–12) for **timestamp** values.

```
openGauss=# SELECT EXTRACT(MONTH FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

2
(1 row)
```

Specifies the number of months after modulo 12 is performed (0–11) for **interval** values.

```
openGauss=# SELECT EXTRACT(MONTH FROM INTERVAL '2 years 13 months');
date_part

1
(1 row)
```

- quarter

Specifies the quarter of the year (1–4) that the date is in.

```
openGauss=# SELECT EXTRACT(QUARTER FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

1
(1 row)
```

- second

Specifies the second column, including fractional parts (0–59).

```
openGauss=# SELECT EXTRACT(SECOND FROM TIME '17:12:28.5');
date_part

28.5
(1 row)
```

- **timezone**  
Time zone offset from UTC, measured in seconds. Positive values correspond to time zones east of UTC, negative values to zones west of UTC.
- **timezone\_hour**  
Hour component of the time zone offset.
- **timezone\_minute**  
Minute component of the time zone offset.
- **week**  
Number of the week of the year that the day is in. By definition (ISO 8601), the first week of a year contains January 4 of that year. (The ISO-8601 week starts on Monday.) In other words, the first Thursday of a year is in week 1 of that year.

Because of this, it is possible for early January dates to be part of the 52nd or 53rd week of the previous year, and late December dates to be part of the 1st week of the next year. For example, **2005-01-01** is part of the 53rd week of year 2004, **2006-01-01** is part of the 52nd week of year 2005, and **2012-12-31** is part of the 1st week of year 2013. You are advised to use the columns **isoyear** and **week** together to ensure consistency.

```
openGauss=# SELECT EXTRACT(WEEK FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 7
(1 row)
```

- **year**  
Specifies the year column.

```
openGauss=# SELECT EXTRACT(YEAR FROM TIMESTAMP '2001-02-16 20:38:40');
date_part

 2001
(1 row)
```

## date\_part

The **date\_part** function is modeled on the traditional Ingres equivalent to the SQL-standard function **extract**:

- **date\_part('field', source)**

Note that here the **field** parameter needs to be a string value, not a name. The valid field names are the same as those for **extract**. For details, see [EXTRACT](#).

Example:

```
openGauss=# SELECT date_part('day', TIMESTAMP '2001-02-16 20:38:40');
date_part

 16
(1 row)
openGauss=# SELECT date_part('hour', INTERVAL '4 hours 3 minutes');
date_part

 4
(1 row)
```

**Table 7-30** specifies the schema for formatting date and time values.



**Table 7-30** Schema for formatting date and time

| Category              | Format     | Description                                                                                                                                                                                 |
|-----------------------|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hour                  | HH         | Number of hours in one day (01–12)                                                                                                                                                          |
|                       | HH12       | Number of hours in one day (01–12)                                                                                                                                                          |
|                       | HH24       | Number of hours in one day (00–23)                                                                                                                                                          |
| Minute                | MI         | Minute (00–59)                                                                                                                                                                              |
| Second                | SS         | Second (00–59)                                                                                                                                                                              |
|                       | FF         | Microsecond (000000–999999)                                                                                                                                                                 |
|                       | FF1        | Microsecond (0–9)                                                                                                                                                                           |
|                       | FF2        | Microsecond (00–99)                                                                                                                                                                         |
|                       | FF3        | Microsecond (000–999)                                                                                                                                                                       |
|                       | FF4        | Microsecond (0000–9999)                                                                                                                                                                     |
|                       | FF5        | Microsecond (00000–99999)                                                                                                                                                                   |
|                       | FF6        | Microsecond (000000–999999)                                                                                                                                                                 |
|                       | SSSSS      | Second after midnight (0–86399)                                                                                                                                                             |
| Morning and afternoon | AM or A.M. | Morning identifier                                                                                                                                                                          |
|                       | PM or P.M. | Afternoon identifier                                                                                                                                                                        |
| Year                  | Y,YYY      | Year with comma (with four digits or more)                                                                                                                                                  |
|                       | SYYYY      | Year with four digits BC                                                                                                                                                                    |
|                       | YYYY       | Year (with four digits or more)                                                                                                                                                             |
|                       | YYY        | Last three digits of a year                                                                                                                                                                 |
|                       | YY         | Last two digits of a year                                                                                                                                                                   |
|                       | Y          | Last one digit of a year                                                                                                                                                                    |
|                       | IYYY       | ISO year (with four digits or more)                                                                                                                                                         |
|                       | IYY        | Last three digits of an ISO year                                                                                                                                                            |
|                       | IY         | Last two digits of an ISO year                                                                                                                                                              |
|                       | I          | Last one digit of an ISO year                                                                                                                                                               |
|                       | RR         | Last two digits of a year (A year of the 20th century can be stored in the 21st century.)                                                                                                   |
|                       | RRRR       | Capable of receiving a year with four digits or two digits. If there are 2 digits, the value is the same as the returned value of RR. If there are 4 digits, the value is the same as YYYY. |

| Category    | Format                                                                            | Description                                                                           |
|-------------|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
|             | <ul style="list-style-type: none"><li>• BC or B.C.</li><li>• AD or A.D.</li></ul> | Era indicator Before Christ (BC) and After Christ (AD)                                |
| Month       | MONTH                                                                             | Full name of a month in uppercase (9 characters are filled in if the value is empty.) |
|             | MON                                                                               | Month in abbreviated format in uppercase (with three characters)                      |
|             | MM                                                                                | Month (01–12)                                                                         |
|             | RM                                                                                | Month in Roman numerals (I–XII; I=JAN) and uppercase                                  |
| Day         | DAY                                                                               | Full name of a date in uppercase (9 characters are filled in if the value is empty.)  |
|             | DY                                                                                | Day in abbreviated format in uppercase (with three characters)                        |
|             | DDD                                                                               | Day in a year (001–366)                                                               |
|             | DD                                                                                | Day in a month (01–31)                                                                |
|             | D                                                                                 | Day in a week (1–7).                                                                  |
| Week        | W                                                                                 | Week in a month (1–5) (The first week starts from the first day of the month.)        |
|             | WW                                                                                | Week in a year (1–53) (The first week starts from the first day of the year.)         |
|             | IW                                                                                | Week in an ISO year (The first Thursday is in the first week.)                        |
| Century     | CC                                                                                | Century (with two digits) (The 21st century starts from 2001-01-01.)                  |
| Julian date | J                                                                                 | Julian date (starting from January 1 of 4712 BC)                                      |
| Quarter     | Q                                                                                 | Quarter                                                                               |

 **NOTE**

In the table, the rules for RR to calculate years are as follows:

- If the range of the input two-digit year is between 00 and 49:  
If the last two digits of the current year are between 00 and 49, the first two digits of the returned year are the same as the first two digits of the current year.  
If the last two digits of the current year are between 50 and 99, the first two digits of the returned year equal to the first two digits of the current year plus 1.
- If the range of the input two-digit year is between 50 and 99:  
If the last two digits of the current year are between 00 and 49, the first two digits of the returned year equal to the first two digits of the current year minus 1.  
If the last two digits of the current year are between 50 and 99, the first two digits of the returned year are the same as the first two digits of the current year.

## 7.5.9 Type Conversion Functions

### Type Conversion Functions

- `cash_words(money)`

Description: Type conversion function, which converts money into text.

Example:

```
openGauss=# SELECT cash_words('1.23');
 cash_words

One dollar and twenty three cents
(1 row)
```

- `cast(x as y)`

Description: Converts x into the type specified by y.

Example:

```
openGauss=# SELECT cast('22-oct-1997' as timestamp);
 timestamp

1997-10-22 00:00:00
(1 row)
```

- `hextoraw(text)`

Description: Converts a string in hexadecimal format into raw type.

Return type: raw

Example:

```
openGauss=# SELECT hextoraw('7D');
 hextoraw

7D
(1 row)
```

- `numtoday(numeric)`

Description: Converts values of the number type into the timestamp of the specified type.

Return type: timestamp

Example:

```
openGauss=# SELECT numtoday(2);
 numtoday

2 days
(1 row)
```

- `rawtohex(string)`

Description: Converts a string in binary format into hexadecimal format.

The result is the ASCII code of the input characters in hexadecimal format.

Return type: `varchar`

Example:

```
openGauss=# SELECT rawtohex('1234567');
rawtohex

31323334353637
(1 row)
```

- `to_bigint(varchar)`

Description: Converts the character type to the bigint type.

Return type: `bigint`

Example:

```
openGauss=# SELECT to_bigint('123364545554455');
to_bigint

123364545554455
(1 row)
```

- `to_char(datetime/interval [, fmt])`

Description: Converts a DATETIME or INTERVAL value of the DATE/TIMESTAMP/TIMESTAMP WITH TIME ZONE/TIMESTAMP WITH LOCAL TIME ZONE type into the TEXT type according to the format specified by **fmt**.

- The optional parameter **fmt** allows for the following types: date, time, week, quarter, and century. Each type has a unique template. The templates can be combined together. Common templates include HH, MI, SS, YYYY, MM, and DD.
- A template may have a modification word. FM is a common modification word and is used to suppress the preceding zero or the following blank spaces.

Return type: `text`

Example:

```
openGauss=# SELECT to_char(current_timestamp,'HH12:MI:SS');
to_char

10:19:26
(1 row)
openGauss=# SELECT to_char(current_timestamp,'FMHH12:FMMI:FMSS');
to_char

10:19:46
(1 row)
```

- `to_char(double precision/real, text)`

Description: Converts the values of the floating point type into the strings in the specified format.

Return type: `text`

Example:

```
openGauss=# SELECT to_char(125.8::real, '999D99');
to_char

125.80
(1 row)
```

- `to_char(numeric/smallint/integer/bigint/double precision/real[, fmt])`

Descriptions: Converts an integer or a value in floating point format into a string in specified format.

- The optional parameter **fmt** allows for the following types: decimal characters, grouping characters, positive/negative sign and currency sign. Each type has a unique template. The templates can be combined together. Common templates include: 9, 0, millesimal sign (,), and decimal point (.).
- A template can have a modification word, similar to FM. However, FM does not suppress 0 which is output according to the **template0**.
- Use the template X or x to convert an integer value into a string in hexadecimal format.

Return type: text

Example:

```
openGauss=# SELECT to_char(1485,'9,999');
to_char

1,485
(1 row)
openGauss=# SELECT to_char(1148.5,'9,999.999');
to_char

1,148.500
(1 row)
openGauss=# SELECT to_char(148.5,'990999.909');
to_char

0148.500
(1 row)
openGauss=# SELECT to_char(123,'XXX');
to_char

7B
(1 row)
```

- `to_char(interval, text)`

Description: Converts the values of the time interval type into the strings in the specified format.

Return type: text

Example:

```
openGauss=# SELECT to_char(interval '15h 2m 12s', 'HH24:MI:SS');
to_char

15:02:12
(1 row)
```

- `to_char(integer, text)`

Description: Converts the values of the integer type into the strings in the specified format.

Return type: text

Example:

```
openGauss=# SELECT to_char(125, '999');
to_char

125
(1 row)
```

- `to_char(numeric, text)`

Description: Converts the values of the numeric type into the strings in the specified format.

Return type: text

Example:

```
openGauss=# SELECT to_char(-125.8, '999D99S');
to_char

125.80-
(1 row)
```

- `to_char(string)`

Description: Converts the CHAR/VARCHAR/VARCHAR2/CLOB type into the TEXT type.

If this function is used to convert data of the CLOB type, and the value to be converted exceeds the value range of the target type, an error is returned.

Return type: text

Example:

```
openGauss=# SELECT to_char('01110');
to_char

01110
(1 row)
```

- `to_char(timestamp, text)`

Description: Converts the values of the timestamp type into the strings in the specified format.

Return type: text

Example:

```
openGauss=# SELECT to_char(current_timestamp, 'HH12:MI:SS');
to_char

10:55:59
(1 row)
```

#### NOTE

The `to_char` function outputs the incorrect format (**fmt**) without change. For example, if the **fmt** is **FF10**, **FF1** is matched for formatted output, and then **0** is output without change.

- `to_clob(char/nchar/varchar/varchar2/nvarchar/nvarchar2/text/raw)`

Description: Converts the raw type or text character set type CHAR, NCHAR, VARCHAR, VARCHAR2, NVARCHAR, NVARCHAR2, or TEXT to the CLOB type.

Return type: clob

Example:

```
openGauss=# SELECT to_clob('ABCDEF'::RAW(10));
to_clob

ABCDEF
(1 row)
openGauss=# SELECT to_clob('hello111'::CHAR(15));
to_clob

hello111
(1 row)
openGauss=# SELECT to_clob('gauss123'::NCHAR(10));
to_clob

gauss123
(1 row)
```

```

gauss123
(1 row)
openGauss=# SELECT to_clob('gauss234'::VARCHAR(10));
to_clob

gauss234
(1 row)
openGauss=# SELECT to_clob('gauss345'::VARCHAR2(10));
to_clob

gauss345
(1 row)
openGauss=# SELECT to_clob('gauss456'::NVARCHAR2(10));
to_clob

gauss456
(1 row)
openGauss=# SELECT to_clob('World222!'::TEXT);
to_clob

World222!
(1 row)
```

- **to\_date(text)**

Description: Converts values of the text type into the timestamp in the specified format. Currently, only the following two formats are supported:

- Format 1: Date without separators, for example, 20150814. The value must contain the complete year, month, and day.
- Format 2: Date with separators, for example, 2014-08-14. The separator can be any non-digit character.

Return type: timestamp without time zone

Example:

```
openGauss=# SELECT to_date('2015-08-14');
to_date

2015-08-14 00:00:00
(1 row)
```

- **to\_date(text, text)**

Description: Converts the values of the string type into the dates in the specified format.

Return type: timestamp without time zone

Example:

```
openGauss=# SELECT to_date('05 Dec 2000', 'DD Mon YYYY');
to_date

2000-12-05 00:00:00
(1 row)
```

- **to\_number ( expr [, fmt])**

Description: Converts **expr** into a value of the NUMBER type according to the specified format.

For details about the type conversion formats, see [Table 7-31](#).

If a hexadecimal string is converted into a decimal number, the hexadecimal string can include a maximum of 16 bytes if it is to be converted into a sign-free number.

During the conversion from a hexadecimal string to a decimal digit, the format string cannot have a character other than x or X. Otherwise, an error is reported.

Return type: number

Example:

```
openGauss=# SELECT to_number('12,454.8-', '99G999D9S');
to_number

-12454.8
(1 row)
```

- `to_number(text, text)`

Description: Converts the values of the string type into the numbers in the specified format.

Return type: numeric

Example:

```
openGauss=# SELECT to_number('12,454.8-', '99G999D9S');
to_number

-12454.8
(1 row)
```

- `to_timestamp(double precision)`

Description: Converts a Unix century into a timestamp.

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT to_timestamp(1284352323);
to_timestamp

2010-09-13 12:32:03+08
(1 row)
```

- `to_timestamp(string [,fmt])`

Description: Converts a string into a value of the timestamp type according to the format specified by **fmt**. When **fmt** is not specified, perform the conversion according to the format specified by **nls\_timestamp\_format**.

In **to\_timestamp** in GaussDB:

- If the input year *YYYY* is 0, an error will be reported.
- If the input year *YYYY* is less than 0, specify *SYYYYY* in **fmt**. The year with the value of *n* (an absolute value) BC will be output correctly.

Characters in the **fmt** must match the schema for formatting the date and time. Otherwise, an error is reported.

Return type: timestamp without time zone

Example:

```
openGauss=# SHOW nls_timestamp_format;
nls_timestamp_format

DD-Mon-YYYY HH:MI:SS.FF AM
(1 row)

openGauss=# SELECT to_timestamp('12-sep-2014');
to_timestamp

2014-09-12 00:00:00
(1 row)
```



```
openGauss=# SELECT to_timestamp('12-Sep-10 14:10:10.123000','DD-Mon-YY HH24:MI:SS.FF');
to_timestamp

2010-09-12 14:10:10.123
(1 row)
openGauss=# SELECT to_timestamp('-1','YYYYY');
to_timestamp

0001-01-01 00:00:00 BC
(1 row)
openGauss=# SELECT to_timestamp('98','RR');
to_timestamp

1998-01-01 00:00:00
(1 row)
openGauss=# SELECT to_timestamp('01','RR');
to_timestamp

2001-01-01 00:00:00
(1 row)
```

- `to_timestamp(text, text)`

Description: Converts values of the string type into the timestamp of the specified type.

Return type: timestamp

Example:

```
openGauss=# SELECT to_timestamp('05 Dec 2000', 'DD Mon YYYY');
to_timestamp

2000-12-05 00:00:00
(1 row)
```

**Table 7-31** Template patterns for numeric formatting

| Format     | Description                                                           |
|------------|-----------------------------------------------------------------------|
| 9          | Value with specified digits                                           |
| 0          | Values with leading zeros                                             |
| Period (.) | Decimal point                                                         |
| Comma (,)  | Group (thousand) separator                                            |
| PR         | Negative values in angle brackets                                     |
| S          | Signed number (depending on the locale setting)                       |
| L          | Currency symbol (depending on the locale setting)                     |
| D          | Decimal point (depending on the locale setting)                       |
| G          | Group separator (depending on the locale setting)                     |
| MI         | Minus sign in the specified position (if the number is less than 0)   |
| PL         | Plus sign in the specified position (if the number is greater than 0) |
| SG         | Plus or minus sign in the specified position                          |

| Format   | Description                                     |
|----------|-------------------------------------------------|
| RN       | Roman numeral (ranging from 1 to 3999)          |
| TH or th | Ordinal number suffix                           |
| V        | Decimal with specified number of digits shifted |

- `abstime_text`  
Description: Converts abstime to text.  
Parameter: `abstime`  
Return type: `text`
- `abstime_to_smalldatetime`  
Description: Converts abstime to smalldatetime.  
Parameter: `abstime`  
Return type: `smalldatetime`
- `bigint_tid`  
Description: Converts bigint to tid.  
Parameter: `bigint`  
Return type: `tid`
- `bool_int1`  
Description: Converts bool to int1.  
Parameter: `Boolean`  
Return type: `tinyint`
- `bool_int2`  
Description: Converts bool to int2.  
Parameter: `Boolean`  
Return type: `smallint`
- `bool_int8`  
Description: Converts bool to int8.  
Parameter: `Boolean`  
Return type: `bigint`
- `bpchar_date`  
Description: Converts a string to a date.  
Parameter: `character`  
Return type: `date`
- `bpchar_float4`  
Description: Converts a string to float4.  
Parameter: `character`  
Return type: `real`
- `bpchar_float8`

Description: Converts a string to float8.

Parameter: character

Return type: double precision

- `bpchar_int4`

Description: Converts a string to int4.

Parameter: character

Return type: integer

- `bpchar_int8`

Description: Converts a string to int8.

Parameter: character

Return type: bigint

- `bpchar_numeric`

Description: Converts a string to numeric.

Parameter: character

Return type: numeric

- `bpchar_timestamp`

Description: Converts a string to a timestamp.

Parameter: character

Return type: timestamp without time zone

- `bpchar_to_smalldatetime`

Description: Converts a string to smalldatetime.

Parameter: character

Return type: smalldatetime

- `complex_array_in`

Description: Converts the external `complex_array` type to the internal `anyarray` array type.

Parameter: `cstring`, `oid`, `int2vector`

Return type: `anyarray`

- `date_bpchar`

Description: Converts the date type to `bpchar`.

Parameter: `date`

Return type: `character`

- `date_text`

Description: Converts date to text.

Parameter: `date`

Return type: `text`

- `date_varchar`

Description: Converts date to `varchar`.

Parameter: `date`

Return type: `character varying`

- **f4toi1**  
Description: Forcibly converts float4 to uint8.  
Parameter: real  
Return type: tinyint
- **f8toi1**  
Description: Forcibly converts float8 to uint8.  
Parameter: double precision  
Return type: tinyint
- **float4\_bpchar**  
Description: Converts float4 to bpchar.  
Parameter: real  
Return type: character
- **float4\_text**  
Description: Converts float4 to text.  
Parameter: real  
Return type: text
- **float4\_varchar**  
Description: Converts float4 to varchar.  
Parameter: real  
Return type: character varying
- **float8\_bpchar**  
Description: Converts float8 to bpchar.  
Parameter: double precision  
Return type: character
- **float8\_interval**  
Description: Converts float8 to interval.  
Parameter: double precision  
Return type: interval
- **float8\_text**  
Description: Converts float8 to text.  
Parameter: double precision  
Return type: text
- **float8\_varchar**  
Description: Converts float8 to varchar.  
Parameter: double precision  
Return type: character varying
- **i1tof4**  
Description: Converts uint8 to float4.  
Parameter: tinyint  
Return type: real

- **i1tof8**  
Description: Converts uint8 to float8.  
Parameter: tinyint  
Return type: double precision
- **i1toi2**  
Description: Converts uint8 to int16.  
Parameter: tinyint  
Return type: smallint
- **i1toi4**  
Description: Converts uint8 to int32.  
Parameter: tinyint  
Return type: integer
- **i1toi8**  
Description: Converts uint8 to int64.  
Parameter: tinyint  
Return type: bigint
- **i2toi1**  
Description: Converts int16 to uint8.  
Parameter: smallint  
Return type: tinyint
- **i4toi1**  
Description: Converts int32 to uint8.  
Parameter: integer  
Return type: tinyint
- **i8toi1**  
Description: Converts int64 to uint8.  
Parameter: bigint  
Return type: tinyint
- **int1\_avg\_accum**  
Description: Adds the second parameter of the uint8 type to the first parameter. The first parameter is an array of the bigint type.  
Parameter: bigint[], tinyint  
Return type: bigint[]
- **int1\_bool**  
Description: Converts uint8 to bool.  
Parameter: tinyint  
Return type: Boolean
- **int1\_bpchar**  
Description: Converts uint8 to bpchar.  
Parameter: tinyint  
Return type: character

- `int1_mul_cash`  
Description: Returns the product of a parameter of the `int8` type and a parameter of the `cash` type. The return type is `cash`.  
Parameter: `tinyint`, `money`  
Return type: `money`
- `int1_numeric`  
Description: Converts `uint8` to `numeric`.  
Parameter: `tinyint`  
Return type: `numeric`
- `int1_nvarchar2`  
Description: Converts `uint8` to `nvarchar2`.  
Parameter: `tinyint`  
Return type: `nvarchar2`
- `int1_text`  
Description: Converts `uint8` to `text`.  
Parameter: `tinyint`  
Return type: `text`
- `int1_varchar`  
Description: Converts `uint8` to `varchar`.  
Parameter: `tinyint`  
Return type: `character varying`
- `int1in`  
Description: Converts a string into an unsigned 1-byte integer.  
Parameter: `cstring`  
Return type: `tinyint`
- `int1out`  
Description: Converts an unsigned 1-byte integer into a string.  
Parameter: `tinyint`  
Return type: `cstring`
- `int1up`  
Description: Converts an input integer to an unsigned 1-byte integer.  
Parameter: `tinyint`  
Return type: `tinyint`
- `int2_bool`  
Description: Converts a signed two-byte integer to the `bool` type.  
Parameter: `smallint`  
Return type: `Boolean`
- `int2_bpchar`  
Description: Converts a signed two-byte integer to the `bpchar` type.  
Parameter: `smallint`  
Return type: `character`

- `int2_text`  
Description: Converts a signed two-byte integer to the text type.  
Parameter: `smallint`  
Return type: `text`
- `int2_varchar`  
Description: Converts a signed two-byte integer to the varchar type.  
Parameter: `smallint`  
Return type: `character varying`
- `int4_bpchar`  
Description: Converts a signed four-byte integer to `bpchar`.  
Parameter: `integer`  
Return type: `character`
- `int4_text`  
Description: Converts a signed four-byte integer to the text type.  
Parameter: `integer`  
Return type: `text`
- `int4_varchar`  
Description: Converts a signed four-byte integer into `varchar`.  
Parameter: `integer`  
Return type: `character varying`
- `int8_bool`  
Description: Converts an eight-byte signed integer to a Boolean value.  
Parameter: `bigint`  
Return type: `Boolean`
- `int8_bpchar`  
Description: Converts an 8-byte signed integer to `bpchar`.  
Parameter: `bigint`  
Return type: `character`
- `int8_text`  
Description: Converts an eight-byte signed integer to the text type.  
Parameter: `bigint`  
Return type: `text`
- `int8_varchar`  
Description: Converts an eight-byte signed integer to `varchar`.  
Parameter: `bigint`  
Return type: `character varying`
- `intervaltonum`  
Description: Converts the internal `date` type to numeric.  
Parameter: `interval`  
Return type: `numeric`

- `numeric_bpchar`  
Description: Converts numeric to bpchar.  
Parameter: numeric  
Return type: character
- `numeric_int1`  
Description: Converts numeric to a signed one-byte integer.  
Parameter: numeric  
Return type: tinyint
- `numeric_text`  
Description: Converts numeric to text.  
Parameter: numeric  
Return type: text
- `numeric_varchar`  
Description: Converts numeric to varchar.  
Parameter: numeric  
Return type: character varying
- `nvarchar2in`  
Description: Converts c string to varchar.  
Parameter: cstring, oid, integer  
Return type: nvarchar2
- `nvarchar2out`  
Description: Converts text into a c string.  
Parameter: nvarchar2  
Return type: cstring
- `nvarchar2send`  
Description: Converts varchar to binary.  
Parameter: nvarchar2  
Return type: bytea
- `oidvectorin_extend`  
Description: Converts a string to oidvector.  
Parameter: cstring  
Return type: oidvector\_extend
- `oidvectorout_extend`  
Description: Converts oidvector to a string.  
Parameter: oidvector\_extend  
Return type: cstring
- `oidvectorsend_extend`  
Description: Converts oidvector to a string.  
Parameter: oidvector\_extend  
Return type: bytea



- reltime\_text  
Description: Converts reltime to text.  
Parameter: reltime  
Return type: text
- text\_date  
Description: Converts the text type to the date type.  
Parameter: text  
Return type: date
- text\_float4  
Description: Converts text to float4.  
Parameter: text  
Return type: real
- text\_float8  
Description: Converts the text type to float8.  
Parameter: text  
Return type: double precision
- text\_int1  
Description: Converts the text type to int1.  
Parameter: text  
Return type: tinyint
- text\_int2  
Description: Converts the text type to the int2 type.  
Parameter: text  
Return type: smallint
- text\_int4  
Description: Converts the text type to int4.  
Parameter: text  
Return type: integer
- text\_int8  
Description: Converts the text type to the int8 type.  
Parameter: text  
Return type: bigint
- text\_numeric  
Description: Converts the text type to the numeric type.  
Parameter: text  
Return type: numeric
- text\_timestamp  
Description: Converts the text type to the timestamp type.  
Parameter: text  
Return type: timestamp without time zone

- `time_text`  
Description: Converts the time type to the text type.  
Parameter: time without time zone  
Return type: text
- `timestamp_text`  
Description: Converts the timestamp type to the text type.  
Parameter: timestamp without time zone  
Return type: text
- `timestamp_to_smalldatetime`  
Description: Converts the timestamp type to the smalldatetime type.  
Parameter: timestamp without time zone  
Return type: smalldatetime
- `timestamp_varchar`  
Description: Converts the timestamp type to varchar.  
Parameter: timestamp without time zone  
Return type: character varying
- `timestamptz_to_smalldatetime`  
Description: Converts timestamptz to smalldatetime.  
Parameter: timestamp with time zone  
Return type: smalldatetime
- `timestampzone_text`  
Description: Converts the timestampzone type to the text type.  
Parameter: timestamp with time zone  
Return type: text
- `timetz_text`  
Description: Converts the timetz type to the text type.  
Parameter: time with time zone  
Return type: text
- `to_integer`  
Description: Converts data to the integer type.  
Parameter: character varying  
Return type: integer
- `to_interval`  
Description: Converts to the interval type.  
Parameter: character varying  
Return type: interval
- `to_numeric`  
Description: Converts to the numeric type.  
Parameter: character varying  
Return type: numeric

- `to_nvarchar2`  
Description: Converts to the `nvarchar2` type.  
Parameter: numeric  
Return type: `nvarchar2`
- `to_text`  
Description: Converts to the `text` type.  
Parameter: `smallint`  
Return type: `text`
- `to_ts`  
Description: Converts to the `ts` type.  
Parameter: character varying  
Return type: timestamp without time zone
- `to_varchar2`  
Description: Converts to the `varchar2` type.  
Parameter: timestamp without time zone  
Return type: character varying
- `varchar_date`  
Description: Converts `varchar` to `date`.  
Parameter: character varying  
Return type: `date`
- `varchar_float4`  
Description: Converts `varchar` to `float4`.  
Parameter: character varying  
Return type: `real`
- `varchar_float8`  
Description: Converts the `varchar` type to the `float8` type.  
Parameter: character varying  
Return type: double precision
- `varchar_int4`  
Description: Converts the type from `varchar` to `int4`.  
Parameter: character varying  
Return type: integer
- `varchar_int8`  
Description: Converts the `varchar` type to the `int8` type.  
Parameter: character varying  
Return type: `bigint`
- `varchar_numeric`  
Description: Converts `varchar` to `numeric`.  
Parameter: character varying  
Return type: `numeric`

- `varchar_timestamp`  
Description: Converts varchar to timestamp.  
Parameter: character varying  
Return type: timestamp without time zone
- `varchar2_to_smlldatetime`  
Description: Converts varchar2 to smlldatetime.  
Parameter: character varying  
Return type: smalldatetime
- `xidout4`  
Description: The xid output is a four-byte number.  
Parameter: xid32  
Return type: cstring
- `xidsend4`  
Description: Converts xid to the binary format.  
Parameter: xid32  
Return type: bytea

## Encoding Type Conversion

- `convert_to_nocase(text, text)`  
Description: Converts a string into a specified encoding type.  
Return type: bytea  
Example:

```
openGauss=# SELECT convert_to_nocase('12345', 'GBK');
convert_to_nocase

\x3132333435
(1 row)
```

## 7.5.10 Geometric Functions and Operators

### Geometric Operators

- `+`  
Description: Translation  
Example:  

```
openGauss=# SELECT box '((0,0),(1,1))' + point '(2,0,0)' AS RESULT;
result

(3,1),(2,0)
(1 row)
```
- `-`  
Description: Translation  
Example:  

```
openGauss=# SELECT box '((0,0),(1,1))' - point '(2,0,0)' AS RESULT;
result

(-1,1),(-2,0)
(1 row)
```

- **\***  
Description: Scaling out/Rotation  
Example:  

```
openGauss=# SELECT box '((0,0),(1,1))' * point '(2.0,0)' AS RESULT;
result

(2,2),(0,0)
(1 row)
```
- **/**  
Description: Scaling in/Rotation  
Example:  

```
openGauss=# SELECT box '((0,0),(2,2))' / point '(2.0,0)' AS RESULT;
result

(1,1),(0,0)
(1 row)
```
- **#**  
Description: Intersection of two figures  
Example:  

```
openGauss=# SELECT box '((1,-1),(-1,1))' # box '((1,1),(-2,-2))' AS RESULT;
result

(1,1),(-1,-1)
(1 row)
```
- **#**  
Description: Number of paths or polygon vertexes of a figure  
Example:  

```
openGauss=# SELECT # path'((1,0),(0,1),(-1,0))' AS RESULT;
result

3
(1 row)
```
- **@-@**  
Description: Length or circumference of a figure  
Example:  

```
openGauss=# SELECT @-@ path'((0,0),(1,0))' AS RESULT;
result

2
(1 row)
```
- **@@**  
Description: Center of a figure  
Example:  

```
openGauss=# SELECT @@ circle'((0,0),10)' AS RESULT;
result

(0,0)
(1 row)
```
- **<->**  
Description: Distance between two figures  
Example:  

```
openGauss=# SELECT circle'((0,0),1)' <-> circle'((5,0),1)' AS RESULT;
result
```

```

3
(1 row)
```

- **&&**

Description: Overlaps? (One point in common makes this true.)

Example:

```
openGauss=# SELECT box '((0,0),(1,1))' && box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- **<<**

Description: Is strictly left of (no common horizontal coordinate)?

Example:

```
openGauss=# SELECT circle '((0,0),1)' << circle '((5,0),1)' AS RESULT;
result

t
(1 row)
```

- **>>**

Description: Is strictly right of (no common horizontal coordinate)?

Example:

```
openGauss=# SELECT circle '((5,0),1)' >> circle '((0,0),1)' AS RESULT;
result

t
(1 row)
```

- **&<**

Description: Does not extend to the right of?

Example:

```
openGauss=# SELECT box '((0,0),(1,1))' &< box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- **&>**

Description: Does not extend to the left of?

Example:

```
openGauss=# SELECT box '((0,0),(3,3))' &> box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- **<<|**

Description: Is strictly below (no common horizontal coordinate)?

Example:

```
openGauss=# SELECT box '((0,0),(3,3))' <<| box '((3,4),(5,5))' AS RESULT;
result

t
(1 row)
```

- **|>>**

Description: Is strictly above (no common horizontal coordinate)?

Example:

```
openGauss=# SELECT box '((3,4),(5,5))' |>> box '((0,0),(3,3))' AS RESULT;
result

t
(1 row)
```

- &<|

Description: Does not extend above?

Example:

```
openGauss=# SELECT box '((0,0),(1,1))' &<| box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- |&>

Description: Does not extend below?

Example:

```
openGauss=# SELECT box '((0,0),(3,3))' |&> box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- <^

Description: Is below (allows touching)?

Example:

```
openGauss=# SELECT box '((0,0),(-3,-3))' <^ box '((0,0),(2,2))' AS RESULT;
result

t
(1 row)
```

- >^

Description: Is above (allows touching)?

Example:

```
openGauss=# SELECT box '((0,0),(2,2))' >^ box '((0,0),(-3,-3))' AS RESULT;
result

t
(1 row)
```

- ?#

Description: Intersect?

Example:

```
openGauss=# SELECT lseg '((-1,0),(1,0))' ?# box '((-2,-2),(2,2))' AS RESULT;
result

t
(1 row)
```

- ?-

Description: Is horizontal?

Example:

```
openGauss=# SELECT ?- lseg '((-1,0),(1,0))' AS RESULT;
result

t
(1 row)
```

- ?-

Description: Are horizontally aligned?

Example:

```
openGauss=# SELECT point '(1,0)' ?- point '(0,0)' AS RESULT;
result

t
(1 row)
```

- ?|

Description: Is vertical?

Example:

```
openGauss=# SELECT ?| lseg '((-1,0),(1,0))' AS RESULT;
result

f
(1 row)
```

- ?|

Description: Are vertically aligned?

Example:

```
openGauss=# SELECT point '(0,1)' ?| point '(0,0)' AS RESULT;
result

t
(1 row)
```

- ?-|

Description: Are perpendicular?

Example:

```
openGauss=# SELECT lseg '((0,0),(0,1))' ?-| lseg '((0,0),(1,0))' AS RESULT;
result

t
(1 row)
```

- ?||

Description: Are parallel?

Example:

```
openGauss=# SELECT lseg '((-1,0),(1,0))' ?|| lseg '((-1,2),(1,2))' AS RESULT;
result

t
(1 row)
```

- @>

Description: Contains?

Example:

```
openGauss=# SELECT circle '((0,0),2)' @> point '(1,1)' AS RESULT;
result

t
(1 row)
```

- <@

Description: Contained in or on?

Example:

```
openGauss=# SELECT point '(1,1)' <@ circle '((0,0),2)' AS RESULT;
result
```



- ```
-----
t
(1 row)
```
- `~=`
Description: Same as?
Example:

```
openGauss=# SELECT polygon '((0,0),(1,1))' ~= polygon '((1,1),(0,0))' AS RESULT;
result
-----
t
(1 row)
```

Geometric Functions

- `area(object)`
Description: Area calculation
Return type: double precision
Example:

```
openGauss=# SELECT area(box '((0,0),(1,1))') AS RESULT;
result
-----
1
(1 row)
```
- `center(object)`
Description: Figure center calculation
Return type: point
Example:

```
openGauss=# SELECT center(box '((0,0),(1,2))') AS RESULT;
result
-----
(0.5,1)
(1 row)
```
- `diameter(circle)`
Description: Circle diameter calculation
Return type: double precision
Example:

```
openGauss=# SELECT diameter(circle '((0,0),2.0)') AS RESULT;
result
-----
4
(1 row)
```
- `height(box)`
Description: Vertical size of box
Return type: double precision
Example:

```
openGauss=# SELECT height(box '((0,0),(1,1))') AS RESULT;
result
-----
1
(1 row)
```
- `isclosed(path)`
Description: A closed path?

Return type: Boolean

Example:

```
openGauss=# SELECTisclosed(path '((0,0),(1,1),(2,0)))' AS RESULT;
result
-----
t
(1 row)
```

- **isopen(path)**

Description: An open path?

Return type: Boolean

Example:

```
openGauss=# SELECT isopen(path '[(0,0),(1,1),(2,0)]') AS RESULT;
result
-----
t
(1 row)
```

- **length(object)**

Description: Length calculation

Return type: double precision

Example:

```
openGauss=# SELECT length(path '((-1,0),(1,0))') AS RESULT;
result
-----
4
(1 row)
```

- **npoints(path)**

Description: Number of points in a path

Return type: int

Example:

```
openGauss=# SELECT npoints(path '[(0,0),(1,1),(2,0)]') AS RESULT;
result
-----
3
(1 row)
```

- **npoints(polygon)**

Description: Number of points in a polygon

Return type: int

Example:

```
openGauss=# SELECT npoints(polygon '((1,1),(0,0))') AS RESULT;
result
-----
2
(1 row)
```

- **pclose(path)**

Description: Converts a path to closed.

Return type: path

Example:

```
openGauss=# SELECT pclose(path '[(0,0),(1,1),(2,0)]') AS RESULT;
result
-----
((0,0),(1,1),(2,0))
(1 row)
```

- **popen(path)**

Description: Converts a path to open.

Return type: path

Example:

```
openGauss=# SELECT popen(path '((0,0),(1,1),(2,0)))' AS RESULT;
result
-----
[(0,0),(1,1),(2,0)]
(1 row)
```

- **radius(circle)**

Description: Circle radius calculation

Return type: double precision

Example:

```
openGauss=# SELECT radius(circle '((0,0),2.0)') AS RESULT;
result
-----
2
(1 row)
```

- **width(box)**

Description: Horizontal size of a box

Return type: double precision

Example:

```
openGauss=# SELECT width(box '((0,0),(1,1))') AS RESULT;
result
-----
1
(1 row)
```

Geometric Type Conversion Functions

- **box(circle)**

Description: Circle to box

Return type: box

Example:

```
openGauss=# SELECT box(circle '((0,0),2.0)') AS RESULT;
result
-----
(1.41421356237309,1.41421356237309),(-1.41421356237309,-1.41421356237309)
(1 row)
```

- **box(point, point)**

Description: Points to box

Return type: box

Example:

```
openGauss=# SELECT box(point '(0,0)', point '(1,1)') AS RESULT;
result
-----
(1,1),(0,0)
(1 row)
```

- **box(polygon)**

Description: Polygon to box

Return type: box

Example:

```
openGauss=# SELECT box(polygon '((0,0),(1,1),(2,0))) AS RESULT;
result
-----
(2,1),(0,0)
(1 row)
```

- circle(box)

Description: Box to circle

Return type: circle

Example:

```
openGauss=# SELECT circle(box '((0,0),(1,1))) AS RESULT;
result
-----
<(0.5,0.5),0.707106781186548>
(1 row)
```

- circle(point, double precision)

Description: Center and radius to circle

Return type: circle

Example:

```
openGauss=# SELECT circle(point '(0,0)', 2.0) AS RESULT;
result
-----
<(0,0),2>
(1 row)
```

- circle(polygon)

Description: Polygon to circle

Return type: circle

Example:

```
openGauss=# SELECT circle(polygon '((0,0),(1,1),(2,0))) AS RESULT;
result
-----
<(1,0.3333333333333333),0.924950591148529>
(1 row)
```

- lseg(box)

Description: Box diagonal to line segment

Return type: lseg

Example:

```
openGauss=# SELECT lseg(box '((-1,0),(1,0))') AS RESULT;
result
-----
[(1,0),(-1,0)]
(1 row)
```

- lseg(point, point)

Description: Points to line segment

Return type: lseg

Example:

```
openGauss=# SELECT lseg(point '(-1,0)', point '(1,0)') AS RESULT;
result
-----
[(-1,0),(1,0)]
(1 row)
```

- **slope(point, point)**
Description: Calculates the slope of a straight line formed by two points.
Return type: double

Example:

```
openGauss=# SELECT slope(point '(1,1)', point '(0,0)') AS RESULT;
result
-----
1
(1 row)
```

- **path(polygon)**
Description: Polygon to path
Return type: path

Example:

```
openGauss=# SELECT path(polygon '((0,0),(1,1),(2,0))') AS RESULT;
result
-----
((0,0),(1,1),(2,0))
(1 row)
```

- **point(double precision, double precision)**
Description: Points
Return type: point

Example:

```
openGauss=# SELECT point(23.4, -44.5) AS RESULT;
result
-----
(23.4,-44.5)
(1 row)
```

- **point(box)**
Description: Center of a box
Return type: point

Example:

```
openGauss=# SELECT point(box '((-1,0),(1,0))') AS RESULT;
result
-----
(0,0)
(1 row)
```

- **point(circle)**
Description: Center of a circle
Return type: point

Example:

```
openGauss=# SELECT point(circle '((0,0),2.0)') AS RESULT;
result
-----
(0,0)
(1 row)
```

- **point(lseg)**
Description: Center of a line segment
Return type: point

Example:

```
openGauss=# SELECT point(lseg '((-1,0),(1,0))') AS RESULT;
result
```

```
-----  
(0,0)  
(1 row)
```

- point(polygon)

Description: Center of a polygon

Return type: point

Example:

```
openGauss=# SELECT point(polygon '((0,0),(1,1),(2,0))') AS RESULT;  
result  
-----  
(1,0.3333333333333333)  
(1 row)
```

- polygon(box)

Description: Box to 4-point polygon

Return type: polygon

Example:

```
openGauss=# SELECT polygon(box '((0,0),(1,1))') AS RESULT;  
result  
-----  
((0,0),(0,1),(1,1),(1,0))  
(1 row)
```

- polygon(circle)

Description: Circle to 12-point polygon

Return type: polygon

Example:

```
openGauss=# SELECT polygon(circle '((0,0),2.0)') AS RESULT;  
result  
-----  
-----  
((-2,0),(-1.73205080756888,1),(-1,1.73205080756888),(-1.22464679914735e-16,2),  
(1,1.73205080756888),(1.73205080756888,1),(2,2.44929359829471e-16),  
(1.73205080756888,-0.9999999999999999),(1,-1.73205080756888),(3.67394039744206e-16,-2),  
(-0.9999999999999999,-1.73205080756888),(-1.73205080756888,-1))  
(1 row)
```

- polygon(npts, circle)

Description: Circle to **npts**-point polygon

Return type: polygon

Example:

```
openGauss=# SELECT polygon(12, circle '((0,0),2.0)') AS RESULT;  
result  
-----  
-----  
((-2,0),(-1.73205080756888,1),(-1,1.73205080756888),(-1.22464679914735e-16,2),  
(1,1.73205080756888),(1.73205080756888,1),(2,2.44929359829471e-16),  
(1.73205080756888,-0.9999999999999999),(1,-1.73205080756888),(3.67394039744206e-16,-2),  
(-0.9999999999999999,-1.73205080756888),(-1.73205080756888,-1))  
(1 row)
```

- polygon(path)

Description: Path to polygon

Return type: polygon

Example:

```
openGauss=# SELECT polygon(path '((0,0),(1,1),(2,0)))' AS RESULT;  
result  
-----  
((0,0),(1,1),(2,0))  
(1 row)
```

7.5.11 Network Address Functions and Operators

cidr and inet Operators

The operators <<, <=, >>, and >= test for subnet inclusion. They consider only the network parts of the two addresses (ignoring any host part) and determine whether one network is identical to or a subnet of the other.

- <

Description: Is less than

Example:

```
openGauss=# SELECT inet '192.168.1.5' < inet '192.168.1.6' AS RESULT;  
result  
-----  
t  
(1 row)
```

- <=

Description: Is less than or equals

Example:

```
openGauss=# SELECT inet '192.168.1.5' <= inet '192.168.1.5' AS RESULT;  
result  
-----  
t  
(1 row)
```

- =

Description: Equals

Example:

```
openGauss=# SELECT inet '192.168.1.5' = inet '192.168.1.5' AS RESULT;  
result  
-----  
t  
(1 row)
```

- >=

Description: Is greater than or equals

Example:

```
openGauss=# SELECT inet '192.168.1.5' >= inet '192.168.1.5' AS RESULT;  
result  
-----  
t  
(1 row)
```

- >

Description: Is greater than

Example:

```
openGauss=# SELECT inet '192.168.1.5' > inet '192.168.1.4' AS RESULT;  
result
```

- ```

t
(1 row)
```
- <>

Description: Does not equal to

Example:

```
openGauss=# SELECT inet '192.168.1.5' <> inet '192.168.1.4' AS RESULT;
result

t
(1 row)
```
- <<

Description: Is contained in

Example:

```
openGauss=# SELECT inet '192.168.1.5' << inet '192.168.1/24' AS RESULT;
result

t
(1 row)
```
- <<=

Description: Is contained in or equals

Example:

```
openGauss=# SELECT inet '192.168.1/24' <<= inet '192.168.1/24' AS RESULT;
result

t
(1 row)
```
- >>

Description: Contains

Example:

```
openGauss=# SELECT inet '192.168.1/24' >> inet '192.168.1.5' AS RESULT;
result

t
(1 row)
```
- >>=

Description: Contains or equals

Example:

```
openGauss=# SELECT inet '192.168.1/24' >>= inet '192.168.1/24' AS RESULT;
result

t
(1 row)
```
- ~

Description: Bitwise NOT

Example:

```
openGauss=# SELECT ~ inet '192.168.1.6' AS RESULT;
result

63.87.254.249
(1 row)
```
- &



Description: Performs an AND operation on each bit of the two network addresses.

Example:

```
openGauss=# SELECT inet '192.168.1.6' & inet '10.0.0.0' AS RESULT;
result

0.0.0.0
(1 row)
```

- |

Description: Performs an OR operation on each bit of the two network addresses.

Example:

```
openGauss=# SELECT inet '192.168.1.6' | inet '10.0.0.0' AS RESULT;
result

202.168.1.6
(1 row)
```

- +

Description: Addition

Example:

```
openGauss=# SELECT inet '192.168.1.6' + 25 AS RESULT;
result

192.168.1.31
(1 row)
```

- -

Description: Subtraction

Example:

```
openGauss=# SELECT inet '192.168.1.43' - 36 AS RESULT;
result

192.168.1.7
(1 row)
```

- -

Description: Subtraction

Example:

```
openGauss=# SELECT inet '192.168.1.43' - inet '192.168.1.19' AS RESULT;
result

24
(1 row)
```

## cidr and inet Functions

The **abbrev**, **host**, and **text** functions are primarily intended to offer alternative display formats.

- abbrev(inet)

Description: Abbreviated display format as text

Return type: text

Example:

```
openGauss=# SELECT abbrev(inet '10.1.0.0/16') AS RESULT;
result
```

- ```
-----
10.1.0.0/16
(1 row)
```
- **abbrev(cidr)**
Description: Abbreviated display format as text
Return type: text
Example:

```
openGauss=# SELECT abbrev(cidr '10.1.0.0/16') AS RESULT;
result
-----
10.1/16
(1 row)
```
 - **broadcast(inet)**
Description: Broadcast address for networks
Return type: inet
Example:

```
openGauss=# SELECT broadcast('192.168.1.5/24') AS RESULT;
result
-----
192.168.1.255/24
(1 row)
```
 - **family(inet)**
Description: Extracts family of addresses, **4** for IPv4.
Return type: int
Example:

```
openGauss=# SELECT family('127.0.0.1') AS RESULT;
result
-----
4
(1 row)
```
 - **host(inet)**
Description: Extracts IP addresses as text.
Return type: text
Example:

```
openGauss=# SELECT host('192.168.1.5/24') AS RESULT;
result
-----
192.168.1.5
(1 row)
```
 - **hostmask(inet)**
Description: Constructs the host mask for a network.
Return type: inet
Example:

```
openGauss=# SELECT hostmask('192.168.23.20/30') AS RESULT;
result
-----
0.0.0.3
(1 row)
```
 - **masklen(inet)**
Description: Extracts subnet mask length.
Return type: int

Example:

```
openGauss=# SELECT masklen('192.168.1.5/24') AS RESULT;
result
-----
      24
(1 row)
```

- **netmask(inet)**

Description: Constructs the subnet mask for a network.

Return type: inet

Example:

```
openGauss=# SELECT netmask('192.168.1.5/24') AS RESULT;
result
-----
255.255.255.0
(1 row)
```

- **network(inet)**

Description: Extracts the network part of an address.

Return type: cidr

Example:

```
openGauss=# SELECT network('192.168.1.5/24') AS RESULT;
result
-----
192.168.1.0/24
(1 row)
```

- **set_masklen(inet, int)**

Description: Sets subnet mask length for the **inet** value.

Return type: inet

Example:

```
openGauss=# SELECT set_masklen('192.168.1.5/24', 16) AS RESULT;
result
-----
192.168.1.5/16
(1 row)
```

- **set_masklen(cidr, int)**

Description: Sets subnet mask length for the **cidr** value.

Return type: cidr

Example:

```
openGauss=# SELECT set_masklen('192.168.1.0/24::cidr, 16) AS RESULT;
result
-----
192.168.0.0/16
(1 row)
```

- **text(inet)**

Description: Extracts IP addresses and subnet mask length as text.

Return type: text

Example:

```
openGauss=# SELECT text(inet '192.168.1.5') AS RESULT;
result
-----
192.168.1.5/32
(1 row)
```

Any **cidr** value can be cast to **inet** implicitly or explicitly; therefore, the functions shown above as operating on **inet** also work on **cidr** values. An **inet** value can be cast to **cidr**. After the conversion, any bits to the right of the subnet mask are silently zeroed to create a valid **cidr** value. In addition, you can cast a text string to **inet** or **cidr** using normal casting syntax. For example, **inet(expression)** or **colname::cidr**.

macaddr Functions

The function **trunc(macaddr)** returns a MAC address with the last 3 bytes set to zero.

- trunc(macaddr)**
Description: Sets last 3 bytes to zero.
Return type: macaddr
Example:

```
openGauss=# SELECT trunc(macaddr '12:34:56:78:90:ab') AS RESULT;
          result
-----
12:34:56:00:00:00
(1 row)
```

The **macaddr** type also supports the standard relational operators (such as > and <=) for lexicographical ordering, and the bitwise arithmetic operators (~, & and |) for NOT, AND and OR.

7.5.12 JSON/JSONB Functions and Operators

For details about the JSON/JSONB data type, see [JSON/JSONB Types](#).

Table 7-32 JSON/JSONB common operators

Operator	Left Operand Type	Right Operand Type	Return Type	Description	Example
->	Array - json(b)	int	json(b)	Obtains the array-json element. If the index does not exist, NULL is returned.	<pre>SELECT '["a":"foo"],["b":"bar"], {"c":"baz"}'::json->2; ?column? ----- {"c":"baz"} (1 row)</pre>
->	object-json(b)	text	json(b)	Obtains the value by a key. If no record is found, NULL is returned.	<pre>SELECT '{"a": {"b":"foo"}}'::json->'a'; ?column? ----- {"b":"foo"} (1 row)</pre>

O p e r a t o r	Left Oper and Type	Rig ht Ope rand Type	Re tu rn Ty pe	Description	Example
- > >	Array - json(b)	int	te xt	Obtains the JSON array element. If the index does not exist, NULL is returned.	<pre>SELECT '[1,2,3]::json->>2; ?column? ----- 3 (1 row)</pre>
- > >	objec t- json(b)	text	te xt	Obtains the value by a key. If no record is found, NULL is returned.	<pre>SELECT '{"a":1,"b":2}::json->>'b'; ?column? ----- 2 (1 row)</pre>
# >	conta iner- json (b)	tex t[]	jso n(b)	Obtains the JSON object in the specified path. If the path does not exist, NULL is returned.	<pre>SELECT '{"a": {"b":{"c": "foo"}}}::json #>'a,b'; ?column? ----- {"c": "foo"} (1 row)</pre>
# > >	conta iner- json (b)	tex t[]	te xt	Obtains the JSON object in the specified path. If the path does not exist, NULL is returned.	<pre>SELECT '{"a":[1,2,3],"b":[4,5,6]}::json #>>'a,2'; ?column? ----- 3 (1 row)</pre>

 CAUTION

For the #> and #>> operators, if no data can be found in the specified path, no error is reported and a **NULL** value is returned.

Table 7-33 Additional JSONB support for operators

Operator	Right Operand Type	Description	Example
@>	jsonb	Specifies whether the top layer of the JSON on the left contains all items of the top layer of the JSON on the right.	'{"a":1, "b":2}':::jsonb @> '{"b":2}':::jsonb
<@	jsonb	Specifies whether all items in the JSON file on the left exist at the top layer of the JSON file on the right.	'{"b":2}':::jsonb <@ '{"a":1, "b":2}':::jsonb
?	text	Specifies whether the string of the key or element exists at the top layer of the JSON value.	'{"a":1, "b":2}':::jsonb ? 'b'
?	text[]	Specifies whether any of these array strings exists as top-layer keys.	'{"a":1, "b":2, "c":3}':::jsonb ? array['b', 'c']
?&	text[]	Specifies whether all these array strings exist as top-layer keys.	'["a", "b"]':::jsonb ? & array['a', 'b']
=	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_eq function.	/
<>	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_ne function.	/

Operator	Right Operand Type	Description	Example
<	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_lt function.	/
>	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_gt function.	/
<=	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_le function.	/
>=	jsonb	Determines the size between two JSONB files, which is the same as the jsonb_ge function.	/

Functions Supported by JSON/JSONB

- array_to_json(anyarray [, pretty_bool])

Description: Returns an array as JSON. It combines a multi-dimensional array into a JSON array. Line feeds will be added between one-dimensional elements if **pretty_bool** is **true**.

Return type: json

For example:

```
openGauss=# SELECT array_to_json('{{1,5},{99,100}}':int[]);
 array_to_json
-----
 [[1,5],[99,100]]
 (1 row)
```

- row_to_json(record [, pretty_bool])

Description: Returns a row as JSON. Line feeds will be added between level-1 elements if **pretty_bool** is **true**.

Return type: json

For example:

```
openGauss=# SELECT row_to_json(row(1,'foo'));
 row_to_json
-----
 [{"1":"foo"}]
```

```
-----  
{ "f1":1,"f2":"foo"} (1 row)
```

- `json_array_element(array-json, integer)`, `jsonb_array_element(array-jsonb, integer)`

Description: Same as the operator ``->``, which returns the element with the specified index in the array.

Return type: json, jsonb

For example:

```
openGauss=# SELECT json_array_element('[1,true,[1,[2,3  
]],null]',2);  
json_array_element  
-----  
[1,[2,3]]  
(1 row)
```

- `json_array_element_text(array-json, integer)`, `jsonb_array_element_text(array-jsonb, integer)`

Description: Same as the operator ``->>``, which returns the element with the specified index in the array.

Return type: text, text

For example:

```
openGauss=# SELECT json_array_element_text('[1,true,[1,[2,3]],null]',2);  
json_array_element_text  
-----  
[1,[2,3]]  
(1 row)
```

- `json_object_field(object-json, text)`, `jsonb_object_field(object-jsonb, text)`

Description: Same as the operator ``->``, which returns the value of a specified key in an object.

Return type: json, json

For example:

```
openGauss=# SELECT json_object_field('{ "a": { "b": "foo" } }', 'a');  
json_object_field  
-----  
{ "b": "foo" }  
(1 row)
```

- `json_object_field_text(object-json, text)`, `jsonb_object_field_text(object-jsonb, text)`

Description: Same as the operator ``->``, which returns the value of a specified key in an object.

Return type: text, text

For example:

```
openGauss=# SELECT json_object_field_text('{ "a": { "b": "foo" } }', 'a');  
json_object_field_text  
-----  
{ "b": "foo" }  
(1 row)
```

- `json_extract_path(json, VARIADIC text[])`, `jsonb_extract_path((jsonb, VARIADIC text[])`

Description: Equivalent to the operator ``#>`` searches for JSON based on the path specified by `$2` and returns the result.

Return type: json, jsonb

For example:


```
openGauss=# SELECT json_extract_path('{\"f2\":{\"f3\":1},\"f4\":{\"f5\":99,\"f6\":\"stringy\"}}', 'f4','f6');
json_extract_path
-----
"stringy"
(1 row)
```

- `json_extract_path_op(json, text[]), jsonb_extract_path_op(jsonb, text[])`

Description: It is equivalent to the operator ``#>``, searches for JSON based on the path specified by `$2`, and returns the result.

Return type: json, jsonb

For example:

```
openGauss=# SELECT json_extract_path_op('{\"f2\":{\"f3\":1},\"f4\":{\"f5\":99,\"f6\":\"stringy\"}}',
ARRAY['f4','f6']);
json_extract_path_op
-----
"stringy"
(1 row)
```

- `json_extract_path_text(json, VARIADIC text[]), jsonb_extract_path_text(jsonb, VARIADIC text[])`

Description: It is equivalent to the operator ``#>>``, searches for JSON based on the path specified by `$2`, and returns the result.

Return type: text, text

For example:

```
openGauss=# SELECT json_extract_path_text('{\"f2\":{\"f3\":1},\"f4\":{\"f5\":99,\"f6\":\"stringy\"}}', 'f4','f6');
json_extract_path_text
-----
stringy
(1 row)
```

- `json_extract_path_text_op(json, text[]), jsonb_extract_path_text_op(jsonb, text[])`

Description: It is equivalent to the operator ``#>>>``, searches for JSON based on the path specified by `$2`, and returns the result.

Return type: text, text

For example:

```
openGauss=# SELECT json_extract_path_text_op('{\"f2\":{\"f3\":1},\"f4\":{\"f5\":99,\"f6\":\"stringy\"}}',
ARRAY['f4','f6']);
json_extract_path_text_op
-----
stringy
(1 row)
```

- `Json_array_elements(array-json), jsonb_array_elements(array-jsonb)`

Description: Splits an array. Each element returns a row.

Return type: json, jsonb

For example:

```
openGauss=# SELECT json_array_elements('[1,true,[1,2,3],null]');
json_array_elements
-----
1
true
[1,[2,3]]
null
(4 rows)
```

- `Json_array_elements_text(array-json), jsonb_array_elements_text(array-jsonb)`

Description: Splits an array. Each element returns a row.

Return type: text, text

For example:

```
openGauss=# SELECT * FROM json_array_elements_text('[1,true,[1,[2,3]],null]');
 value
-----
 1
 true
[1,[2,3]]
(4 rows)
```

- `json_array_length(array-json), jsonb_array_length(array-jsonb)`

Description: Returns the array length.

Return type: integer

For example:

```
openGauss=# SELECT json_array_length('[1,2,3,{"f1":1,"f2":[5,6]},4,null]');
 json_array_length
-----
                6
(1 row)
```

- `json_each(object-json), jsonb_each(object-jsonb)`

Description: Splits each key-value pair of an object into one row and two columns.

Return type: `setof(key text, value json), setof(key text, value jsonb)`

For example:

```
openGauss=# SELECT * FROM json_each('{\"f1\":[1,2,3],\"f2\":{\"f3\":1},\"f4\":null}');
 key | value
-----+-----
 f1  | [1,2,3]
 f2  | {\"f3\":1}
 f4  | null
(3 rows)
```

- `json_each_text(object-json), jsonb_each_text(object-jsonb)`

Description: Splits each key-value pair of an object into one row and two columns.

Return type: `setof(key text, value text), setof(key text, value text)`

For example:

```
openGauss=# SELECT * FROM json_each_text('{\"f1\":[1,2,3],\"f2\":{\"f3\":1},\"f4\":null}');
 key | value
-----+-----
 f1  | [1,2,3]
 f2  | {\"f3\":1}
 f4  | 
(3 rows)
```

- `json_object_keys(object-json), jsonb_object_keys(object-jsonb)`

Description: Returns all keys at the top layer of the object.

Return type: SETOF text

For example:

```
openGauss=# SELECT json_object_keys('{\"f1\":\"abc\", \"f2\":{\"f3\":\"a\", \"f4\":\"b\"}, \"f1\":\"abcd\"}');
 json_object_keys
-----
 f1
 f2
 f1
(3 rows)
```

- JSONB deduplication operations:

```
openGauss=# SELECT jsonb_object_keys('{\"f1\":\"abc\", \"f2\":{\"f3\":\"a\", \"f4\":\"b\"}, \"f1\":\"abcd\"}');
 jsonb_object_keys
-----
```

```
f1  
f2  
(2 rows)
```

- `json_populate_record(anyelement, object-json [, bool])`,
`jsonb_populate_record(anyelement, object-jsonb [, bool])`

Description: *\$1* must be a compound parameter. Each key-value in the **object-json** file is split. The key is used as the column name to match the column name in *\$1* and fill in the *\$1* format.

Return type: anyelement, anyelement

For example:

```
openGauss=# CREATE TYPE jpop as (a text, b int, c bool);  
CREATE TYPE  
openGauss=# SELECT * FROM json_populate_record(null::jpop, '{"a":"blurfl","x":43.2}');  
  a  | b | c  
-----+-----  
blurfl |  |  
(1 row)
```

```
openGauss=# SELECT * FROM json_populate_record((1,1,null)::jpop, '{"a":"blurfl","x":43.2}');  
  a  | b | c  
-----+-----  
blurfl | 1 |  
(1 row)
```

- `json_populate_record_set(anyelement, array-json [, bool])`,
`jsonb_populate_record_set(anyelement, array-jsonb [, bool])`

Description: Performs the preceding operations on each element in the *\$2* array by referring to the **json_populate_record** and **jsonb_populate_record** functions. Therefore, each element in the *\$2* array must be of the **object-json** type.

Return type: setof anyelement, setof anyelement

For example:

```
openGauss=# CREATE TYPE jpop as (a text, b int, c bool);  
CREATE TYPE  
openGauss=# SELECT * FROM json_populate_recordset(null::jpop, '{"a":1,"b":2},{a":3,"b":4}');  
  a  | b | c  
----+----  
  1  | 2 |  
  3  | 4 |  
(2 rows)
```

- `json_typeof(json)`, `jsonb_typeof(jsonb)`

Description: Checks the JSON type.

Return type: text, text

For example:

```
openGauss=# SELECT value, json_typeof(value) FROM (values (json '123.4'), (json '"foo"'), (json  
'true'), (json 'null'), (json '[1, 2, 3]'), (json '{"x":"foo", "y":123}'), (NULL::json)) AS data(value);  
  value  | json_typeof  
-----+-----  
  123.4  | number  
  "foo"   | string  
  true    | boolean  
  null    | null  
  [1, 2, 3] | array  
  {"x":"foo", "y":123} | object  
  |  
(7 rows)
```

- `json_build_array( [VARIADIC "any"] )`

Description: Constructs a JSON array from a variable parameter list.

Return type: array-json

For example:

```
openGauss=# SELECT json_build_array('a',1,'b',1.2,'c',true,'d',null,'e',json '{"x": 3, "y": [1,2,3]}','');
               json_build_array
-----
["a", 1, "b", 1.2, "c", true, "d", null, "e", {"x": 3, "y": [1,2,3]}, ""]
(1 row)
```

- `json_build_object( [VARIADIC "any"] )`

Description: Constructs a JSON object from a variable parameter list. The number of input parameters must be an even number. Every two input parameters form a key-value pair. Note that the value of a key cannot be null.

Return type: object-json

For example:

```
openGauss=# SELECT json_build_object(1,2);
               json_build_object
-----
{"1" : 2}
(1 row)
```

- `json_to_record(object-json, bool)`

Description: Like all functions that return **record**, the caller must explicitly define the structure of the record with an AS clause. The key-value pair of **object-json** is split and reassembled. The key is used as a column name to match and fill in the structure of the specified record.

Return type: record

For example:

```
openGauss=# SELECT * FROM json_to_record('{"a":1,"b":"foo","c":"bar"}',true) AS x(a int, b text, d text);
 a | b | d
---+---+---
 1 | foo | 
(1 row)
```

- `json_to_recordset(array-json, bool)`

Description: Executes the preceding function on each element in the array by referring to the `json_to_record` function. Therefore, each element in the array must be object-json.

Return type: SETOF record

For example:

```
openGauss=# SELECT * FROM json_to_recordset(
  openGauss( # '{"a":1,"b":"foo","d":false},{a":2,"b":"bar","c":true}]' ,
  openGauss( # false
  openGauss( # ) as x(a int, b text, c boolean);
 a | b | c
---+---+---
 1 | foo | 
 2 | bar | t
(2 rows)
```

- `json_object(text[]), json_object(text[], text[])`

Description: Constructs an **object-json** from a text array. This is an overloaded function. When the input parameter is a text array, the array length must be an even number, and members are considered as alternate key-value pairs. When two text arrays are used, the first array is considered as a key, and the second array a value. The lengths of the two arrays must be the same. Note that the value of a key cannot be null.

Return type: object-json

For example:

```
openGauss=# SELECT json_object('{a,1,b,2,3,NULL,"d e f","a b c"}');
               json_object
-----
{"a": "1", "b": "2", "3": null, "d e f": "a b c"}
(1 row)
openGauss=# SELECT json_object('{a,b,"a b c"}', '{a,1,1}');
               json_object
-----
{"a": "a", "b": "1", "a b c": "1"}
(1 row)
```

- **json_agg(any)**

Description: Aggregates values into a JSON array.

Return type: array-json

For example:

```
openGauss=# SELECT * FROM classes;
 name | score
-----+-----
 A    |    2
 A    |    3
 D    |    5
 D    |
(4 rows)
openGauss=# SELECT name, json_agg(score) score FROM classes GROUP BY name ORDER BY name;
 name |      score
-----+-----
 A    | [2, 3]
 D    | [5, null]
      | [null]
(3 rows)
```

- **json_object_agg(any, any)**

Description: Aggregates values into a JSON object.

Return type: object-json

For example:

```
openGauss=# SELECT * FROM classes;
 name | score
-----+-----
 A    |    2
 A    |    3
 D    |    5
 D    |
(4 rows)

openGauss=# SELECT json_object_agg(name, score) FROM classes GROUP BY name ORDER BY
name;
               json_object_agg
-----
{"A": 2, "A": 3 }
{"D": 5, "D": null }
(2 rows)
```

- **jsonb_contained(jsonb, jsonb)**

Description: Same as the operator `<@`, determines whether all elements in \$1 exist at the top layer of \$2.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_contained('[1,2,3]', '[1,2,3,4]');
 jsonb_contained
-----
t
(1 row)
```

- **jsonb_contains(jsonb, jsonb)**

Description: Same as the operator `@>`, checks whether all top-layer elements in \$1 are contained in \$2.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_contains('[1,2,3,4]', '[1,2,3]');
 jsonb_contains
-----
t
(1 row)
```

- `jsonb_exists(jsonb, text)`

Description: Same as the operator `?`, determines whether all elements in the string array \$2 exist at the top layer of \$1 in the form of **key\elem\scalar**.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_exists('["1",2,3]', '1');
 jsonb_exists
-----
t
(1 row)
```

- `jsonb_exists_all(jsonb, text[])`

Description: Same as the operator `?&`, checks whether all elements in the string array \$2 exist at the top layer of \$1 in the form of **key\elem\scalar**.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_exists_all('["1","2",3]', '{1, 2}');
 jsonb_exists_all
-----
t
(1 row)
```

- `jsonb_exists_any(jsonb, text[])`

Description: Same as the operator `?|`, checks whether all elements in the string array \$2 exist at the top layer of \$1 in the form of **key\elem\scalar**.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_exists_any('["1","2",3]', '{1, 2, 4}');
 jsonb_exists_any
-----
t
(1 row)
```

- `jsonb_cmp(jsonb, jsonb)`

Description: Compares values. A positive value indicates greater than, a negative value indicates less than, and 0 indicates equal.

Return type: integer

For example:

```
openGauss=# SELECT jsonb_cmp('["a", "b"]', '{"a":1, "b":2}');
 jsonb_cmp
-----
-1
(1 row)
```

- `jsonb_eq(jsonb, jsonb)`

Description: Same as the operator `=`, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_eq('["a", "b"]', '{"a":1, "b":2}');
 jsonb_eq
-----
f
(1 row)
```

- `jsonb_ne(jsonb, jsonb)`

Description: Same as the operator ``<>``, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_ne('["a", "b"]', '{"a":1, "b":2}');
 jsonb_ne
-----
t
(1 row)
```

- `jsonb_gt(jsonb, jsonb)`

Description: Same as the operator ``>``, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_gt('["a", "b"]', '{"a":1, "b":2}');
 jsonb_gt
-----
f
(1 row)
```

- `jsonb_ge(jsonb, jsonb)`

Description: Same as the operator ``>=``, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_ge('["a", "b"]', '{"a":1, "b":2}');
 jsonb_ge
-----
f
(1 row)
```

- `jsonb_lt(jsonb, jsonb)`

Description: Same as the operator ``<``, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_lt('["a", "b"]', '{"a":1, "b":2}');
 jsonb_lt
-----
t
(1 row)
```

- `jsonb_le(jsonb, jsonb)`

Description: Same as the operator ``<=``, compares two values.

Return type: Boolean

For example:

```
openGauss=# SELECT jsonb_le('["a", "b"]', '{"a":1, "b":2}');
 jsonb_le
-----
t
(1 row)
```

- `to_json(anyelement)`

Description: Converts parameters to ``json``.

Return type: json

For example:

```
openGauss=# SELECT to_json('{1,5}::text[]);
to_json
-----
["1","5"]
(1 row)
```

- jsonb_hash(jsonb)

Description: Performs the hash operation on JSONB.

Return type: integer

For example:

```
openGauss=# SELECT jsonb_hash('[1,2,3]');
jsonb_hash
-----
-559968547
(1 row)
```

- Other functions

Description: Internal functions used by JSON and JSONB aggregate functions.

```
json_agg_transfn
json_agg_finalfn
json_object_agg_transfn
json_object_agg_finalfn
```

7.5.13 HLL Functions and Operators

Hash Functions

- hll_hash_boolean(bool)

Description: Hashes data of the Boolean type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_boolean(FALSE);
hll_hash_boolean
-----
-5451962507482445012
(1 row)
```

- hll_hash_boolean(bool, int32)

Description: Configures a hash seed (that is, change the hash policy) and hashes data of the bool type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_boolean(FALSE, 10);
hll_hash_boolean
-----
-1169037589280886076
(1 row)
```

- hll_hash_smallint(smallint)

Description: Hashes data of the smallint type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_smallint(100::smallint);
hll_hash_smallint
-----
```



```
962727970174027904
(1 row)
```

NOTE

If parameters with the same numeric value are hashed using different data types, the data will differ, because hash functions select different calculation policies for each type.

- `hll_hash_smallint(smallint, int32)`

Description: Configures a hash seed (that is, change the hash policy) and hashes data of the smallint type.

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_smallint(100::smallint, 10);
hll_hash_smallint
-----
-9056177146160443041
(1 row)
```

- `hll_hash_integer(integer)`

Description: Hashes data of the integer type.

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_integer(0);
hll_hash_integer
-----
5156626420896634997
(1 row)
```

- `hll_hash_integer(integer, int32)`

Description: Hashes data of the integer type and configures a hash seed (that is, change the hash policy).

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_integer(0, 10);
hll_hash_integer
-----
-5035020264353794276
(1 row)
```

- `hll_hash_bigint(bigint)`

Description: Hashes data of the bigint type.

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_bigint(100::bigint);
hll_hash_bigint
-----
-2401963681423227794
(1 row)
```

- `hll_hash_bigint(bigint, int32)`

Description: Hashes data of the bigint type and configures a hash seed (that is, change the hash policy).

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_bigint(100::bigint, 10);
hll_hash_bigint
```

```
-----  
-2305749404374433531  
(1 row)
```

- **hll_hash_bytea(bytea)**

Description: Hashes data of the bytea type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_bytea(E'\x');  
hll_hash_bytea  
-----  
0  
(1 row)
```

- **hll_hash_bytea(bytea, int32)**

Description: Hashes data of the bytea type and configures a hash seed (that is, change the hash policy).

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_bytea(E'\x', 10);  
hll_hash_bytea  
-----  
7233188113542599437  
(1 row)
```

- **hll_hash_text(text)**

Description: Hashes data of the text type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_text('AB');  
hll_hash_text  
-----  
-5666002586880275174  
(1 row)
```

- **hll_hash_text(text, int32)**

Description: Hashes data of the text type and configures a hash seed (that is, change the hash policy).

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_text('AB', 10);  
hll_hash_text  
-----  
-2215507121143724132  
(1 row)
```

- **hll_hash_any(anytype)**

Description: Hashes data of any type.

Return type: hll_hashval

Example:

```
openGauss=# SELECT hll_hash_any(1);  
hll_hash_any  
-----  
-1316670585935156930  
(1 row)  
  
openGauss=# SELECT hll_hash_any('08:00:2b:01:02:03':macaddr);  
hll_hash_any
```

```
-----  
-3719950434455589360  
(1 row)
```

- `hll_hash_any(anytype, int32)`

Description: Hashes data of any type and configures a hash seed (that is, change the hash policy).

Return type: `hll_hashval`

Example:

```
openGauss=# SELECT hll_hash_any(1, 10);  
hll_hash_any  
-----  
7048553517657992351  
(1 row)
```

- `hll_hashval_eq(hll_hashval, hll_hashval)`

Description: Compares two pieces of data of the **`hll_hashval`** type to check whether they are the same.

Return type: Boolean

Example:

```
openGauss=# SELECT hll_hashval_eq(hll_hash_integer(1), hll_hash_integer(1));  
hll_hashval_eq  
-----  
t  
(1 row)
```

- `hll_hashval_ne(hll_hashval, hll_hashval)`

Description: Compares two pieces of data of the **`hll_hashval`** type to check whether they are different.

Return type: Boolean

Example:

```
openGauss=# SELECT hll_hashval_ne(hll_hash_integer(1), hll_hash_integer(1));  
hll_hashval_ne  
-----  
f  
(1 row)
```

Log Functions

There are three HLL modes: explicit, sparse, and full. When the data size is small, the explicit mode is used. In this mode, distinct values are calculated without errors. As the number of distinct values increases, the HLL mode is switched to the sparse and full modes in sequence. The two modes have no difference in the calculation result, but vary in the calculation efficiency of HLL functions and the storage space of HLL objects. The following functions can be used to view some HLL parameters:

- `hll_print(hll)`

Description: Prints some debugging parameters of an HLL.

Example:

```
openGauss=# SELECT hll_print(hll_empty());  
hll_print  
-----  
type=1(HLL_EMPTY), log2m=14, log2explicit=10, log2sparse=12, duplicatecheck=0  
(1 row)
```

- `hll_type(hll)`

Description: Checks the type of the current HLL. The return values are described as follows: **0** indicates **HLL_UNINIT**, an HLL object that is not initialized. **1** indicates **HLL_EMPTY**, an empty HLL object. **2** indicates **HLL_EXPLICIT**, an HLL object in explicit mode. **3** indicates **HLL_SPARSE**, an HLL object in sparse mode. **4** indicates **HLL_FULL**, an HLL object in full mode. **5** indicates **HLL_UNDEFINED**, an invalid HLL object.

Example:

```
openGauss=# SELECT hll_type(hll_empty());
 hll_type
-----
      1
(1 row)
```

- **hll_log2m(hll)**

Description: Checks the value of **log2m** in the current HLL data structure. **log2m** is the logarithm of the number of buckets. This value affects the error rate of calculating distinct values by HLL. The error rate = $\pm 1.04/\sqrt{2^{\log 2m}}$. If the value of **log2m** ranges from 10 to 16, HLL sets the number of buckets to $2^{\log 2m}$. When the value of **log2explicit** is explicitly set to **-1**, the built-in default value is used.

Example:

```
openGauss=# SELECT hll_log2m(hll_empty());
 hll_log2m
-----
      14
(1 row)

openGauss=# SELECT hll_log2m(hll_empty(10));
 hll_log2m
-----
      10
(1 row)

openGauss=# SELECT hll_log2m(hll_empty(-1));
 hll_log2m
-----
      14
(1 row)
```

- **hll_log2explicit(hll)**

Description: Queries the **log2explicit** value in the current HLL data structure. Generally, the HLL changes from the explicit mode to the sparse mode and then to the full mode. This process is called the promotion hierarchy policy. You can change the value of **log2explicit** to change the policy. For example, if the value of **log2explicit** is **0**, the explicit mode will be skipped and the sparse mode is directly entered. When the value of **log2explicit** is explicitly set to a value ranging from 1 to 12, HLL will switch to the sparse mode when the length of the data segment exceeds $2^{\log 2explicit}$. When the value of **log2explicit** is explicitly set to **-1**, the built-in default value is used.

Example:

```
openGauss=# SELECT hll_log2explicit(hll_empty());
 hll_log2explicit
-----
      10
(1 row)

openGauss=# SELECT hll_log2explicit(hll_empty(12, 8));
 hll_log2explicit
-----
```

```
      8
(1 row)

openGauss=# SELECT hll_log2explicit(hll_empty(12, -1));
 hll_log2explicit 
-----
      10
(1 row)
```

- **hll_log2sparse(hll)**

Description: Queries the value of **log2sparse** in the current HLL data structure. Generally, the HLL changes from the explicit mode to the sparse mode and then to the full mode. This process is called the promotion hierarchy policy. You can adjust the value of **log2sparse** to change the policy. For example, if the value of **log2sparse** is **0**, the system skips the sparse mode and directly enters the full mode. If the value of **log2sparse** is explicitly set to a value ranging from 1 to 14, HLL will switch to the full mode when the length of the data segment exceeds $2^{\text{log2sparse}}$. When the value of **log2sparse** is explicitly set to **-1**, the built-in default value is used.

Example:

```
openGauss=# SELECT hll_log2sparse(hll_empty());
 hll_log2sparse 
-----
      12
(1 row)

openGauss=# SELECT hll_log2sparse(hll_empty(12, 8, 10));
 hll_log2sparse 
-----
      10
(1 row)

openGauss=# SELECT hll_log2sparse(hll_empty(12, 8, -1));
 hll_log2sparse 
-----
      12
(1 row)
```

- **hll_duplicatecheck(hll)**

Description: Specifies whether duplicate check is enabled. The value **0** indicates that it is disabled and the value **1** indicates that it is enabled. This function is disabled by default. If there are many duplicate values, you can enable this function to improve efficiency. When the value of **duplicatecheck** is explicitly set to **-1**, the built-in default value is used.

Example:

```
openGauss=# SELECT hll_duplicatecheck(hll_empty());
 hll_duplicatecheck 
-----
      0
(1 row)

openGauss=# SELECT hll_duplicatecheck(hll_empty(12, 8, 10, 1));
 hll_duplicatecheck 
-----
      1
(1 row)

openGauss=# SELECT hll_duplicatecheck(hll_empty(12, 8, 10, -1));
 hll_duplicatecheck 
-----
      0
(1 row)
```

Functional Functions

- `hll_empty()`

Description: Creates an empty HLL.

Return type: hll

Example:

```
openGauss=# SELECT hll_empty();
          hll_empty
-----
\x484c4c0000000002b0500000000000000000000000000000000000000000000000
(1 row)
```

- `hll_empty(int32 log2m)`

Description: Creates an empty HLL and sets the **log2m** parameter. The parameter value ranges from 10 to 16. If the input is **-1**, the built-in default value is used.

Return type: HLL

Example:

[illegible]

- `hll_empty(int32 log2m, int32 log2explicit)`

Description: Creates an empty HLL and sets the **log2m** and **log2explicit** parameters in sequence. The value of **log2explicit** ranges from 0 to 12. The value **0** indicates that the explicit mode is skipped. This parameter is used to set the threshold of the explicit mode. When the length of the data segment reaches $2^{\text{log2explicit}}$, the mode is switched to the sparse or full mode. If the input is **-1**, the built-in default value of **log2explicit** is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_empty(10, 4);  
                hll_empty  
-----  
\x484c4c00000000000130400000000000000000000000000000  
(1 row)  
  
openGauss=# SELECT hll_empty(10, -1);  
                hll_empty  
-----  
\x484c4c00000000002b040000000000000000000000000000  
(1 row)
```

- `hll_empty(int32 log2m, int32 log2explicit, int64 log2sparse)`

Description: Creates an empty HLL and sets the **log2m**, **log2explicit** and **log2sparse** parameters in sequence. The value of **log2sparse** ranges from 0 to 14. The value **0** indicates that the sparse mode is skipped. This parameter is used to set the threshold of the sparse mode. When the length of the data segment reaches $2^{\text{log2sparse}}$, the mode is switched to the full mode. If the input is **-1**, the built-in default value of **log2sparse** is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_empty(10, 4, 8);
```

```
\x484c4c000000000012040000000000000000000000000000000000000000000000  
(1 row)
```

```
openGauss=# SELECT hll_empty(10, 4, -1);
          hll_empty
-----
\x484c4c000000000013040000000000000000000000000000000000000000000000
(1 row)
```

- `hll_empty(int32 log2m, int32 log2explicit, int64 log2sparse, int32 duplicatecheck)`

Description: Creates an empty HLL and sets the **log2m**, **log2explicit**, **log2sparse**, and **duplicatecheck** parameters in sequence. The value of **duplicatecheck** is **0** or **1**, indicating whether the duplicate check mode is enabled. By default, this mode is disabled. If the input is **-1**, the built-in default value of **duplicatecheck** is used.

Return type: HLL

Example:

[illegible]

```
openGauss=# SELECT hll_empty(10, 4, 8, -1);
           hll_empty
-----
\x484c4c0000000000120400000000000000000000000000000000000000
(1 row)
```

- `hll add(hll, hll hashval)`

Description: Adds **hll_hashval** to an HLL.

Return type: HLL

Example:

```
openGauss=# SELECT hll_add(hll_empty(), hll_hash_integer(1));
          hll_add
-----
\x484c4c08000002002b09000000000000f03f3e2921ff133fbaed3e2921ff133fbaed00
(1 row)
```

- `hll add rev(hll hashval, hll)`

Description: Adds **hll_hashval** to an HLL. This function works the same as **hll_add**, except that the positions of parameters are switched.

Return type: HLL

Example:

```
openGauss=# SELECT hll_add_rev(hll_hash_integer(1), hll_empty());
               hll_add_rev
-----
\x484c4c08000002002b09000000000000f03f3e2921ff133fbaed3e2921ff133fbaed00
(1 row)
```

- $\text{hll} \text{ eq}(\text{hll}, \text{hll})$

Description: Compares two HLLs to check whether they are the same.

Return type: Boolean

Example:

```
openGauss=# SELECT hll_eq(hll_add(hll_empty(), hll_hash_integer(1)), hll_add(hll_empty(),
hll_hash_integer(2)));
 hll_eq
-----
f
(1 row)
```

- **hll_ne(hll, hll)**

Description: Compares two HLLs to check whether they are different.

Return type: Boolean

Example:

```
openGauss=# SELECT hll_ne(hll_add(hll_empty(), hll_hash_integer(1)), hll_add(hll_empty(),
hll_hash_integer(2)));
 hll_ne
-----
t
(1 row)
```

- **hll_cardinality(hll)**

Description: Calculates the number of distinct values of an HLL.

Return type: int

Example:

```
openGauss=# SELECT hll_cardinality(hll_empty() || hll_hash_integer(1));
 hll_cardinality
-----
1
(1 row)
```

- **hll_union(hll, hll)**

Description: Performs an UNION operation on two HLL data structures to obtain one HLL.

Return type: HLL

Example:

```
openGauss=# SELECT hll_union(hll_add(hll_empty(), hll_hash_integer(1)), hll_add(hll_empty(),
hll_hash_integer(2)));
               hll_union
-----
\x484c4c10002000002b090000000000000004000000000000000b3ccc49320cca1ae3e2921ff133fba
ed00
(1 row)
```

Aggregate Functions

- **hll_add_agg(hll_hashval)**

Description: Groups hashed data into HLL

Return type: HLL

Example:

```
-- Prepare data.
openGauss=# CREATE TABLE t_id(id int);
openGauss=# INSERT INTO t_id VALUES(generate_series(1,500));
openGauss=# CREATE TABLE t_data(a int, c text);
openGauss=# INSERT INTO t_data SELECT mod(id,2), id FROM t_id;

-- Create a table and specify an HLL column.
openGauss=# CREATE TABLE t_a_c_hll(a int, c hll);
```



```
-- Use GROUP BY on column a to group data, and insert the data to the HLL.
openGauss=# INSERT INTO t_a_c_hll SELECT a, hll_add_agg(hll_hash_text(c)) FROM t_data GROUP
BY a;

-- Calculate the number of distinct values for each group in the HLL.
openGauss=# SELECT a, #c as cardinality FROM t_a_c_hll ORDER BY a;
 a | cardinality
---+-----
 0 | 247.862354346299
 1 | 250.908710610377
(2 rows)
```

- **hll_add_agg(hll_hashval, int32 log2m)**

Description: Groups hashed data into HLL and specifies the **log2m** parameter. The value ranges from 10 to 16. If the input is **-1** or **NULL**, the built-in default value is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_cardinality(hll_add_agg(hll_hash_text(c), 12)) FROM t_data;
 hll_cardinality
-----
497.965240179228
(1 row)
```

- **hll_add_agg(hll_hashval, int32 log2m, int32 log2explicit)**

Description: Groups hashed data into HLL and specifies the **log2m** and **log2explicit** parameters in sequence. The value of **log2explicit** ranges from 0 to 12. The value **0** indicates that the explicit mode is skipped. This parameter is used to set the threshold of the explicit mode. When the length of the data segment reaches $2^{\text{log2explicit}}$, the mode is switched to the sparse or full mode. If the input is **-1** or **NULL**, the built-in default value of **log2explicit** is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_cardinality(hll_add_agg(hll_hash_text(c), NULL, 1)) FROM t_data;
 hll_cardinality
-----
498.496062953313
(1 row)
```

- **hll_add_agg(hll_hashval, int32 log2m, int32 log2explicit, int64 log2sparse)**

Description: Groups hashed data into HLL and sets the parameters **log2m**, **log2explicit**, and **log2sparse** in sequence. The value of **log2sparse** ranges from 0 to 14. The value **0** indicates that the sparse mode is skipped. This parameter is used to set the threshold of the sparse mode. When the length of the data segment reaches $2^{\text{log2sparse}}$, the mode is switched to the full mode. If the input is **-1** or **NULL**, the built-in default value of **log2sparse** is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_cardinality(hll_add_agg(hll_hash_text(c), NULL, 6, 10)) FROM t_data;
 hll_cardinality
-----
498.496062953313
(1 row)
```

- **hll_add_agg(hll_hashval, int32 log2m, int32 log2explicit, int64 log2sparse, int32 duplicatecheck)**

Description: Groups hashed data into HLL and specifies the **log2m**, **log2explicit**, **log2sparse**, and **duplicatecheck** parameters. The value of

duplicatecheck can be **0** or **1**, indicating whether to enable this mode. By default, this mode is disabled. If the input is **-1** or **NULL**, the built-in default value of **duplicatecheck** is used.

Return type: HLL

Example:

```
openGauss=# SELECT hll_cardinality(hll_add_agg(hll_hash_text(c), NULL, 6, 10, -1)) FROM t_data;
hll_cardinality
-----
498.496062953313
(1 row)
```

- **hll_union_agg(hll)**

Description: Performs an UNION operation on multiple pieces of data of the HLL type to obtain one HLL.

Return type: HLL

Example:

```
-- Perform the UNION operation on data of the HLL type in each group to obtain one HLL, and
-- calculate the number of distinct values:
openGauss=# SELECT #hll_union_agg(c) as cardinality FROM t_a_c_hll;
cardinality
-----
498.496062953313
(1 row)

-- Drop tables.
openGauss=# DROP TABLE t_id;
openGauss=# DROP TABLE t_data;
openGauss=# DROP TABLE t_a_c_hll;
```

NOTE

To perform the UNION operation on data in multiple HLLs, ensure that the HLLs have the same precision. Otherwise, **UNION** cannot be performed. This constraint also applies to the **hll_union(hll, hll)** function.

Obsolete Functions

Some old HLL functions are deprecated due to version upgrade. You can replace them with similar functions.

- **hll_schema_version(hll)**

Description: Checks the schema version in the current HLL. In earlier versions, the schema version is fixed at **1**, which is used to verify the header of the HLL field. After refactoring, the HLL field is added to the header for verification. The schema version is no longer used.

- **hll_regwidth(hll)**

Description: Queries the bucket size in the HLL data structure. In earlier versions, the value of **regwidth** ranges from 1 to 5, which has a large error and limits the upper limit of the cardinality estimation. After refactoring, the value of **regwidth** is fixed at **6** and the **regwidth** variable is not used.

- **hll_expthresh(hll)**

Description: Obtains the **expthresh** value in the current HLL. The **hll_log2explicit(hll)** function is used to replace similar functions.

- **hll_sparseon(hll)**

Description: Specifies whether to enable the sparse mode. Use **hll_log2sparse(hll)** to replace similar functions. The value **0** indicates that the sparse mode is disabled.

Built-in Functions

HyperLogLog (HLL) has a series of built-in functions for internal data processing. Generally, users do not need to know how to use these functions. For details, see [Table 7-34](#).

Table 7-34 Built-in Functions

Function	Description
hll_in	Receives hll data in string format.
hll_out	Sends hll data in string format.
hll_recv	Receives hll data in bytea format.
hll_send	Sends hll data in bytea format.
hll_trans_in	Receives hll_trans_type data in string format.
hll_trans_out	Sends hll_trans_type data in string format.
hll_trans_recv	Receives hll_trans_type data in bytea format.
hll_trans_send	Sends hll_trans_type data in bytea format.
hll_tpmmod_in	Receives tpmmod data.
hll_tpmmod_out	Sends tpmmod data.
hll_hashval_in	Receives hll_hashval data.
hll_hashval_out	Sends hll_hashval data.
hll_add_trans0	It is similar to hll_add . No input parameter is specified during initialization. It is usually used in the first phase of DNs in aggregation operations.
hll_add_trans1	It is similar to hll_add . An input parameter is specified during initialization. It is usually used in the first phase of DNs in aggregation operations.
hll_add_trans2	It is similar to hll_add . Two input parameters are specified during initialization. It is usually used in the first phase of DNs in aggregation operations.
hll_add_trans3	It is similar to hll_add . Three input parameters are specified during initialization. It is usually used in the first phase of DNs in aggregation operations.
hll_add_trans4	It is similar to hll_add . Four input parameters are specified during initialization. It is usually used in the first phase of DNs in aggregation operations.

Function	Description
hll_union_trans	It is similar to hll_union and is used in the first phase of DNs in aggregation operations.
hll_union_collect	It is similar to hll_union and is used in the second phase of DNs in aggregation operations to summarize the results of each DN.
hll_pack	It is used in the third phase of DNs in aggregation operations to convert a user-defined type hll_trans_type to the hll type.
hll	Converts an HLL type to another HLL type. Input parameters can be specified.
hll_hashval	Converts the bigint type to the hll_hashval type.
hll_hashval_int4	Converts the int4 type to the hll_hashval type.

Operators

- **=**
Description: Compares the values of HLL and **hll_hashval** types to check whether they are the same.

Return type: Boolean

Example:

```
--hll
openGauss=# SELECT (hll_empty() || hll_hash_integer(1)) = (hll_empty() || hll_hash_integer(1));
column
-----
t
(1 row)

--hll_hashval
openGauss=# SELECT hll_hash_integer(1) = hll_hash_integer(1);
?column?
-----
t
(1 row)
```

- **<> or !=**
Description: Compares the values of HLL and **hll_hashval** types to check whether they are different.

Return type: Boolean

Example:

```
--hll
openGauss=# SELECT (hll_empty() || hll_hash_integer(1)) <> (hll_empty() || hll_hash_integer(2));
?column?
-----
t
(1 row)

--hll_hashval
openGauss=# SELECT hll_hash_integer(1) <> hll_hash_integer(2);
?column?
-----
```

- ```
t
(1 row)
```

• `||`

Description: Represents the functions of **hll\_add**, **hll\_union**, and **hll\_add\_rev**.  
Return type: HLL  
Example:

```
--hll_add
openGauss=# SELECT hll_empty() || hll_hash_integer(1);
 ?column?

\x484c4c08000002002b090000000000000f03f3e2921ff133fbaed3e2921ff133fbaed00
(1 row)

--hll_add_rev
openGauss=# SELECT hll_hash_integer(1) || hll_empty();
 ?column?

\x484c4c08000002002b090000000000000f03f3e2921ff133fbaed3e2921ff133fbaed00
(1 row)

--hll_union
openGauss=# SELECT (hll_empty() || hll_hash_integer(1)) || (hll_empty() || hll_hash_integer(2));
 ?column?

\x484c4c10002000002b090000000000000004000000000000000b3ccc49320cca1ae3e2921ff133fbaed00
(1 row)
```

• `#`

Description: Calculates the number of distinct values of an HLL. It works the same as the **hll\_cardinality** function.  
Return type: int  
Example:

```
openGauss=# SELECT #(hll_empty() || hll_hash_integer(1));
 ?column?

 1
(1 row)
```

## 7.5.14 SEQUENCE Functions

The sequence functions provide a simple method to ensure security of multiple users for users to obtain sequence values from sequence objects.

- `nextval(regclass)`

Description: Specifies an increasing sequence and returns a new value.

### NOTE

To avoid blocking of concurrent transactions that obtain numbers from the same sequence, a `nextval` operation is never rolled back; that is, once a value has been fetched it is considered used, even if the transaction that did the `nextval` later aborts. This means that aborted transactions may leave unused "holes" in the sequence of assigned values. Therefore, GaussDB sequences cannot be used to obtain sequence without gaps.

**NOTICE**

The nextval function can be executed only on the primary node. It is not supported on standby nodes.

Return type: numeric

The nextval function can be called in either of the following ways: (In example 2, the sequence name cannot contain a dot.)

```
openGauss=# CREATE SEQUENCE seqDemo;
-- Example 1:
openGauss=# SELECT nextval('seqDemo');
nextval

1
(1 row)
-- Example 2:
openGauss=# SELECT seqDemo.nextval;
nextval

2
(1 row)
openGauss=# DROP SEQUENCE seqDemo;
```

- **currval(regclass)**

Description: Returns the last value returned by nextval in the current session. If nextval has not been called for the specified sequence in the current session, an error is reported when currval is called.

Return type: numeric

The currval function can be called in either of the following ways: (In example 2, the sequence name cannot contain a dot.)

```
openGauss=# CREATE SEQUENCE seq1;
openGauss=# SELECT nextval('seq1');
-- Example 1:
openGauss=# SELECT currval('seq1');
currval

1
(1 row)
-- Example 2:
openGauss=# SELECT seq1.currval;
currval

1
(1 row)
openGauss=# DROP SEQUENCE seq1;
```

- **lastval()**

Description: Returns the last value returned by nextval in the current session. This function is equivalent to currval, except that it does not use sequence names as parameters. It fetches the sequence used by nextval last time in the current session. If nextval has not been called in the current session, an error is reported when lastval is called.

Return type: numeric

Example:

```
openGauss=# CREATE SEQUENCE seq1;
openGauss=# SELECT nextval('seq1');
openGauss=# SELECT lastval();
lastval

```

```
1
(1 row)
openGauss=# DROP SEQUENCE seq1;
```

- **setval(regclass, numeric)**

Description: Sets the current value of a sequence.

Return type: numeric

Example:

```
openGauss=# CREATE SEQUENCE seqDemo;
openGauss=# SELECT nextval('seqDemo');
openGauss=# SELECT setval('seqDemo',5);
setval
```

```
5
(1 row)
openGauss=# DROP SEQUENCE seqDemo;
```

- **setval(regclass, numeric, Boolean)**

Description: Sets the current value of a sequence and the is\_called sign.

Return type: numeric

Example:

```
openGauss=# CREATE SEQUENCE seqDemo;
openGauss=# SELECT nextval('seqDemo');
openGauss=# SELECT setval('seqDemo',5,true);
setval
```

```
5
(1 row)
openGauss=# DROP SEQUENCE seqDemo;
```

#### NOTE

The current session will take effect immediately after **setval** is performed. If other sessions have buffered sequence values, **setval** will take effect only after the values are used up. Therefore, to prevent sequence value conflicts, you are advised to use **setval** with caution.

Because the sequence is non-transactional, the change caused by **setval** will not be undone by transaction rollback.

#### NOTICE

The **nextval** function can be executed only on the primary node. It is not supported on standby nodes.

- **pg\_sequence\_last\_value(sequence\_oid oid, OUT cache\_value int16, OUT last\_value int16)**

Description: Obtains the parameters of a specified sequence, including the cache value and current value.

Return type: int16, int16

## 7.5.15 Array Functions and Operators

### Array Operators

- **=**

Description: Specifies whether two arrays are equal.

Example:

```
openGauss=# SELECT ARRAY[1.1,2.1,3.1]::int[] = ARRAY[1,2,3] AS RESULT ;
result

t
(1 row)
```

- <>

Description: Specifies whether two arrays are not equal.

Example:

```
openGauss=# SELECT ARRAY[1,2,3] <> ARRAY[1,2,4] AS RESULT;
result

t
(1 row)
```

- <

Description: Specifies whether an array is less than another.

Example:

```
openGauss=# SELECT ARRAY[1,2,3] < ARRAY[1,2,4] AS RESULT;
result

t
(1 row)
```

- >

Description: Specifies whether an array is greater than another.

Example:

```
openGauss=# SELECT ARRAY[1,4,3] > ARRAY[1,2,4] AS RESULT;
result

t
(1 row)
```

- <=

Description: Specifies whether an array is less than another.

Example:

```
openGauss=# SELECT ARRAY[1,2,3] <= ARRAY[1,2,3] AS RESULT;
result

t
(1 row)
```

- >=

Description: Specifies whether an array is greater than or equal to another.

Example:

```
openGauss=# SELECT ARRAY[1,4,3] >= ARRAY[1,4,3] AS RESULT;
result

t
(1 row)
```

- @>

Description: Specifies whether an array contains another.

Example:

```
openGauss=# SELECT ARRAY[1,4,3] @> ARRAY[3,1] AS RESULT;
result

t
(1 row)
```



- <@

Description: Specifies whether an array is contained in another.

Example:

```
openGauss=# SELECT ARRAY[2,7] <@ ARRAY[1,7,4,2,6] AS RESULT;
result

t
(1 row)
```

- &&

Description: Specifies whether an array overlaps another (have common elements).

Example:

```
openGauss=# SELECT ARRAY[1,4,3] && ARRAY[2,1] AS RESULT;
result

t
(1 row)
```

- ||

Description: Array-to-array concatenation

Example:

```
openGauss=# SELECT ARRAY[1,2,3] || ARRAY[4,5,6] AS RESULT;
result

{1,2,3,4,5,6}
(1 row)
openGauss=# SELECT ARRAY[1,2,3] || ARRAY[[4,5,6],[7,8,9]] AS RESULT;
result

{{1,2,3},{4,5,6},{7,8,9}}
(1 row)
```

- ||

Description: Element-to-array concatenation

Example:

```
openGauss=# SELECT 3 || ARRAY[4,5,6] AS RESULT;
result

{3,4,5,6}
(1 row)
```

- ||

Description: Array-to-element concatenation

Example:

```
openGauss=# SELECT ARRAY[4,5,6] || 7 AS RESULT;
result

{4,5,6,7}
(1 row)
```

Array comparisons compare the array contents element-by-element, using the default B-tree comparison function for the element data type. In multidimensional arrays, the elements are accessed in row-major order. If the contents of two arrays are equal but the dimensionality is different, the first difference in the dimensionality information determines the sort order.

## Array Functions

- `array_append(anyarray, anyelement)`

Description: Appends an element to the end of an array, and only supports one-dimensional arrays.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_append(ARRAY[1,2], 3) AS RESULT;
result

{1,2,3}
(1 row)
```

- `array_prepend(anyelement, anyarray)`

Description: Appends an element to the beginning of an array, and only supports one-dimensional arrays.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_prepend(1, ARRAY[2,3]) AS RESULT;
result

{1,2,3}
(1 row)
```

- `array_cat(anyarray, anyarray)`

Description: Concatenates two arrays, and supports multi-dimensional arrays.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_cat(ARRAY[1,2,3], ARRAY[4,5]) AS RESULT;
result

{1,2,3,4,5}
(1 row)

openGauss=# SELECT array_cat(ARRAY[[1,2],[4,5]], ARRAY[6,7]) AS RESULT;
result

{{1,2},{4,5},{6,7}}
(1 row)
```

- `array_union(anyarray, anyarray)`

Description: Concatenates two arrays, and supports only one-dimensional arrays.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_union(ARRAY[1,2,3], ARRAY[3,4,5]) AS RESULT;
result

{1,2,3,3,4,5}
(1 row)
```

- `array_union_distinct(anyarray, anyarray)`

Description: Concatenates two arrays and deduplicates them. Only one-dimensional arrays are supported.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_union_distinct(ARRAY[1,2,3], ARRAY[3,4,5]) AS RESULT;
result

{1,2,3,4,5}
(1 row)
```

- `array_intersect(anyarray, anyarray)`

Description: Intersects two arrays. Only one-dimensional arrays are supported.

Return type: anyarray

Example:

```
openGauss=# SELECT array_intersect(ARRAY[1,2,3], ARRAY[3,4,5]) AS RESULT;
result

{3}
(1 row)
```

- `array_intersect_distinct(anyarray, anyarray)`

Description: Intersects two arrays and deduplicates them. Only one-dimensional arrays are supported.

Return type: anyarray

Example:

```
openGauss=# SELECT array_intersect_distinct(ARRAY[1,2,2], ARRAY[2,2,4,5]) AS RESULT;
result

{2}
(1 row)
```

- `array_except(anyarray, anyarray)`

Description: Calculates the difference between two arrays. Only one-dimensional arrays are supported.

Return type: anyarray

Example:

```
openGauss=# SELECT array_except(ARRAY[1,2,3], ARRAY[3,4,5]) AS RESULT;
result

{1,2}
(1 row)
```

- `array_except_distinct(anyarray, anyarray)`

Description: Calculates the difference between two arrays and deduplicates them. Only one-dimensional arrays are supported.

Return type: anyarray

Example:

```
openGauss=# SELECT array_except_distinct(ARRAY[1,2,2,3], ARRAY[3,4,5]) AS RESULT;
result

{1,2}
(1 row)
```

- `array_ndims(anyarray)`

Description: Returns the number of dimensions of an array.

Return type: int

Example:

```
openGauss=# SELECT array_ndims(ARRAY[[1,2,3], [4,5,6]]) AS RESULT;
result

2
(1 row)
```

- `array_dims(anyarray)`

Description: Returns the low-order flag bits and high-order flag bits of each dimension in an array.

Return type: text

Example:

```
openGauss=# SELECT array_dims(ARRAY[[1,2,3], [4,5,6]]) AS RESULT;
result

[1:2][1:3]
(1 row)
```

- `array_length(anyarray, int)`

Description: Returns the length of the requested array dimension. **int** is the requested array dimension.

Return type: int

Example:

```
openGauss=# SELECT array_length(array[1,2,3], 1) AS RESULT;
result

3
(1 row)

openGauss=# SELECT array_length(array[[1,2,3],[4,5,6]], 2) AS RESULT;
result

3
(1 row)
```

- `array_lower(anyarray, int)`

Description: Returns lower bound of the requested array dimension. **int** is the requested array dimension.

Return type: int

Example:

```
openGauss=# SELECT array_lower('[0:2]={1,2,3}::int[]', 1) AS RESULT;
result

0
(1 row)
```

- `array_upper(anyarray, int)`

Description: Returns upper bound of the requested array dimension. **int** is the requested array dimension.

Return type: int

Example:

```
openGauss=# SELECT array_upper(ARRAY[1,8,3,7], 1) AS RESULT;
result

4
(1 row)
```

- `array_upper(anyarray, int)`

Description: Returns upper bound of the requested array dimension. **int** is the requested array dimension.

Return type: int

Example:

```
openGauss=# SELECT array_upper(ARRAY[1,8,3,7], 1) AS RESULT;
result
```

```

4
(1 row)
```

- `array_remove(anyarray, anyelement)`

Description: Removes all specified elements from an array. Only one-dimensional arrays are supported.

Return type: anyarray

Example:

```
openGauss=# SELECT array_remove(ARRAY[1,8,8,7], 8) AS RESULT;
result

{1,7}
(1 row)
```

- `array_to_string(anyarray, text [, text])`

Description: Uses the first **text** as the new delimiter and the second **text** to replace **NULL** values.

Return type: text

Example:

```
openGauss=# SELECT array_to_string(ARRAY[1, 2, 3, NULL, 5], ',', '**') AS RESULT;
result

1,2,3,*,5
(1 row)
```

- `array_delete(anyarray)`

Description: Clears elements in an array and returns an empty array of the same type.

Return type: anyarray

Example:

```
openGauss=# SELECT array_delete(ARRAY[1,8,3,7]) AS RESULT;
result

{}
(1 row)
```

- `array_deleteidx(anyarray, int)`

Description: Deletes specified index elements from an array and returns an array consisting of the remaining elements.

Return type: anyarray

Example:

```
openGauss=# SELECT array_deleteidx(ARRAY[1,2,3,4,5], 1) AS RESULT;
result

{2,3,4,5}
(1 row)
```

- `array_extendnull(anyarray, int)`

Description: Adds a specified number of null elements to the end of an array.

Return type: anyarray

Example:

```
openGauss=# SELECT array_extendnull(ARRAY[1,8,3,7],1) AS RESULT;
result

{1,8,3,7,null}
(1 row)
```

- `array_trim(anyarray, int)`

Description: Deletes a specified number of elements from the end of an array.

Return type: `anyarray`

Example:

```
openGauss=# SELECT array_trim(ARRAY[1,8,3,7],1) AS RESULT;
result

{1,8,3}
(1 row)
```

- `array_exists(anyarray, int)`

Description: Checks whether the second parameter is a valid index of an array.

Return type: `Boolean`

Example:

```
openGauss=# SELECT array_exists(ARRAY[1,8,3,7],1) AS RESULT;
result

t
(1 row)
```

- `array_next(anyarray, int)`

Description: Returns the index of the element following a specified index in an array based on the second input parameter.

Return type: `int`

Example:

```
openGauss=# SELECT array_next(ARRAY[1,8,3,7],1) AS RESULT;
result

2
(1 row)
```

- `array_prior(anyarray, int)`

Description: Returns the index of the element followed by a specified index in an array based on the second input parameter.

Return type: `int`

Example:

```
openGauss=# SELECT array_prior(ARRAY[1,8,3,7],2) AS RESULT;
result

1
(1 row)
```

- `string_to_array(text, text [, text])`

Description: Uses the second **text** as the new delimiter and the third **text** as the substring to be replaced by **NULL** values. A substring can be replaced by **NULL** values only when it is the same as the third **text**.

Return type: `text[]`

Example:

```
openGauss=# SELECT string_to_array('xx~^~yy~^~zz', '~^~', 'yy') AS RESULT;
result

{xx,NULL,zz}
(1 row)
openGauss=# SELECT string_to_array('xx~^~yy~^~zz', '~^~', 'y') AS RESULT;
result

```

```
{xx,yy,zz}
(1 row)
```

- **unnest(anyarray)**

Description: Expands an array to a set of rows.

Return type: setof anyelement

Example:

```
openGauss=# SELECT unnest(ARRAY[1,2]) AS RESULT;
result

1
2
(2 rows)
```

In **string\_to\_array**, if the delimiter parameter is NULL, each character in the input string will become a separate element in the resulting array. If the delimiter is an empty string, then the entire input string is returned as a one-element array. Otherwise the input string is split at each occurrence of the delimiter string.

In **string\_to\_array**, if the null-string parameter is omitted or NULL, none of the substrings of the input will be replaced by NULL.

In **array\_to\_string**, if the null-string parameter is omitted or NULL, any null elements in the array are simply skipped and not represented in the output string.

- **\_pg\_keysequal**

Description: Checks whether two smallint arrays are the same.

Parameter: smallint[], smallint[]

Return type: Boolean

#### NOTE

This function exists in the information\_schema namespace.

## 7.5.16 Range Functions and Operators

### Range Operators

- **=**

Description: Equals

Example:

```
openGauss=# SELECT int4range(1,5) = '[1,4]::int4range AS RESULT;
result

t
(1 row)
```

- **<>**

Description: Does not equal to

Example:

```
openGauss=# SELECT numrange(1.1,2.2) <> numrange(1.1,2.3) AS RESULT;
result

t
(1 row)
```

- **<**

Description: Is less than

Example:

```
openGauss=# SELECT int4range(1,10) < int4range(2,3) AS RESULT;
result

t
(1 row)
```

- >

Description: Is greater than

Example:

```
openGauss=# SELECT int4range(1,10) > int4range(1,5) AS RESULT;
result

t
(1 row)
```

- <=

Description: Is less than or equals

Example:

```
openGauss=# SELECT numrange(1.1,2.2) <= numrange(1.1,2.2) AS RESULT;
result

t
(1 row)
```

- >=

Description: Is greater than or equals

Example:

```
openGauss=# SELECT numrange(1.1,2.2) >= numrange(1.1,2.0) AS RESULT;
result

t
(1 row)
```

- @>

Description: Contains ranges

Example:

```
openGauss=# SELECT int4range(2,4) @> int4range(2,3) AS RESULT;
result

t
(1 row)
```

- @>

Description: Contains elements

Example:

```
openGauss=# SELECT '[2011-01-01,2011-03-01]::tsrange @> '2011-01-10'::timestamp AS RESULT;
result

t
(1 row)
```

- <@

Description: Range is contained by

Example:

```
openGauss=# SELECT int4range(2,4) <@ int4range(1,7) AS RESULT;
result

t
(1 row)
```



- <@

Description: Element is contained by

Example:

```
openGauss=# SELECT 42 <@ int4range(1,7) AS RESULT;
result

f
(1 row)
```
- &&

Description: Overlap (have points in common)

Example:

```
openGauss=# SELECT int8range(3,7) && int8range(4,12) AS RESULT;
result

t
(1 row)
```
- <<

Description: Strictly left of

Example:

```
openGauss=# SELECT int8range(1,10) << int8range(100,110) AS RESULT;
result

t
(1 row)
```
- >>

Description: Strictly right of

Example:

```
openGauss=# SELECT int8range(50,60) >> int8range(20,30) AS RESULT;
result

t
(1 row)
```
- &<

Description: Does not extend to the right of

Example:

```
openGauss=# SELECT int8range(1,20) &< int8range(18,20) AS RESULT;
result

t
(1 row)
```
- &>

Description: Does not extend to the left of

Example:

```
openGauss=# SELECT int8range(7,20) &> int8range(5,10) AS RESULT;
result

t
(1 row)
```
- -|-

Description: Is adjacent to

Example:

```
openGauss=# SELECT numrange(1.1,2.2) -|- numrange(2.2,3.3) AS RESULT;
result
```

- t  
(1 row)
- +  
Description: Union  
Example:  
openGauss=# SELECT numrange(5,15) + numrange(10,20) AS RESULT;  
result  
-----  
[5,20)  
(1 row)
- \*  
Description: Intersection  
Example:  
openGauss=# SELECT int8range(5,15) \* int8range(10,20) AS RESULT;  
result  
-----  
[10,15)  
(1 row)
- -  
Description: Difference  
Example:  
openGauss=# SELECT int8range(5,15) - int8range(10,20) AS RESULT;  
result  
-----  
[5,10)  
(1 row)

The simple comparison operators `<`, `>`, `<=`, and `>=` compare the lower bounds first, and only if those are equal, compare the upper bounds.

The `<<`, `>>`, and `-|-` operators always return false when an empty range is involved; that is, an empty range is not considered to be either before or after any other range.

The union and difference operators will fail if the resulting range would need to contain two disjoint sub-ranges.

## Range Functions

The lower and upper functions return null if the range is empty or the requested bound is infinite. The `lower_inc`, `upper_inc`, `lower_inf`, and `upper_inf` functions all return false for an empty range.

- `numrange(numeric, numeric, [text])`  
Description: Specifies a range.  
Return type: Range's element type  
Example:  
openGauss=# SELECT numrange(1.1,2.2) AS RESULT;  
result  
-----  
[1.1,2.2)  
(1 row)  
openGauss=# SELECT numrange(1.1,2.2, '()') AS RESULT;  
result  
-----

(1.1,2.2)  
(1 row)

- **lower(anyrange)**

Description: Lower bound of a range

Return type: Range's element type

Example:

```
openGauss=# SELECT lower(numrange(1.1,2.2)) AS RESULT;
result

1.1
(1 row)
```

- **upper(anyrange)**

Description: Upper bound of a range

Return type: Range's element type

Example:

```
openGauss=# SELECT upper(numrange(1.1,2.2)) AS RESULT;
result

2.2
(1 row)
```

- **isempty(anyrange)**

Description: Is the range empty?

Return type: Boolean

Example:

```
openGauss=# SELECT isempty(numrange(1.1,2.2)) AS RESULT;
result

f
(1 row)
```

- **lower\_inc(anyrange)**

Description: Is the lower bound inclusive?

Return type: Boolean

Example:

```
openGauss=# SELECT lower_inc(numrange(1.1,2.2)) AS RESULT;
result

t
(1 row)
```

- **upper\_inc(anyrange)**

Description: Is the upper bound inclusive?

Return type: Boolean

Example:

```
openGauss=# SELECT upper_inc(numrange(1.1,2.2)) AS RESULT;
result

f
(1 row)
```

- **lower\_inf(anyrange)**

Description: Is the lower bound infinite?

Return type: Boolean

Example:

```
openGauss=# SELECT lower_inf('(',')::daterange) AS RESULT;
result

t
(1 row)
```

- `upper_inf(anyrange)`

Description: Is the upper bound infinite?

Return type: Boolean

Example:

```
openGauss=# SELECT upper_inf('(',')::daterange) AS RESULT;
result

t
(1 row)
```

- `elem_contained_by_range(anyelement, anyrange)`

Description: Determines whether an element is within the range.

Return type: Boolean

Example:

```
openGauss=# SELECT elem_contained_by_range('2', numrange(1.1,2.2));
elem_contained_by_range

t
(1 row)
```

## 7.5.17 Aggregate Functions

### Aggregate Functions

- `sum(expression)`

Description: Specifies the sum of expressions across all input values.

Return type:

Generally, same as the argument data type. In the following cases, type conversion occurs:

- **BIGINT** for **SMALLINT** or **INT** arguments
- **NUMBER** for **BIGINT** arguments
- **DOUBLE PRECISION** for floating-point arguments

Example:

```
openGauss=# CREATE TABLE tab(a int);
CREATE TABLE
openGauss=# INSERT INTO tab values(1);
INSERT 0 1
openGauss=# INSERT INTO tab values(2);
INSERT 0 1
openGauss=# SELECT sum(a) FROM tab;
sum

3
(1 row)
```

- `max(expression)`

Description: Specifies the maximum value of expression across all input values.

Parameter type: any array, numeric, string, or date/time type

Return type: same as the argument type

Example:

```
openGauss=# CREATE TABLE max_t1(a int, b int);

openGauss=# INSERT INTO max_t1 VALUES(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT MAX(a) FROM max_t1;
max

4
(1 row)

openGauss=# DROP TABLE max_t1;
```

- **min(expression)**

Description: Specifies the minimum value of expression across all input values.

Parameter type: any array, numeric, string, or date/time type

Return type: same as the argument type

Example:

```
openGauss=# CREATE TABLE min_t1(a int, b int);

openGauss=# INSERT INTO min_t1 VALUES(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT MIN(a) FROM min_t1;
min

1
(1 row)

openGauss=# DROP TABLE min_t1;
```

- **avg(expression)**

Description: Specifies the average (arithmetic mean) of all input values.

Return type:

**NUMBER** for any integer-type argument.

**DOUBLE PRECISION** for floating-point arguments.

Otherwise, it is the same as the input data type.

Example:

```
openGauss=# CREATE TABLE avg_t1(a int, b int);

openGauss=# INSERT INTO avg_t1 VALUES(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT AVG(a) FROM avg_t1;
avg

2.5000000000000000
(1 row)

openGauss=# DROP TABLE avg_t1;
```

- **count(expression)**

Description: Specifies the number of input rows for which the value of the expression is not null.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE count_t1(a int, b int);

openGauss=# INSERT INTO count_t1 VALUES (NULL,1),(1,2),(2,3),(3,4),(4,5);
```

```
openGauss=# SELECT COUNT(a) FROM count_t1;
count

 4
(1 row)

openGauss=# DROP TABLE count_t1;
```

- **count(\*)**

Description: Returns the number of input rows.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE count_t1(a int, b int);

openGauss=# INSERT INTO count_t1 VALUES (NULL,1),(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT COUNT(*) FROM count_t1;
count

 5
(1 row)

openGauss=# DROP TABLE count_t1;
```

- **median(expression) [over (query partition clause)]**

Description: Returns the median of an expression. **NULL** will be ignored by the median function during calculation. The **DISTINCT** keyword can be used to exclude duplicate records in an expression. The data type of the input expression can be numeric (including integer, double, and bigint) or interval. For other data types, the median cannot be calculated.

Return type: double or interval

Example:

```
SELECT median(id) FROM (values(1), (2), (3), (4), (null)) test(id);
median

 2.5
(1 row)
```

- **array\_agg(expression)**

Description: Concatenates input values, including nulls, into an array.

Return type: array of the argument type

Example:

```
openGauss=# CREATE TABLE array_agg_t1(a int, b int);

openGauss=# INSERT INTO array_agg_t1 VALUES (NULL,1),(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT ARRAY_AGG(a) FROM array_agg_t1;
array_agg

{NULL,1,2,3,4}
(1 row)

openGauss=# DROP TABLE array_agg_t1;
```

- **string\_agg(expression, delimiter)**

Description: Concatenates input values into a string, separated by delimiter.

Return type: same as the argument type

Example:

```
openGauss=# CREATE TABLE string_agg_t1(a int, b int);

openGauss=# INSERT INTO string_agg_t1 VALUES (NULL,1),(1,2),(2,3),(3,4),(4,5);

openGauss=# SELECT STRING_AGG(a,',') FROM string_agg_t1;
 string_agg

1;2;3;4
(1 row)

openGauss=# DROP TABLE string_agg_t1;
```

- **listagg(expression [, delimiter]) WITHIN GROUP(ORDER BY order-list)**

Description: Sorts aggregation column data according to the mode specified by **WITHIN GROUP** and concatenates the data to a string using the specified delimiter.

- **expression**: Required. It specifies an aggregation column name or a column-based valid expression. It does not support the **DISTINCT** keyword and the **VARIADIC** parameter.
- **delimiter**: Optional. It specifies a delimiter, which can be a string constant or a deterministic expression based on a group of columns. The default value is empty.
- **order-list**: Required. It specifies the sorting mode in a group.

Return type: text

Example:

The aggregation column is of the text character set type.

```
openGauss=# CREATE TABLE listagg_t1(a int, b text);

openGauss=# INSERT INTO listagg_t1 VALUES (NULL,'a1'),(1,'b2'),(1,'c3'),(2,'d4'),(2,'e5'),(3,'f6');

openGauss=# SELECT a,LISTAGG(b,',') WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
 a | listagg
---+-----
 1 | b2;c3
 2 | d4;e5
 3 | f6
 | a1
(4 rows)

openGauss=# DROP TABLE listagg_t1;
```

The aggregation column is of the integer type.

```
openGauss=# CREATE TABLE listagg_t1(a int, b int);

openGauss=# INSERT INTO listagg_t1 VALUES (NULL,1),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT a,LISTAGG(b,',') WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
 a | listagg
---+-----
 1 | 2;3
 2 | 4;5
 3 | 6
 | 1
(4 rows)

openGauss=# DROP TABLE listagg_t1;
```

The aggregation column is of the floating point type.

```
openGauss=# CREATE TABLE listagg_t1(a int, b float);

openGauss=# INSERT INTO listagg_t1 VALUES (NULL,1.111),(1,2.222),(1,3.333),(2,4.444),(2,5.555),
(3,6.666);
```

```
openGauss=# SELECT a,LISTAGG(b,',') WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
a | listagg
-----+-----
1 | 2.222000;3.333000
2 | 4.444000;5.555000
3 | 6.666000
 | 1.111000
(4 rows)
```

```
openGauss=# DROP TABLE listagg_t1;
```

The aggregation column is of the time type.

```
openGauss=# CREATE TABLE listagg_t1(a int, b timestamp);
```

```
openGauss=# INSERT INTO listagg_t1 VALUES (NULL,'2000-01-01'),(1,'2000-02-02'),(1,'2000-03-03'),
(2,'2000-04-04'),(2,'2000-05-05'),(3,'2000-06-06');
```

```
openGauss=# SELECT a,LISTAGG(b,',') WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
a | listagg
-----+-----
1 | 2000-02-02 00:00:00;2000-03-03 00:00:00
2 | 2000-04-04 00:00:00;2000-05-05 00:00:00
3 | 2000-06-06 00:00:00
 | 2000-01-01 00:00:00
(4 rows)
```

```
openGauss=# DROP TABLE listagg_t1;
```

The aggregation column is of the time interval type.

```
openGauss=# CREATE TABLE listagg_t1(a int, b interval);
```

```
openGauss=# INSERT INTO listagg_t1 VALUES (NULL,'1 days'),(1,'2 days'),(1,'3 days'),(2,'4 days'),(2,'5
days'),(3,'6 days');
```

```
openGauss=# SELECT a,LISTAGG(b,',') WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
a | listagg
-----+-----
1 | 2 days;3 days
2 | 4 days;5 days
3 | 6 days
 | 1 day
(4 rows)
```

```
openGauss=# DROP TABLE listagg_t1;
```

By default, the delimiter is empty.

```
openGauss=# CREATE TABLE listagg_t1(a int, b interval);
```

```
openGauss=# INSERT INTO listagg_t1 VALUES (NULL,'1 days'),(1,'2 days'),(1,'3 days'),(2,'4 days'),(2,'5
days'),(3,'6 days');
```

```
openGauss=# SELECT a,LISTAGG(b) WITHIN GROUP(ORDER BY b) FROM listagg_t1 group by a;
a | listagg
-----+-----
1 | 2 days3 days
2 | 4 days5 days
3 | 6 days
 | 1 day
(4 rows)
```

```
openGauss=# DROP TABLE listagg_t1;
```

When **listagg** is used as a window function, the **OVER** clause does not support the window sorting of **ORDER BY**, and the **listagg** column is an ordered aggregation of the corresponding groups.

```
openGauss=# CREATE TABLE listagg_t1(a int, b interval);
```

```
openGauss=# INSERT INTO listagg_t1 VALUES (NULL,'1 days'),(1,'2 days'),(1,'3 days'),(2,'4 days'),(2,'5
days'),(3,'6 days');
```



```
openGauss=# SELECT a,LISTAGG(b) WITHIN GROUP(ORDER BY b) OVER(PARTITION BY a) FROM
listagg_t1;
a | listagg
-----+-----
1 | 2 days3 days
1 | 2 days3 days
2 | 4 days5 days
2 | 4 days5 days
3 | 6 days
 | 1 day
(6 rows)

openGauss=# DROP TABLE listagg_t1;
```

- covar\_pop(Y, X)

Description: Specifies the overall covariance.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE covar_pop_t1(a int, b int);

openGauss=# INSERT INTO covar_pop_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT COVAR_POP(a,b) FROM covar_pop_t1;
covar_pop

 100
(1 row)

openGauss=# DROP TABLE covar_pop_t1;
```

- covar\_samp(Y, X)

Description: Specifies the sample covariance.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE covar_samp_t1(a int, b int);

openGauss=# INSERT INTO covar_samp_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT COVAR_SAMP(a,b) FROM covar_samp_t1;
covar_samp

 125
(1 row)

openGauss=# DROP TABLE covar_samp_t1;
```

- stddev\_pop(expression)

Description: Specifies the overall standard deviation.

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE stddev_pop_t1(a int, b int);

openGauss=# INSERT INTO stddev_pop_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT STDDEV_POP(a) FROM stddev_pop_t1;
stddev_pop

7.4833147735478828
(1 row)

openGauss=# DROP TABLE stddev_pop_t1;
```

- **stddev\_samp(expression)**

Description: Specifies the sample standard deviation of the input values.

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE stddev_samp_t1(a int, b int);

openGauss=# INSERT INTO stddev_samp_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT STDDEV_SAMP(a) FROM stddev_samp_t1;
 stddev_samp

8.3666002653407555
(1 row)

openGauss=# DROP TABLE stddev_samp_t1;
```

- **var\_pop(expression)**

Description: Specifies the population variance of the input values (square of the population standard deviation).

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE var_pop_t1(a int, b int);

openGauss=# INSERT INTO var_pop_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT VAR_POP(a) FROM var_pop_t1;
 var_pop

56.0000000000000000
(1 row)

openGauss=# DROP TABLE var_pop_t1;
```

- **var\_samp(expression)**

Description: Specifies the sample variance of the input values (square of the sample standard deviation).

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE var_samp_t1(a int, b int);

openGauss=# INSERT INTO var_samp_t1 VALUES (NULL,11),(11,21),(11,31),(21,41),(21,51),(31,61);

openGauss=# SELECT VAR_SAMP(a) FROM var_samp_t1;
 var_samp

70.0000000000000000
(1 row)

openGauss=# DROP TABLE var_samp_t1;
```

- **bit\_and(expression)**

Description: bitwise AND of all non-null input values, or null if none

Return type: same as the argument type

Example:

```
openGauss=# CREATE TABLE bit_and_t1(a int, b int);
```

```
openGauss=# INSERT INTO bit_and_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT BIT_AND(a) FROM bit_and_t1;
 bit_and

 0
(1 row)

openGauss=# DROP TABLE bit_and_t1;
```

- **bit\_or(expression)**

Description: bitwise OR of all non-null input values, or null if none

Return type: same as the argument type

Example:

```
openGauss=# CREATE TABLE bit_or_t1(a int, b int);

openGauss=# INSERT INTO bit_or_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT BIT_OR(a) FROM bit_or_t1;
 bit_or

 3
(1 row)

openGauss=# DROP TABLE bit_or_t1;
```

- **bool\_and(expression)**

Description: Its value is **true** if all input values are **true**, otherwise **false**.

Return type: Boolean

Example:

```
openGauss=# SELECT bool_and(100 < 2500);
 bool_and

 t
(1 row)
```

- **bool\_or(expression)**

Description: Its value is **true** if at least one input value is **true**, otherwise **false**.

Return type: Boolean

Example:

```
openGauss=# SELECT bool_or(100 < 2500);
 bool_or

 t
(1 row)
```

- **corr(Y, X)**

Description: Specifies the correlation coefficient.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE corr_t1(a int, b int);

openGauss=# INSERT INTO corr_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT CORR(a,b) FROM corr_t1;
 corr

.944911182523068
(1 row)
```

```
openGauss=# DROP TABLE corr_t1;
```

- **every(expression)**

Description: Equivalent to **bool\_and**

Return type: Boolean

Example:

```
openGauss=# SELECT every(100 <2500);
every

t
(1 row)
```

- **regr\_avgx(Y, X)**

Description: Specifies the average of the independent variable (**sum(X)/Y**).

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_t1(a int, b int);

openGauss=# INSERT INTO regr_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_AVGX(a,b) FROM regr_t1;
regr_avgx

4
(1 row)

openGauss=# DROP TABLE regr_t1;
```

- **regr\_avgy(Y, X)**

Description: Specifies the average of the dependent variable (**sum(Y)/X**).

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_avgy_t1(a int, b int);

openGauss=# INSERT INTO regr_avgy_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_AVGY(a,b) FROM regr_avgy_t1;
regr_avgy

1.8
(1 row)

openGauss=# DROP TABLE regr_avgy_t1;
```

- **regr\_count(Y, X)**

Description: Specifies the number of input rows in which both expressions are non-null.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE regr_count_t1(a int, b int);

openGauss=# INSERT INTO regr_count_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_COUNT(a,b) FROM regr_count_t1;
regr_count

5
(1 row)

openGauss=# DROP TABLE regr_count_t1;
```

- **regr\_intercept(Y, X)**

Description: Specifies the y-intercept of the least-squares-fit linear equation determined by the (X, Y) pairs.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_intercept_t1(a int, b int);

openGauss=# INSERT INTO regr_intercept_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_INTERCEPT(b,a) FROM regr_intercept_t1;
regr_intercept

.785714285714286
(1 row)

openGauss=# DROP TABLE regr_intercept_t1;
```

- **regr\_r2(Y, X)**

Description: Specifies the square of the correlation coefficient.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_r2_t1(a int, b int);

openGauss=# INSERT INTO regr_r2_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_R2(b,a) FROM regr_r2_t1;
regr_r2

.892857142857143
(1 row)

openGauss=# DROP TABLE regr_r2_t1;
```

- **regr\_slope(Y, X)**

Description: Specifies the slope of the least-squares-fit linear equation determined by the (X, Y) pairs.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_slope_t1(a int, b int);

openGauss=# INSERT INTO regr_slope_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_SLOPE(b,a) FROM regr_slope_t1;
regr_slope

1.78571428571429
(1 row)

openGauss=# DROP TABLE regr_slope_t1;
```

- **regr\_sxx(Y, X)**

Description:  $\text{sum}(Y^2) - \text{sum}(X)^2/N$  (sum of squares of the independent variables)

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_sxx_t1(a int, b int);

openGauss=# INSERT INTO regr_sxx_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);
```

```
openGauss=# SELECT REGR_SXX(b,a) FROM regr_sxx_t1;
regr_sxx

 2.8
(1 row)
```

```
openGauss=# DROP TABLE regr_sxx_t1;
```

- **regr\_sxy(Y, X)**

Description:  **$\text{sum}(X*Y) - \text{sum}(X) * \text{sum}(Y)/N$**  ("sum of products" of independent times dependent variable)

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_sxy_t1(a int, b int);

openGauss=# INSERT INTO regr_sxy_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_SXY(b,a) FROM regr_sxy_t1;
regr_sxy

 5
(1 row)

openGauss=# DROP TABLE regr_sxy_t1;
```

- **regr\_syy(Y, X)**

Description:  **$\text{sum}(Y^2) - \text{sum}(X)^2/N$**  ("sum of squares" of the dependent variable)

Return type: double precision

Example:

```
openGauss=# CREATE TABLE regr_syy_t1(a int, b int);

openGauss=# INSERT INTO regr_syy_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT REGR_SYY(b,a) FROM regr_syy_t1;
regr_syy

 10
(1 row)

openGauss=# DROP TABLE regr_syy_t1;
```

- **stddev(expression)**

Description: Specifies the alias of **stddev\_samp**.

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE stddev_t1(a int, b int);

openGauss=# INSERT INTO stddev_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT STDDEV(a) FROM stddev_t1;
stddev

.83666002653407554798
(1 row)

openGauss=# DROP TABLE stddev_t1;
```

- **variance(expexpression,ression)**

Description: Specifies the alias of **var\_samp**.

Return type: **double precision** for floating-point arguments, otherwise **numeric**

Example:

```
openGauss=# CREATE TABLE variance_t1(a int, b int);

openGauss=# INSERT INTO variance_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT VARIANCE(a) FROM variance_t1;
 variance

.700000000000000000000000
(1 row)

openGauss=# DROP TABLE variance_t1;
```

- **delta**

Description: Returns the difference between the current row and the previous row.

Parameter: numeric

Return type: numeric

- **checksum(expression)**

Description: Returns the **CHECKSUM** value of all input values. This function can be used to check whether the data in the tables is the same before and after the backup, restoration, or migration of the GaussDB database (databases other than GaussDB are not supported). Before and after database backup, database restoration, or data migration, you need to manually run SQL commands to obtain the execution results. Compare the obtained execution results to check whether the data in the tables before and after the backup or migration is the same.

 **NOTE**

- For large tables, the execution of the **CHECKSUM** function may take a long time.
- If the **CHECKSUM** values of two tables are different, it indicates that the contents of the two tables are different. Using the hash function in the **CHECKSUM** function may incur conflicts. There is low possibility that two tables with different contents may have the same **CHECKSUM** value. The same problem may occur when **CHECKSUM** is used for columns.
- If the time type is timestamp, timestamptz, or smalldatetime, ensure that the time zone settings are the same when calculating the **CHECKSUM** value.
- If the **CHECKSUM** value of a column is calculated and the column type can be changed to TEXT by default, set *expression* to the column name.
- If the **CHECKSUM** value of a column is calculated and the column type cannot be converted to TEXT by default, set *expression* to *Column name::TEXT*.
- If the **CHECKSUM** value of all columns is calculated, set *expression* to *Table name::TEXT*.

The following types of data can be converted into TEXT types by default: char, name, int8, int2, int1, int4, raw, pg\_node\_tree, float4, float8, bpchar, varchar, nvarchar, nvarchar2, date, timestamp, timestamptz, numeric, and smalldatetime. Other types need to be forcibly converted to TEXT.

Return type: numeric

Example:

The following shows the **CHECKSUM** value of a column that can be converted to the TEXT type by default:

```
openGauss=# CREATE TABLE checksum_t1(a int, b int);

openGauss=# INSERT INTO checksum_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT CHECKSUM(a) FROM checksum_t1;
checksum

18126842830
(1 row)

openGauss=# DROP TABLE checksum_t1;
```

The following shows the **CHECKSUM** value of a column that cannot be converted to the TEXT type by default. Note that the **CHECKSUM** parameter is set to *Column name::TEXT*.

```
openGauss=# CREATE TABLE checksum_t1(a int, b int);

openGauss=# INSERT INTO checksum_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT CHECKSUM(a::TEXT) FROM checksum_t1;
checksum

18126842830
(1 row)

openGauss=# DROP TABLE checksum_t1;
```

The following shows the **CHECKSUM** value of all columns in a table. Note that the **CHECKSUM** parameter is set to *Table name::TEXT*. The table name is not modified by its schema.

```
openGauss=# CREATE TABLE checksum_t1(a int, b int);

openGauss=# INSERT INTO checksum_t1 VALUES (NULL,11),(1,2),(1,3),(2,4),(2,5),(3,6);

openGauss=# SELECT CHECKSUM(checksum_t1::TEXT) FROM checksum_t1;
checksum

11160522226
(1 row)

openGauss=# DROP TABLE checksum_t1;
```

- **mode()** WITHIN GROUP (ORDER BY value anyelement)

Description: Returns the value with the highest occurrence frequency in a column. If multiple values have the same frequency, the smallest value is returned. The sorting mode is the same as the default sorting mode of the column type. **value** is an input parameter and can be of any type.

Return type: same as the input parameter type

Example:

```
openGauss=# SELECT mode() WITHIN GROUP (ORDER BY value) FROM (values(1, 'a'), (2, 'b'), (2, 'c')) v(value, tag);
mode

2
(1 row)
openGauss=# SELECT mode() WITHIN GROUP (ORDER BY tag) FROM (values(1, 'a'), (2, 'b'), (2, 'c')) v(value, tag);
mode

a
(1 row)
```



## 7.5.18 Window Functions

### Window Functions

Window functions and the **OVER** clause are used together. The **OVER** clause is used for grouping data and sorting the elements in a group. Window functions are used for generating sequence numbers for the values in the group.

#### NOTE

**order by** in a window function must be followed by a column name. If it is followed by a number, the number is processed as a constant value and the target column is not ranked.

- **RANK()**

Description: Generates non-consecutive sequence numbers for the values in each group. The same values have the same sequence number.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE rank_t1(a int, b int);

openGauss=# INSERT INTO rank_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,RANK() OVER(PARTITION BY a ORDER BY b) FROM rank_t1;
a | b | rank
---+---+-----
1 | 1 | 1
1 | 1 | 1
1 | 2 | 3
1 | 3 | 4
2 | 4 | 1
2 | 5 | 2
3 | 6 | 1
(7 rows)

openGauss=# DROP TABLE rank_t1;
```

- **ROW\_NUMBER()**

Description: Generates consecutive sequence numbers for the values in each group. The same values have different sequence numbers.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE row_number_t1(a int, b int);

openGauss=# INSERT INTO row_number_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,ROW_NUMBER() OVER(PARTITION BY a ORDER BY b) FROM
row_number_t1;
a | b | row_number
---+---+-----
1 | 1 | 1
1 | 1 | 2
1 | 2 | 3
1 | 3 | 4
2 | 4 | 1
2 | 5 | 2
3 | 6 | 1
(7 rows)

openGauss=# DROP TABLE row_number_t1;
```

- **DENSE\_RANK()**

Description: Generates consecutive sequence numbers for the values in each group. The same values have the same sequence number.

Return type: bigint

Example:

```
openGauss=# CREATE TABLE dense_rank_t1(a int, b int);

openGauss=# INSERT INTO dense_rank_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,DENSE_RANK() OVER(PARTITION BY a ORDER BY b) FROM dense_rank_t1;
 a | b | dense_rank
---+---+-----
 1 | 1 | 1
 1 | 1 | 1
 1 | 2 | 2
 1 | 3 | 3
 2 | 4 | 1
 2 | 5 | 2
 3 | 6 | 1
(7 rows)

openGauss=# DROP TABLE dense_rank_t1;
```

- **PERCENT\_RANK()**

Description: The PERCENT\_RANK function is used for generating corresponding sequence numbers for the values in each group. That is, the function calculates the value according to the formula: Sequence number =  $(\text{rank} - 1) / (\text{totalrows} - 1)$ . **rank** is the corresponding sequence number generated based on the **RANK** function for the value and **totalrows** is the total number of elements in a group.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE percent_rank_t1(a int, b int);

openGauss=# INSERT INTO percent_rank_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,PERCENT_RANK() OVER(PARTITION BY a ORDER BY b) FROM
percent_rank_t1;
 a | b | percent_rank
---+---+-----
 1 | 1 | 0
 1 | 1 | 0
 1 | 2 | .6666666666666667
 1 | 3 | 1
 2 | 4 | 0
 2 | 5 | 1
 3 | 6 | 0
(7 rows)

openGauss=# DROP TABLE percent_rank_t1;
```

- **CUME\_DIST()**

Description: Generates accumulative distribution sequence numbers for the values in each group. That is, the function calculates the value according to the following formula: Sequence number = Number of rows preceding or peer with current row/Total rows.

Return type: double precision

Example:

```
openGauss=# CREATE TABLE cume_dist_t1(a int, b int);

openGauss=# INSERT INTO cume_dist_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);
```

```
openGauss=# SELECT a,b,CUME_DIST() OVER(PARTITION BY a ORDER BY b) FROM cume_dist_t1;
a | b | cume_dist
---+---+-----
1 | 1 | .5
1 | 1 | .5
1 | 2 | .75
1 | 3 | 1
2 | 4 | .5
2 | 5 | 1
3 | 6 | 1
(7 rows)

openGauss=# DROP TABLE cume_dist_t1;
```

- NTILE(num\_buckets integer)

Description: Equally allocates sequential data sets to the buckets whose quantity is specified by **num\_buckets** according to **num\_buckets integer** and allocates the bucket number to each row. Divide the partition as evenly as possible.

Return type: integer

Example:

```
openGauss=# CREATE TABLE ntile_t1(a int, b int);

openGauss=# INSERT INTO ntile_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,NTILE(2) OVER(PARTITION BY a ORDER BY b) FROM ntile_t1;
a | b | ntile
---+---+-----
1 | 1 | 1
1 | 1 | 1
1 | 2 | 2
1 | 3 | 2
2 | 4 | 1
2 | 5 | 2
3 | 6 | 1
(7 rows)

openGauss=# DROP TABLE ntile_t1;
```

- LAG(value any [, offset integer [, default any ]])

Description: Generates lag values for the corresponding values in each group. That is, the value of the row obtained by moving forward the row corresponding to the current value by **offset** (integer) is the sequence number. If the row does not exist after the moving, the result value is the default value. If omitted, **offset** defaults to **1** and **default** to **NULL**. The type of the **default** value must be the same as that of the **value** value.

Return type: same as the parameter type

Example:

```
-- Create a table and insert data into the table.
openGauss=# CREATE TABLE ta1 (hire_date date, last_name varchar(20), department_id int);
CREATE TABLE
openGauss=# INSERT INTO ta1 values('07-DEC-02', 'Raphaely', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('24-JUL-05', 'Tobias', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('24-DEC-05', 'Baida', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('18-MAY-03', 'Khoo', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('15-NOV-06', 'Himuro', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-AUG-07', 'Colmenares', 30);
```

```
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-MAY-07', 'yq', 11);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-MAY-08', 'zi', 11);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('', 'yq1', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, 'yq2', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-DEC-07', 'yq3', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, null, 11);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, null, 11);
INSERT 0 1

-- Call LAG and set offset to 3 and default to null.
openGauss=# SELECT hire_date, last_name, department_id, lag(hire_date, 3, null) OVER (PARTITION
BY department_id ORDER BY last_name) AS "NextHired" FROM ta1 ORDER BY department_id;
 hire_date | last_name | department_id | NextHired
-----+-----+-----+-----
2007-05-10 00:00:00 | yq | 11 |
2008-05-10 00:00:00 | zi | 11 |
 | | 11 |
 | | 11 | 2007-05-10 00:00:00
2005-12-24 00:00:00 | Baida | 30 |
2007-08-10 00:00:00 | Colmenares | 30 |
2006-11-15 00:00:00 | Himuro | 30 |
2003-05-18 00:00:00 | Khoo | 30 | 2005-12-24 00:00:00
2002-12-07 00:00:00 | Raphaely | 30 | 2007-08-10 00:00:00
2005-07-24 00:00:00 | Tobias | 30 | 2006-11-15 00:00:00
 | yq1 | 30 | 2003-05-18 00:00:00
 | yq2 | 30 | 2002-12-07 00:00:00
2007-12-10 00:00:00 | yq3 | 30 | 2005-07-24 00:00:00
(13 rows)

openGauss=# DROP TABLE ta1;
```

- LEAD(value any [, offset integer [, default any ]])

Description: Generates leading values for the corresponding values in each group. That is, the value of the row obtained by moving backward the row corresponding to the current value by **offset** (integer) is the sequence number. If the row after the moving exceeds the total number of rows for the current group, the result value is the default value. If omitted, **offset** defaults to **1** and **default** to **NULL**. The type of the **default** value must be the same as that of the **value** value.

Return type: same as the parameter type

Example:

```
openGauss=# CREATE TABLE ta1 (hire_date date, last_name varchar(20), department_id int);
CREATE TABLE
openGauss=# INSERT INTO ta1 values('07-DEC-02', 'Raphaely', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('24-JUL-05', 'Tobias', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('24-DEC-05', 'Baida', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('18-MAY-03', 'Khoo', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('15-NOV-06', 'Himuro', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-AUG-07', 'Colmenares', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-MAY-07', 'yq', 11);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-MAY-08', 'zi', 11);
INSERT 0 1
```

```
openGauss=# INSERT INTO ta1 values('yq1', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, 'yq2', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values('10-DEC-07', 'yq3', 30);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, null, 11);
INSERT 0 1
openGauss=# INSERT INTO ta1 values(null, null, 11);
INSERT 0 1

-- Call LEAD and set offset to 2.
openGauss=# SELECT hire_date, last_name, department_id, lead(hire_date, 2) OVER (PARTITION BY
department_id ORDER BY last_name) AS "NextHired" FROM ta1 ORDER BY department_id;
 hire_date | last_name | department_id | NextHired
-----+-----+-----+-----
2007-05-10 00:00:00 | yq | 11 |
2008-05-10 00:00:00 | zi | 11 |
 | | 11 |
 | | 11 |
2005-12-24 00:00:00 | Baida | 30 | 2006-11-15 00:00:00
2007-08-10 00:00:00 | Colmenares | 30 | 2003-05-18 00:00:00
2006-11-15 00:00:00 | Himuro | 30 | 2002-12-07 00:00:00
2003-05-18 00:00:00 | Khoo | 30 | 2005-07-24 00:00:00
2002-12-07 00:00:00 | Raphaely | 30 |
2005-07-24 00:00:00 | Tobias | 30 |
 | yq1 | 30 | 2007-12-10 00:00:00
 | yq2 | 30 |
2007-12-10 00:00:00 | yq3 | 30 |
(13 rows)

openGauss=# DROP TABLE ta1;
```

- **FIRST\_VALUE(value any)**

Description: Returns the first value of each group.

Return type: same as the parameter type

Example:

```
openGauss=# CREATE TABLE first_value_t1(a int, b int);

openGauss=# INSERT INTO first_value_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,FIRST_VALUE(b) OVER(PARTITION BY a ORDER BY b) FROM first_value_t1;
 a | b | first_value
---+---+-----
 1 | 1 | 1
 1 | 1 | 1
 1 | 2 | 1
 1 | 3 | 1
 2 | 4 | 4
 2 | 5 | 4
 3 | 6 | 6
(7 rows)

openGauss=# DROP TABLE first_value_t1;
```

- **LAST\_VALUE(value any)**

Description: Returns the last value of each group.

Return type: same as the parameter type

Example:

```
openGauss=# CREATE TABLE last_value_t1(a int, b int);

openGauss=# INSERT INTO last_value_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);

openGauss=# SELECT a,b,LAST_VALUE(b) OVER(PARTITION BY a ORDER BY b) FROM last_value_t1;
 a | b | last_value
---+---+-----
 1 | 1 | 6
 1 | 1 | 6
 1 | 2 | 6
 1 | 3 | 6
 2 | 4 | 6
 2 | 5 | 6
 3 | 6 | 6
```

```
-----+-----
1 | 1 | 1
1 | 1 | 1
1 | 2 | 2
1 | 3 | 3
2 | 4 | 4
2 | 5 | 5
3 | 6 | 6
(7 rows)
```

```
openGauss=# DROP TABLE last_value_t1;
```

- **DELTA**

Description: Returns the difference between the current row and the previous row.

Parameter: numeric

Return type: numeric

- **NTH\_VALUE(value any, nth integer)**

Description: Returns the *n*th row for a group. If the row does not exist, **NULL** is returned by default.

Return type: same as the parameter type

Example:

```
openGauss=# CREATE TABLE nth_value_t1(a int, b int);
```

```
openGauss=# INSERT INTO nth_value_t1 VALUES(1,1),(1,1),(1, 2),(1, 3),(2, 4),(2, 5),(3,6);
```

```
openGauss=# SELECT a,b,NTH_VALUE(b, 2) OVER(PARTITION BY a order by b) FROM nth_value_t1;
a | b | nth_value
```

```
-----+-----
1 | 1 | 1
1 | 1 | 1
1 | 2 | 1
1 | 3 | 1
2 | 4 | 5
2 | 5 | 5
3 | 6 |
(7 rows)
```

```
openGauss=# DROP TABLE nth_value_t1;
```

## 7.5.19 Security Functions

### Security Functions

- **gs\_encrypt\_aes128(encryptstr,keyststr)**

Description: Encrypts **encryptstr** strings using **keyststr** as the encryption password and returns encrypted strings. The value of **keyststr** ranges from 8 to 16 bytes and contains at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters.

Return type: text

Length of the return value: At least 92 bytes and no more than  $(4 * [Len / 3] + 68)$  bytes, where *Len* indicates the length of the data before encryption (unit: byte).

Example:

```
openGauss=# SELECT gs_encrypt_aes128('MPPDB','Asdf1234');
```

```
gs_encrypt_aes128
```

```

gwditQLQG8NhFw4OuoKhhQJoXojhFLYkjeG0aYdSctLCnIUgkNwvYI04KbuhmcGZp8jWizBdR1vU9Cspjuzl
0lbz12A=
(1 row)
```

**NOTE**

An encryption password is required during the execution of this function. For security purposes, the gsql tool does not record the SQL statements containing the function name in the execution history. That is, the execution history of this function cannot be found in gsql by paging up and down.

- **gs\_encrypt(encryptstr,keyst, encrypttype)**

Description: Encrypts **encryptstr** strings using **keyst** as the encryption password and returns encrypted strings based on **encrypttype**. The value of **keyst** contains 8 to 16 bytes and at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters. The value of **encrypttype** can be **aes128** or **sm4**.

Return type: text

Example:

```
openGauss=# SELECT gs_encrypt('MPPDB','Asdf1234','sm4');
gs_encrypt

ZBzOmaGA4Bb+coyucJ0B8AkIShqc
(1 row)
```

**NOTE**

An encryption password is required during the execution of this function. For security purposes, the gsql tool does not record the SQL statements containing the function name in the execution history. That is, the execution history of this function cannot be found in gsql by paging up and down.

- **gs\_decrypt\_aes128(decryptstr,keyst)**

Description: Decrypts **decrypt** strings using **keyst** as the decryption password and returns decrypted strings. The **keyst** used for decryption must be consistent with that used for encryption. **keyst** cannot be empty.

**NOTE**

This parameter needs to be used with the **gs\_encrypt\_aes128** encryption function.

Return type: text

Example:

```
openGauss=# SELECT
gs_decrypt_aes128('gwditQLQG8NhFw4OuoKhhQJoXojhFLYkjeG0aYdSctLCnIUgkNwvYI04KbuhmcGZp8j
WizBdR1vU9Cspjuzl0lbz12A=','1234');
gs_decrypt_aes128

MPPDB
(1 row)
```

**NOTE**

A decryption password is required during the execution of this function. For security purposes, the gsql tool does not record the SQL statements containing the function name in the execution history. That is, the execution history of this function cannot be found in gsql by paging up and down.

- **gs\_decrypt(decryptstr,keyst,decrypttype)**

Description: Decrypts **decrypt** strings using **keyst** as the decryption password and returns decrypted strings based on **decrypttype**. The **decrypttype** and

**keystr** used for decryption must be consistent with **encrypttype** and **keystr** used for encryption. The value of **keystr** cannot be empty. The value of **decrypttype** can be **aes128** or **sm4**.

This function needs to be used with the **gs\_encrypt** encryption function.

Return type: text

Example:

```
openGauss=# SELECT gs_decrypt('ZBzOmaGA4Bb+coyucJ0B8AkIShq', 'Asdf1234', 'sm4');
gs_decrypt

MPPDB
(1 row)
```

#### NOTE

A decryption password is required during the execution of this function. For security purposes, the gsqr tool does not record the SQL statements containing the function name in the execution history. That is, the execution history of this function cannot be found in gsqr by paging up and down.

- **gs\_password\_deadline()**

Description: Indicates the number of remaining days before the password of the current user expires.

Return type: interval

Example:

```
openGauss=# SELECT gs_password_deadline();
gs_password_deadline

83 days 17:44:32.196094
(1 row)
```

- **gs\_password\_notifytime()**

Description: Specifies the number of days prior to password expiration that a user will receive a reminder.

Return type: int32

- **login\_audit\_messages(BOOLEAN)**

Description: Queries login information about a login user.

Return type: tuple

Example:

- Check the date, time, and IP address of the last successful login.

```
openGauss=> SELECT * FROM login_audit_messages(true);
username | database | logintime | mytype | result | client_conninfo
-----+-----+-----+-----+-----+-----
omm | openGauss | 2020-06-29 21:56:40+08 | login_success | ok | gsqr@[local]
(1 row)
```

- Check the number, date, and time of failed attempts since the previous successful login.

```
openGauss=> SELECT * FROM login_audit_messages(false);
username | database | logintime | mytype | result | client_conninfo
-----+-----+-----+-----+-----+-----
omm | openGauss | 2020-06-29 21:57:55+08 | login_failed | failed | [unknown]@[local]
omm | openGauss | 2020-06-29 21:57:53+08 | login_failed | failed | [unknown]@[local]
(2 rows)
```

- **login\_audit\_messages\_pid**

Description: Queries login information about a login user. Different from **login\_audit\_messages**, this function queries login information based on



**backendid**. Information about subsequent logins of the same user does not alter the query result of previous logins and cannot be found using this function.

Return type: tuple

 **NOTE**

When the thread pool is enabled, **backendid** obtained in the same session may change due to thread switchover. As a result, the return values are different when the function is called for multiple times. You are advised not to call this function when the thread pool is enabled.

Example:

- Check the date, time, and IP address of the last successful login.

```
openGauss=> SELECT * FROM login_audit_messages_pid(true);
username | database | logintime | mytype | result | client_conninfo | backendid
-----+-----+-----+-----+-----+-----+-----
omm | openGauss | 2020-06-29 21:56:40+08 | login_success | ok | gsqr@[local] |
139823109633792
(1 row)
```

- Check the number, date, and time of failed attempts since the previous successful login.

```
openGauss=> SELECT * FROM login_audit_messages_pid(false);
username | database | logintime | mytype | result | client_conninfo | backendid
-----+-----+-----+-----+-----+-----+-----
omm | openGauss | 2020-06-29 21:57:55+08 | login_failed | failed | [unknown]@[local] |
139823109633792
omm | openGauss | 2020-06-29 21:57:53+08 | login_failed | failed | [unknown]@[local] |
139823109633792
(2 rows)
```

- **inet\_server\_addr**

Description: Displays the server IP address.

Return type: inet

Example:

```
openGauss=# SELECT inet_server_addr();
inet_server_addr

10.10.0.13
(1 row)
```

 **NOTE**

- The client IP address 10.10.0.50 and server IP address 10.10.0.13 are used as an example.
- If the database is connected to the local PC, the value is empty.
- **inet\_client\_addr**

Description: Displays the client IP address.

Return type: inet

Example:

```
openGauss=# SELECT inet_client_addr();
inet_client_addr

10.10.0.50
(1 row)
```

 **NOTE**

- The client IP address 10.10.0.50 and server IP address 10.10.0.13 are used as an example.
- If the database is connected to the local PC, the value is empty.

- **pg\_query\_audit**

Description: Views audit logs of the primary database node.

Return type: record

The following table describes return fields.

| Name            | Type                     | Description                                 |
|-----------------|--------------------------|---------------------------------------------|
| time            | timestamp with time zone | Operation time                              |
| type            | text                     | Operation                                   |
| result          | text                     | Operation result                            |
| userid          | oid                      | User ID                                     |
| username        | text                     | Name of the user who performs the operation |
| database        | text                     | Database name                               |
| client_conninfo | text                     | Client connection information               |
| object_name     | text                     | Object name                                 |
| detail_info     | text                     | Operation details                           |
| node_name       | text                     | Node name                                   |
| thread_id       | text                     | Thread ID                                   |
| local_port      | text                     | Local port                                  |
| remote_port     | text                     | Remote port                                 |

- **pg\_delete\_audit**

Description: Deletes audit logs in a specified period.

Return type: void

- **alldigitsmasking**

Description: Specifies the internal function of the masking policy, which is used to anonymize all characters.

Parameter: col text, letter character default '0'

Return type: text

- **creditcardmasking**

Description: Specifies the internal function of the masking policy, which is used to anonymize all credit card information.

Parameter: col text, letter character default 'x'

Return type: text

- randommasking

Description: Specifies the internal function of the masking policy. The random policy is used.

Parameter: col text

Return type: text

- fullemailmasking

Description: Specifies the internal function of the masking policy, which is used to anonymize the text before the last period (.) except the at sign (@).

Parameter: col text, letter character default 'x'

Return type: text

- basicemailmasking

Description: Specifies the internal function of the masking policy, which is used to anonymize the text before the first at sign (@).

Parameter: col text, letter character default 'x'

Return type: text

- shufflemasking

Description: Specifies the internal function of the masking policy, which is used to sort characters out of order.

Parameter: col text

Return type: text

## 7.5.20 Set Returning Functions

### Series Generating Functions

- generate\_series(start, stop)

Description: Generates a series of values, from **start** to **stop** with a step size of one.

Parameter type: int, bigint, numeric

Return type: setof int, setof bigint, setof numeric (same as the parameter type)

- generate\_series(start, stop, step)

Description: Generates a series of values, from **start** to **stop** with a step size of **step**.

Parameter type: int, bigint, numeric

Return type: setof int, setof bigint, setof numeric (same as the parameter type)

- generate\_series(start, stop, step interval)

Description: Generates a series of values, from **start** to **stop** with a step size of **step**.

Parameter type: timestamp or timestamp with time zone

Return type: setof timestamp or setof timestamp with time zone (same as parameter type)

When **step** is positive, zero rows are returned if **start** is greater than **stop**. Conversely, when **step** is negative, zero rows are returned if **start** is less than **stop**. Zero rows are also returned for **NULL** inputs. It is an error for **step** to be zero.

Example:

```
openGauss=# SELECT * FROM generate_series(2,4);
generate_series

 2
 3
 4
(3 rows)

openGauss=# SELECT * FROM generate_series(5,1,-2);
generate_series

 5
 3
 1
(3 rows)

openGauss=# SELECT * FROM generate_series(4,3);
generate_series

(0 rows)

-- This example applies to the date-plus-integer operator.
openGauss=# SELECT current_date + s.a AS dates FROM generate_series(0,14,7) AS s(a);
dates

2017-06-02
2017-06-09
2017-06-16
(3 rows)

openGauss=# SELECT * FROM generate_series('2008-03-01 00:00'::timestamp, '2008-03-04 12:00', '10
hours');
generate_series

2008-03-01 00:00:00
2008-03-01 10:00:00
2008-03-01 20:00:00
2008-03-02 06:00:00
2008-03-02 16:00:00
2008-03-03 02:00:00
2008-03-03 12:00:00
2008-03-03 22:00:00
2008-03-04 08:00:00
(9 rows)
```

## Subscript Generating Functions

- `generate_subscripts(array anyarray, dim int)`  
Description: Generates a series comprising the given array's subscripts.  
Return type: setof int
- `generate_subscripts(array anyarray, dim int, reverse boolean)`  
Description: Generates a series comprising the given array's subscripts. When **reverse** is true, the series is returned in reverse order.  
Return type: setof int

**generate\_subscripts** is a function that generates the set of valid subscripts for the specified dimension of the given array. Zero rows are returned for arrays that do

not have the requested dimension, or for NULL arrays (but valid subscripts are returned for NULL array elements). Example:

```
-- Basic usage
openGauss=# SELECT generate_subscripts('{NULL,1,NULL,2}'::int[], 1) AS s;
s

1
2
3
4
(4 rows)
-- Unnest a 2D array:
openGauss=# CREATE OR REPLACE FUNCTION unnest2(anyarray)
RETURNS SETOF anyelement AS $$
SELECT $1[i][j]
 FROM generate_subscripts($1,1) g1(i),
 generate_subscripts($1,2) g2(j);
$$ LANGUAGE sql IMMUTABLE;

openGauss=# SELECT * FROM unnest2(ARRAY[[1,2],[3,4]]);
unnest2

1
2
3
4
(4 rows)

-- Delete the function.
openGauss=# DROP FUNCTION unnest2;
```

## 7.5.21 Conditional Expression Functions

### Conditional Expression Functions

- `coalesce(expr1, expr2, ..., exprn)`

Description:

Returns the first of its parameters that are not null.

**COALESCE(expr1, expr2)** is equivalent to **CASE WHEN expr1 IS NOT NULL THEN expr1 ELSE expr2 END**.

Example:

```
openGauss=# SELECT coalesce(NULL,'hello');
coalesce

hello
(1 row)
```

Note:

- If all the expressions are equivalent to NULL in the expression list, this function returns **NULL**.
- This value is replaced by the default value when data is displayed.
- Like a **CASE** expression, **COALESCE** only evaluates the parameters that are needed to determine the result. That is, parameters to the right of the first not-**NULL** parameter are not evaluated.

- `decode(base_expr, compare1, value1, Compare2,value2, ... default)`

Description: Compares **base\_expr** with each **compare(n)** and **returns value(n)** if they are matched. If no match occurs, **default** is returned.

Example:

```
openGauss=# SELECT decode('A','A',1,'B',2,0);
case

1
(1 row)
```

- **nullif(expr1, expr2)**

Description: Returns **NULL** only when **expr1** is equal to **expr2**. Otherwise, **expr1** is returned.

**nullif(expr1, expr2)** is equivalent to **CASE WHEN expr1 = expr2 THEN NULL ELSE expr1 END**.

Example:

```
openGauss=# SELECT nullif('hello','world');
nullif

hello
(1 row)
```

Note:

Assume the two parameter data types are different:

- If implicit conversion exists between the two data types, implicitly convert the parameter of lower priority to this data type using the data type of higher priority. If the conversion succeeds, computation is performed.

Otherwise, an error is returned. Example:

```
openGauss=# SELECT nullif('1234'::VARCHAR,123::INT4);
nullif

1234
(1 row)
```

```
openGauss=# SELECT nullif('1234'::VARCHAR,'2012-12-24'::DATE);
ERROR: invalid input syntax for type timestamp: "1234"
```

- If implicit conversion is not applied between two data types, an error is returned. Example:

```
openGauss=# SELECT nullif(TRUE::BOOLEAN,'2012-12-24'::DATE);
ERROR: operator does not exist: boolean = timestamp without time zone
LINE 1: SELECT nullif(TRUE::BOOLEAN,'2012-12-24'::DATE) FROM sys_dummy;
 ^
HINT: No operator matches the given name and argument type(s). You might need to add explicit type casts.
```

- **nvl( expr1 , expr2 )**

Description:

- If **expr1** is **NULL**, **expr2** is returned.
- If **expr1** is not **NULL**, **expr1** is returned.

Example:

```
openGauss=# SELECT nvl('hello','world');
nvl

hello
(1 row)
```

Note: Parameters **expr1** and **expr2** can be of any data type. If **expr1** and **expr2** are of different data types, NVL checks whether **expr2** can be implicitly converted to **expr1**. If it can, the data type of **expr1** is returned. Otherwise, an error is returned.

- **greatest(expr1 [, ...])**

Description: Selects the largest value from a list of any number of expressions.

Return type:

Example:

```
openGauss=# SELECT greatest(1*2,2-3,4-1);
greatest

 3
(1 row)
openGauss=# SELECT greatest('HARRY', 'HARRIOT', 'HAROLD');
greatest

HARRY
(1 row)
```

- `least(expr1 [, ...])`

Description: Selects the smallest value from a list of any number of expressions.

Example:

```
openGauss=# SELECT least(1*2,2-3,4-1);
least

 -1
(1 row)
openGauss=# SELECT least('HARRY','HARRIOT','HAROLD');
least

HAROLD
(1 row)
```

- `EMPTY_BLOB()`

Description: Initiates a BLOB variable in an **INSERT** or an **UPDATE** statement to a **NULL** value.

Return type: BLOB

Example:

```
-- Create a table.
openGauss=# CREATE TABLE blob_tb(b blob,id int);
-- Insert data.
openGauss=# INSERT INTO blob_tb VALUES (empty_blob(),1);
-- Drop the table.
openGauss=# DROP TABLE blob_tb;
```

Note: The length is 0 obtained using **DBE\_LOB.GET\_LENGTH**.

## 7.5.22 System Information Functions

### Session Information Functions

- `current_catalog`  
Description: Name of the current database (called "catalog" in the SQL standard)

Return type: name

Example:

```
testdb=# SELECT current_catalog;
current_database

testdb
(1 row)
```

- `current_database()`  
Description: Name of the current database  
Return type: name

Example:

```
testdb=# SELECT current_database();
current_database

testdb
(1 row)
```

- `current_query()`

Description: Text of the currently executing query, as committed by the client (might contain more than one statement)

Return type: text

Example:

```
openGauss=# SELECT current_query();
current_query

SELECT current_query();
(1 row)
```

- `current_schema()`

Description: Name of current schema

Return type: name

Example:

```
openGauss=# SELECT current_schema();
current_schema

public
(1 row)
```

Remarks: **current\_schema** returns the first valid schema name in the search path. (If the search path is empty or contains no valid schema name, **NULL** is returned.) This is the schema that will be used for any tables or other named objects that are created without specifying a target schema.

- `current_schemas(Boolean)`

Description: Names of schemas in search path

Return type: name[]

Example:

```
openGauss=# SELECT current_schemas(true);
current_schemas

{pg_catalog,public}
(1 row)
```

Note:

**current\_schemas(Boolean)** returns an array of the names of all schemas presently in the search path. The Boolean option determines whether implicitly included system schemas such as **pg\_catalog** are included in the returned search path.

#### NOTE

The search path can be altered at run time by running the following command:

```
SET search_path TO schema [, schema, ...]
```

- `current_user`

Description: Username in the current execution environment

Return type: name

Example:



```
openGauss=# SELECT current_user;
current_user

omm
(1 row)
```

Note: **current\_user** is the user identifier that is applicable for permission checking. Normally it is equal to the session user, but it can be changed with **SET ROLE**. It also changes during the execution of functions with the attribute **SECURITY DEFINER**.

- **definer\_current\_user**

Description: Username in the current execution environment

Return type: name

Example:

```
openGauss=# SELECT definer_current_user();
definer_current_user

omm
(1 row)
```

- **pg\_current\_sessionid()**

Description: Session ID of the current execution context

Return type: text

Example:

```
openGauss=# SELECT pg_current_sessionid();
pg_current_sessionid

1579228402.140190434944768
(1 row)
```

Note: **pg\_current\_sessionid()** is used to obtain the session ID in the current execution context. The structure of the value is *Timestamp. Session ID*. When **enable\_thread\_pool** is set to **off**, the actual session ID is the thread ID.

- **pg\_current\_sessid**

Description: Session ID of the current execution context

Return type: text

Example:

```
openGauss=# select pg_current_sessid();
pg_current_sessid

140308875015936
(1 row)
```

Note: In thread pool mode, the ID of the current session is obtained. In non-thread pool mode, the backend thread ID of the current session is obtained.

- **pg\_current\_userid**

Description: Current user ID.

Return type: text

```
openGauss=# SELECT pg_current_userid();
pg_current_userid

10
(1 row)
```

- **working\_version\_num()**

Description: Returns a version number regarding system compatibility.

Return type: int

Example:

```
openGauss=# SELECT working_version_num();
working_version_num

 92231
(1 row)
```

- **tablespace\_oid\_name(oid)**

Description: Queries the tablespace name based on the tablespace OID.

Return type: text

Example:

```
openGauss=# select tablespace_oid_name(1663);
tablespace_oid_name

pg_default
(1 row)
```

- **inet\_client\_addr()**

Description: Remote connection address. **inet\_client\_addr** returns the IP address of the current client.

 **NOTE**

It is available only in remote connection mode.

Return type: inet

Example:

```
openGauss=# SELECT inet_client_addr();
inet_client_addr

10.10.0.50
(1 row)
```

- **inet\_client\_port()**

Description: Remote connection port. And **inet\_client\_port** returns the port number of the current client.

 **NOTE**

It is available only in remote connection mode.

Return type: int

Example:

```
openGauss=# SELECT inet_client_port();
inet_client_port

 33143
(1 row)
```

- **inet\_server\_addr()**

Description: Local connection address. **inet\_server\_addr** returns the IP address on which the server accepted the current connection.

 **NOTE**

It is available only in remote connection mode.

Return type: inet

Example:

```
openGauss=# SELECT inet_server_addr();
inet_server_addr
```

```

10.10.0.13
(1 row)
```

- `inet_server_port()`

Description: Local connection port. **inet\_server\_port** returns the port number. All these functions return **NULL** if the current connection is via a UDS.

 **NOTE**

It is available only in remote connection mode.

Return type: int

Example:

```
openGauss=# SELECT inet_server_port();
inet_server_port

8000
(1 row)
```

- `pg_backend_pid()`

Description: Process ID of the server process attached to the current session

Return type: int

Example:

```
openGauss=# SELECT pg_backend_pid();
pg_backend_pid

140229352617744
(1 row)
```

- `pg_conf_load_time()`

Description: Configures load time. **pg\_conf\_load\_time** returns the timestamp with time zone when the server configuration files were last loaded.

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT pg_conf_load_time();
pg_conf_load_time

2017-09-01 16:05:23.89868+08
(1 row)
```

- `pg_my_temp_schema()`

Description: OID of the temporary schema of a session. The value is **0** if the OID does not exist.

Return type: oid

Example:

```
openGauss=# SELECT pg_my_temp_schema();
pg_my_temp_schema

0
(1 row)
```

Note: **pg\_my\_temp\_schema** returns the OID of the current session's temporary schema, or zero if it has none (because it has not created any temporary tables). **pg\_is\_other\_temp\_schema** returns true if the given OID is the OID of another session's temporary schema.

- `pg_is_other_temp_schema(oid)`

Description: Specifies whether the schema is the temporary schema of another session.

Return type: Boolean

Example:

```
openGauss=# SELECT pg_is_other_temp_schema(25356);
pg_is_other_temp_schema

f
(1 row)
```

- `pg_listening_channels()`

Description: Channel names that the session is currently listening on

Return type: SETOF text

Example:

```
openGauss=# SELECT pg_listening_channels();
pg_listening_channels

(0 rows)
```

Note: **pg\_listening\_channels** returns a set of names of channels that the current session is listening to.

- `pg_postmaster_start_time()`

Description: Server start time **pg\_postmaster\_start\_time** returns the **timestamp with time zone** when the server started.

Return type: timestamp with time zone

Example:

```
openGauss=# SELECT pg_postmaster_start_time();
pg_postmaster_start_time

2017-08-30 16:02:54.99854+08
(1 row)
```

- `pg_get_ruledef(rule_oid)`

Description: Obtains the **CREATE RULE** command for a rule.

Return type: text

Example:

```
openGauss=# select * from pg_get_ruledef(24828);
pg_get_ruledef

CREATE RULE t1_ins AS ON INSERT TO t1 DO INSTEAD INSERT INTO t2 (id) VALUES (new.id);
(1 row)
```

- `sessionid2pid()`

Description: Obtains PID information from a session ID (for example, the **sessid** column in **gs\_session\_stat**).

Return type: int8

Example:

```
openGauss=# select sessionid2pid(sessid::cstring) from gs_session_stat limit 2;
sessionid2pid

139973107902208
139973107902208
(2 rows)
```

- `session_context( 'namespace' , 'parameter')`

Description: Obtains and returns the parameter values of a specified namespace.

Return type: VARCHAR

Example:

```
openGauss=# SELECT session_context('USERENV', 'CURRENT_SCHEMA');
session_context

public
(1 row)
```

Note: Currently, the **current\_user**, **current\_schema**, **client\_info**, **ip\_address**, **sessionid**, and **sid** parameters are supported.

- **pg\_trigger\_depth()**

Description: Current nesting level of triggers

Return type: int

Example:

```
openGauss=# SELECT pg_trigger_depth();
pg_trigger_depth

0
(1 row)
```

- **session\_user**

Description: Session username.

Return type: name

Example:

```
openGauss=# SELECT session_user;
session_user

omm
(1 row)
```

Note: **session\_user** is usually the user who initiated the current database connection, but administrators can change this setting with **SET SESSION AUTHORIZATION**.

- **user**

Description: Equivalent to **current\_user**.

Return type: name

Example:

```
openGauss=# SELECT user;
current_user

omm
(1 row)
```

- **getpgusername()**

Description: Obtains the database username.

Return type: name

Example:

```
openGauss=# select getpgusername();
getpgusername

GaussDB_userna
(1 row)
```

- **getdatabaseencoding()**

Description: Obtains the database encoding mode.

Return type: name

Example:

```
openGauss=# select getdatabaseencoding();
getdatabaseencoding

SQL_ASCII
(1 row)
```

- `version()`

Description: Version information. **version** returns a string describing a server's version.

Return type: text

Example:

```
openGauss=# select version();
version

(GaussDB Kernel VxxxRxxxCxx build fab4f5ea) compiled at 2021-10-24 11:58:22 commit 3086 last mr
6592 release
(1 row)
```

- `opengauss_version()`

Description: openGauss version information

Return type: text

The following is an example. Replace *x.x.x* in the query result with the actual value.

```
openGauss=# select opengauss_version();
opengauss_version

x.x.x
(1 row)
```

- `gs_deployment()`

Description: Information about the deployment mode of the current system

Return type: text

Example:

```
openGauss=# select gs_deployment();
gs_deployment

BusinessCentralized
(1 row)
```

- `get_hostname()`

Description: Returns the host name of the current node.

Return type: text

Example:

```
openGauss=# SELECT get_hostname();
get_hostname

linux-user
(1 row)
```

- `get_nodename()`

Description: Returns the name of the current node.

Return type: text

Example:

```
openGauss=# SELECT get_nodename();
get_nodename

```

```
datanode1
(1 row)
```

- `get_schema_oid(cstring)`

Description: Returns the OID of the queried schema.

Return type: oid

Example:

```
openGauss=# SELECT get_schema_oid('public');
get_schema_oid

2200
(1 row)
```

- `get_client_info()`

Description: Returns client information.

Return type: record

## Access privilege inquiry function

The DDL permissions, including ALTER, DROP, COMMENT, INDEX and VACUUM, are inherent permissions implicitly owned by the owner.

The following access permission query function only specifies whether a user has a certain permission on an object. That is, the permission on the object recorded in the **acl** column of the system catalog is returned.

- `has_any_column_privilege(user, table, privilege)`

Description: Queries whether a specified user has permission for any column of table.

**Table 7-35** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| table     | text, oid                  |
| privilege | text                       |

Return type: Boolean

- `has_any_column_privilege(table, privilege)`

Description: Queries whether the current user has permission to access any column of table. For details about the valid parameter types, see [Table 7-35](#).

Return type: Boolean

**has\_any\_column\_privilege** checks whether a user can access any column of a table in a particular way. Its parameter possibilities are analogous to **has\_table\_privilege**, except that the desired access permission must be some combination of SELECT, INSERT, UPDATE, COMMENT or REFERENCES.

 NOTE

Note that having any of these permissions at the table level implicitly grants it for each column of the table, so **has\_any\_column\_privilege** will always return **true** if **has\_table\_privilege** does for the same parameters. A success message is also returned if a column-level permission is granted for at least one column.

- **has\_column\_privilege**(user, table, column, privilege)

Description: Specifies whether a specified user has permission for columns.

**Table 7-36** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| table     | text, oid                  |
| column    | text, smallint             |
| privilege | text                       |

Return type: Boolean

- **has\_column\_privilege**(table, column, privilege)

Description: Specifies whether the current user has permission to access columns. For details about the valid parameter types, see [Table 7-36](#).

Return type: Boolean

**has\_column\_privilege** checks whether a user can access a column in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**, with the addition that the column can be specified either by name or attribute number. The desired access permission must be some combination of SELECT, INSERT, UPDATE, COMMENT or REFERENCES.

 NOTE

Note that having any of these permissions at the table level indicates that the permission is implicitly granted for each column of the table.

- **has\_database\_privilege**(user, database, privilege)

Description: Specifies whether a specified user has permission for accessing databases. The parameters are described as follows:

**Table 7-37** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| database  | text, oid                  |
| privilege | text                       |



Return type: Boolean

- `has_database_privilege(database, privilege)`

Description: Specifies whether the current user has permission to access a database. For details about the valid parameter types, see [Table 7-37](#).

Return type: Boolean

Note: **has\_database\_privilege** checks whether a user can access a database in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The desired access permission must be some combination of CREATE, CONNECT, TEMPORARY, ALTER, DROP, COMMENT or TEMP (which is equivalent to TEMPORARY).

- `has_directory_privilege(user, directory, privilege)`

Description: Specifies whether a specified user has permission for accessing directories.

**Table 7-38** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| directory | text, oid                  |
| privilege | text                       |

Return type: Boolean

- `has_directory_privilege(directory, privilege)`

Description: Specifies whether the current user has permission to access a directory. For details about the valid parameter types, see [Table 7-38](#).

Return type: Boolean

- `has_foreign_data_wrapper_privilege(user, fdw, privilege)`

Description: Specifies whether a specified user has permission for accessing foreign data wrappers.

**Table 7-39** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| fdw       | text, oid                  |
| privilege | text                       |

Return type: Boolean

- `has_foreign_data_wrapper_privilege(fdw, privilege)`

Description: Specifies whether the current user has permission for accessing foreign data wrappers. For details about the valid parameter types, see [Table 7-39](#).

Return type: Boolean

Note: **has\_foreign\_data\_wrapper\_privilege** checks whether a user can access a foreign data wrapper in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The desired access permission must be USAGE.

- **has\_function\_privilege**(user, function, privilege)

Description: Specifies whether a specified user has permission for accessing functions.

**Table 7-40** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| function  | text, oid                  |
| privilege | text                       |

Return type: Boolean

- **has\_function\_privilege**(function, privilege)

Description: Specifies whether the current user has permission for accessing functions. For details about the valid parameter types, see [Table 7-40](#).

Return type: Boolean

Note: **has\_function\_privilege** checks whether a user can access a function in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. When a function is specified by a text string rather than by OID, the allowed input is the same as that for the **regprocedure** data type (see [OID Types](#)). The access permission must be EXECUTE, ALTER, DROP, or COMMENT.

- **has\_language\_privilege**(user, language, privilege)

Description: Specifies whether a specified user has permission for accessing languages.

**Table 7-41** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| language  | text, oid                  |
| privilege | text                       |

Return type: Boolean

- `has_language_privilege(language, privilege)`  
Description: Specifies whether the current user has permission for accessing languages. For details about the valid parameter types, see [Table 7-41](#).  
Return type: Boolean  
Note: **has\_language\_privilege** checks whether a user can access a procedural language in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The desired access permission must be USAGE.
- `has_nodegroup_privilege(user, nodegroup, privilege)`  
Description: Checks whether a user has permission to access a database node.  
Return type: Boolean

**Table 7-42** Parameter type description

| Parameter | Valid Input Parameter Type |
|-----------|----------------------------|
| user      | name, oid                  |
| nodegroup | text, oid                  |
| privilege | text                       |

- `has_nodegroup_privilege(nodegroup, privilege)`  
Description: Checks whether a user has permission to access a database node. The parameter is similar to **has\_table\_privilege**. The access permission must be USAGE, CREATE, COMPUTE, ALTER, or DROP.  
Return type: Boolean
- `has_schema_privilege(user, schema, privilege)`  
Description: Specifies whether a specified user has permission for accessing schemas.  
Return type: Boolean
- `has_schema_privilege(schema, privilege)`  
Description: Specifies whether the current user has permission for accessing schemas.  
Return type: Boolean  
Note: **has\_schema\_privilege** checks whether a user can access a schema in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The desired access permission must be some combination of CREATE, USAGE, ALTER, DROP or COMMENT.
- `has_server_privilege(user, server, privilege)`  
Description: Specifies whether a specified user has permission for accessing foreign servers.  
Return type: Boolean
- `has_server_privilege(server, privilege)`  
Description: Specifies whether the current user has permission for accessing foreign servers.  
Return type: Boolean

Note: **has\_server\_privilege** checks whether a user can access a foreign server in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The access permission must be USAGE, ALTER, DROP, or COMMENT.

- **has\_table\_privilege(user, table, privilege)**

Description: Specifies whether a specified user has permission for accessing tables.

Return type: Boolean

- **has\_table\_privilege(table, privilege)**

Description: Specifies whether the current user has permission for accessing tables.

Return type: Boolean

Note: **has\_table\_privilege** checks whether a user can access a table in a particular way. The user can be specified by name or by OID (**pg\_authid.oid**), or be set to **public** which indicates the PUBLIC role. If this parameter is omitted, **current\_user** is used. The table can be specified by name or by OID. When it is specified by name, the name can be schema-qualified if necessary. The desired access permission is specified by a text string, which must be SELECT, INSERT, UPDATE, DELETE, TRUNCATE, REFERENCES, TRIGGER, ALTER, DROP, COMMENT, INDEX or VACUUM. Optionally, **WITH GRANT OPTION** can be added to a permission type to test whether the permission is held with grant option. Also, multiple permission types can be listed separated by commas, in which case the result will be **true** if any of the listed permissions is held.

Example:

```
openGauss=# SELECT has_table_privilege('tpcds.web_site', 'select');
has_table_privilege

t
(1 row)

openGauss=# SELECT has_table_privilege('omm', 'tpcds.web_site', 'select,INSERT WITH GRANT
OPTION ');
has_table_privilege

t
(1 row)
```

- **has\_tablespace\_privilege(user, tablespace, privilege)**

Description: Specifies whether a specified user has permission for tablespaces.

Return type: Boolean

- **has\_tablespace\_privilege(tablespace, privilege)**

Description: Specifies whether the current user has permission for tablespaces.

Return type: Boolean

Note: **has\_tablespace\_privilege** checks whether a user can access a tablespace in a particular way. Its argument possibilities are analogous to **has\_table\_privilege**. The access permission must be CREATE, ALTER, DROP, or COMMENT.

- **pg\_has\_role(user, role, privilege)**

Description: Specifies whether a specified user has permission for accessing roles.

Return type: Boolean

- `pg_has_role(role, privilege)`

Description: Specifies whether the current user has permission for accessing roles.

Return type: Boolean

Note: **pg\_has\_role** checks whether a user can access a role in a particular way. Its parameter possibilities are analogous to those of **has\_table\_privilege**, except that **public** cannot be used as a username. The desired access permission must be some combination of MEMBER and USAGE. MEMBER denotes direct or indirect membership in the role (that is, the SET ROLE permission), while USAGE denotes the permissions of the role are available without SET ROLE.

- `has_any_privilege(user, privilege)`

Description: Queries whether a specified user has certain ANY permission. If multiple permissions are queried at the same time, **true** is returned as long as one permission is obtained.

Return type: Boolean

**Table 7-43** Parameter type description

| Parameter | Valid Input Parameter Type | Description | Range             |
|-----------|----------------------------|-------------|-------------------|
| user      | name                       | User        | Existing username |

| Parameter | Valid Input Parameter Type | Description    | Range                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------|----------------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| privilege | text                       | ANY permission | Available values:<br>CREATE ANY TABLE [WITH ADMIN OPTION]<br>ALTER ANY TABLE [WITH ADMIN OPTION]<br>DROP ANY TABLE [WITH ADMIN OPTION]<br>SELECT ANY TABLE [WITH ADMIN OPTION]<br>INSERT ANY TABLE [WITH ADMIN OPTION]<br>UPDATE ANY TABLE [WITH ADMIN OPTION]<br>DELETE ANY TABLE [WITH ADMIN OPTION]<br>CREATE ANY SEQUENCE [WITH ADMIN OPTION]<br>CREATE ANY INDEX [WITH ADMIN OPTION]<br>CREATE ANY FUNCTION [WITH ADMIN OPTION]<br>EXECUTE ANY FUNCTION [WITH ADMIN OPTION]<br>CREATE ANY PACKAGE [WITH ADMIN OPTION]<br>EXECUTE ANY PACKAGE [WITH ADMIN OPTION]<br>CREATE ANY TYPE [WITH ADMIN OPTION] |

## Schema Visibility Inquiry Functions

Each function performs the visibility check for one type of database object. For functions and operators, an object in the search path is visible if there is no object of the same name and parameter data type earlier in the path. For operator classes, both name and associated index access methods are considered.

All these functions require OIDs to identify the objects to be checked. If you want to test an object by name, it is convenient to use the OID alias type (**regclass**, **regtype**, **regprocedure**, **regoperator**, **regconfig**, or **regdictionary**).

For example, a table is said to be visible if its containing schema is in the search path and no table of the same name appears earlier in the search path. This is

equivalent to the statement that the table can be referenced by name without explicit schema qualification. For example, to list the names of all visible tables:

```
openGauss=# SELECT relname FROM pg_class WHERE pg_table_is_visible(oid);
```

- `pg_collation_is_visible(collation_oid)`  
Description: Specifies whether the collation is visible in search path.  
Return type: Boolean
- `pg_conversion_is_visible(conversion_oid)`  
Description: Specifies whether the conversion is visible in search path.  
Return type: Boolean
- `pg_function_is_visible(function_oid)`  
Description: Specifies whether the function is visible in search path.  
Return type: Boolean
- `pg_opclass_is_visible(opclass_oid)`  
Description: Specifies whether the operator class is visible in search path.  
Return type: Boolean
- `pg_operator_is_visible(operator_oid)`  
Description: Specifies whether the operator is visible in search path.  
Return type: Boolean
- `pg_opfamily_is_visible(opclass_oid)`  
Description: Specifies whether the operator family is visible in search path.  
Return type: Boolean
- `pg_table_is_visible(table_oid)`  
Description: Specifies whether the table is visible in search path.  
Return type: Boolean
- `pg_ts_config_is_visible(config_oid)`  
Description: Specifies whether the text search configuration is visible in search path.  
Return type: Boolean
- `pg_ts_dict_is_visible(dict_oid)`  
Description: Specifies whether the text search dictionary is visible in search path.  
Return type: Boolean
- `pg_ts_parser_is_visible(parser_oid)`  
Description: Specifies whether the text search parser is visible in search path.  
Return type: Boolean
- `pg_ts_template_is_visible(template_oid)`  
Description: Specifies whether the text search template is visible in search path.  
Return type: Boolean
- `pg_type_is_visible(type_oid)`  
Description: Specifies whether the type (or domain) is visible in search path.  
Return type: Boolean

## System Catalog Information Functions

- `format_type(type_oid, typemod)`

Description: Obtains the SQL name of a data type.

Return type: text

Note: **format\_type** returns the SQL name of a data type based on the OID of the data type and possible modifiers. If the specific modifier is unknown, pass **NULL** at the position of the modifier. Modifiers are generally meaningful only for data types with length restrictions. The SQL name returned by **format\_type** contains the length of the data type, which can be calculated by taking `sizeof(int32)` from actual storage length [actual storage len - `sizeof(int32)`] in the unit of bytes. 32-bit space is required to store the customized length set by users. So the actual storage length contains 4 bytes more than the customized length. In the following example, the SQL name returned from **format\_type** is `character varying(6)`, indicating the length of `varchar` type is 6 bytes. So the actual storage length of `varchar` type is 10 bytes.

```
openGauss=# SELECT format_type((SELECT oid FROM pg_type WHERE typname='varchar'), 10);
format_type

character varying(6)
(1 row)
```

- `getdistributekey(table_name)`

Description: Obtains a distribution key for a hash table. Distribution is not supported in a standalone system and the return value of this function is empty.

- `pg_check_authid(role_oid)`

Description: Checks whether a role name with a given OID exists.

Return type: Boolean

Example:

```
openGauss=# select pg_check_authid(1);
pg_check_authid

f
(1 row)
```

- `pg_describe_object(catalog_id, object_id, object_sub_id)`

Description: Obtains the description of a database object.

Return type: text

Note: **pg\_describe\_object** returns a description of a database object specified by catalog OID, object OID and a (possibly zero) sub-object ID. This is useful to determine the identity of an object as stored in the **pg\_depend** catalog.

- `pg_get_constraintdef(constraint_oid)`

Description: Obtains the definition of a constraint.

Return type: text

- `pg_get_constraintdef(constraint_oid, pretty_bool)`

Description: Obtains the definition of a constraint.

Return type: text

Note: **pg\_get\_constraintdef** and **pg\_get\_indexdef** respectively reconstruct the creating command for a constraint and an index.



- `pg_get_expr(pg_node_tree, relation_oid)`  
Description: Decompiles internal form of an expression, assuming that any Vars in it refer to the relationship indicated by the second parameter.  
Return type: text
- `pg_get_expr(pg_node_tree, relation_oid, pretty_bool)`  
Description: Decompiles internal form of an expression, assuming that any Vars in it refer to the relationship indicated by the second parameter.  
Return type: text  
Note: **pg\_get\_expr** decompiles the internal form of an individual expression, such as the default value for a column. It can be useful when examining the contents of system catalogs. If the expression might contain Vars, specify the OID of the relationship they refer to as the second parameter; if no Vars are expected, zero is sufficient.
- `pg_get_functiondef(func_oid)`  
Description: Obtains the definition of a function.  
Return type: text  
Example:  

```
openGauss=# select * from pg_get_functiondef(598);
headerlines | definition
-----+-----
4 | CREATE OR REPLACE FUNCTION pg_catalog.abbrev(inet)+
 | RETURNS text +
 | LANGUAGE internal +
 | IMMUTABLE STRICT NOT FENCED NOT SHIPPABLE +
 | AS $function$inet_abbrev$function$ +
 |
(1 row)
```
- `pg_get_function_arguments(func_oid)`  
Description: Obtains the parameter list of the function's definition (with default values).  
Return type: text  
Note: **pg\_get\_function\_arguments** returns the parameter list of a function, in the form it would need to appear in within **CREATE FUNCTION**.
- `pg_get_function_identity_arguments(func_oid)`  
Description: Obtains the parameter list to identify a function (without default values).  
Return type: text  
Note: **pg\_get\_function\_identity\_arguments** returns the parameter list necessary to identify a function, in the form it would need to appear in within **ALTER FUNCTION**. This form omits default values.
- `pg_get_function_result(func_oid)`  
Description: Obtains the **RETURNS** clause for a function.  
Return type: text  
Note: **pg\_get\_function\_result** returns the appropriate **RETURNS** clause for the function.
- `pg_get_indexdef(index_oid)`  
Description: Obtains the **CREATE INDEX** command for an index.  
Return type: text

## Example:

```
openGauss=# select * from pg_get_indexdef(16416);
 pg_get_indexdef

CREATE INDEX test3_b_idx ON test3 USING btree (b) TABLESPACE pg_default
(1 row)
```

- `pg_get_indexdef(index_oid, dump_schema_only)`

Description: Obtains the **CREATE INDEX** command for indexes in dump scenarios. For an interval partitioned table that contains a local index, if **dump\_schema\_only** is set to **true**, the returned index creation statement does not contain the local index information of the automatically created partition. If **dump\_schema\_only** is set to **false**, the returned index creation statement contains the local index information of the automatically created partition. For a non-interval partitioned table or an interval partitioned table that does not contain a local index, the value of **dump\_schema\_only** does not affect the returned result of the function.

Return type: text

## Example:

```
openGauss=# CREATE TABLE sales
openGauss=# (prod_id NUMBER(6),
openGauss=# cust_id NUMBER,
openGauss=# time_id DATE,
openGauss=# channel_id CHAR(1),
openGauss=# promo_id NUMBER(6),
openGauss=# quantity_sold NUMBER(3),
openGauss=# amount_sold NUMBER(10,2)
openGauss=#)
PARTITION BY RANGE(time_id) INTERVAL('1 day')
openGauss=# (
openGauss=# partition p1 VALUES LESS THAN ('2019-02-01 00:00:00'),
openGauss=# partition p2 VALUES LESS THAN ('2019-02-02 00:00:00')
openGauss=#);
CREATE TABLE
openGauss=# create index index_sales on sales(prod_id) local (PARTITION idx_p1 ,PARTITION idx_p2);
CREATE INDEX
openGauss=#-- If the data to be inserted does not match any partition, create a partition and insert
the data into the new partition.
openGauss=# INSERT INTO sales VALUES(1, 12, '2019-02-05 00:00:00', 'a', 1, 1, 1);
INSERT 0 1
openGauss=# select oid from pg_class where relname = 'index_sales';
 oid

24632
(1 row)
openGauss=# select * from pg_get_indexdef(24632, true);
 pg_get_indexdef

CREATE INDEX index_sales ON sales USING btree (prod_id) LOCAL(PARTITION idx_p1, PARTITION
idx_p2) TABLESPACE pg_default
(1 row)
openGauss=# select * from pg_get_indexdef(24632, false);
 pg_get_indexdef

CREATE INDEX index_sales ON sales USING btree (prod_id) LOCAL(PARTITION idx_p1, PARTITION
idx_p2, PARTITION sys_p1_prod_id_idx) TA
BLESPEACE pg_default
(1 row)
```

- `pg_get_indexdef(index_oid, column_no, pretty_bool)`

Description: Obtains the **CREATE INDEX** command for an index, or definition of just one index column when **column\_no** is not zero.

Example:

```
openGauss=# select * from pg_get_indexdef(16416, 0, false);
 pg_get_indexdef

CREATE INDEX test3_b_idx ON test3 USING btree (b) TABLESPACE pg_default
(1 row)
openGauss=# select * from pg_get_indexdef(16416, 1, false);
 pg_get_indexdef

b
(1 row)
```

Return type: text

Note: **pg\_get\_functiondef** returns a complete **CREATE OR REPLACE FUNCTION** statement for a function.

- **pg\_get\_keywords()**

Description: Obtains the list of SQL keywords and their categories.

Return type: SETOF record

Note: **pg\_get\_keywords** returns a set of records describing the SQL keywords recognized by the server. The **word** column contains the keyword. The **catcode** column contains a category code: **U** for unreserved, **C** for column name, **T** for type or function name, or **R** for reserved. The **catdesc** column contains a possibly-localized string describing the category.

- **pg\_get\_userbyid(role\_oid)**

Description: Obtains the role name with a given OID.

Return type: name

Note: **pg\_get\_userbyid** extracts a role's name given its OID.

- **pg\_check\_authid(role\_id)**

Description: Checks whether a user exists based on **role\_id**.

Return type: text

Example:

```
openGauss=# select pg_check_authid(20);
 pg_check_authid

f
(1 row)
```

- **pg\_get\_viewdef(view\_name)**

Description: Obtains the underlying **SELECT** command for a view.

Return type: text

- **pg\_get\_viewdef(view\_name, pretty\_bool)**

Description: Obtains the underlying **SELECT** command for a view, lines with columns are wrapped to 80 columns if **pretty\_bool** is set to **true**.

Return type: text

Note: **pg\_get\_viewdef** reconstructs the **SELECT** query that defines a view. Most of these functions come in two variants. When the function has the parameter **pretty\_bool** and the value is true, it can optionally "pretty-print" the result. The pretty-printed format is more readable. The other one is the default format which is more likely to be interpreted in the same way by future versions. Avoid using pretty-printed output for dump purposes. Passing

**false** to the pretty-print parameter generates the same result as a variant without this parameter.

- `pg_get_viewdef(view_oid)`

Description: Obtains the underlying **SELECT** command for a view.

Return type: text

- `pg_get_viewdef(view_oid, pretty_bool)`

Description: Obtains the underlying **SELECT** command for a view, lines with columns are wrapped to 80 columns if **pretty\_bool** is set to **true**.

Return type: text

- `pg_get_viewdef(view_oid, wrap_column_int)`

Description: Obtains the underlying **SELECT** command for a view, wrapping lines with columns as specified, printing is implied.

Return type: text

- `pg_get_tabledef(table_oid)`

Description: Obtains a table definition based on **table\_oid**.

Example:

```
openGauss=# select * from pg_get_tabledef(16384);
 pg_get_tabledef

SET search_path = public;
CREATE TABLE t1 (
 c1 bigint DEFAULT nextval('serial'::regclass)+
)
WITH (orientation=row)
TO GROUP group1;
(1 row)
```

Return type: text

- `pg_get_tabledef(table_name)`

Description: Obtains a table definition based on **table\_name**.

Example:

```
openGauss=# select * from pg_get_tabledef('t1');
 pg_get_tabledef

SET search_path = public;
CREATE TABLE t1 (
 c1 bigint DEFAULT nextval('serial'::regclass)+
)
WITH (orientation=row)
TO GROUP group1;
(1 row)
```

Return type: text

Remarks: **pg\_get\_tabledef** reconstructs the **CREATE** statement of the table definition, including the table definition, index information, and comments. Users need to create the dependent objects of the table, such as groups, schemas, tablespaces, and servers. The table definition does not include the statements for creating these dependent objects.

- `pg_options_to_table(reloptions)`

Description: Obtains the set of storage option name/value pairs.

Return type: SETOF record

Note: **pg\_options\_to\_table** returns the set of storage option name/value pairs (**option\_name/option\_value**) when passed **pg\_class.reloptions** or **pg\_attribute.attoptions**.

- pg\_tablespace\_databases(tablespace\_oid)**

Description: Obtains the set of database OIDs that have objects in the specified tablespace.

Return type: SETOF oid

Note: **pg\_tablespace\_databases** allows a tablespace to be examined. It returns the set of OIDs of databases that have objects stored in the tablespace. If this function returns any rows, the tablespace is not empty and cannot be dropped. To display the specific objects populating the tablespace, you will need to connect to the databases identified by **pg\_tablespace\_databases** and query their **pg\_class** catalogs.
- pg\_tablespace\_location(tablespace\_oid)**

Description: Obtains the path in the file system that this tablespace is located in.

Return type: text
- pg\_typeof(any)**

Description: Obtains the data type of any value.

Return type: regtype

Note: **pg\_typeof** returns the OID of the data type of the value that is passed to it. This can be helpful for troubleshooting or dynamically constructing SQL queries. The function is declared as returning **regtype**, which is an OID alias type (see [OID Types](#)). This means that it is the same as an OID for comparison purposes but displays as a type name.

Example:

```
openGauss=# SELECT pg_typeof(33);
pg_typeof

integer
(1 row)

openGauss=# SELECT typen FROM pg_type WHERE oid = pg_typeof(33);
typen

4
(1 row)
```
- collation for (any)**

Description: Obtains the collation of the parameter.

Return type: text

Note: The expression **collation for** returns the collation of the value that is passed to it.

Example:

```
openGauss=# SELECT collation for (description) FROM pg_description LIMIT 1;
pg_collation_for

"default"
(1 row)
```

The value might be quoted and schema-qualified. If no collation is derived for the argument expression, then a null value is returned. If the parameter is not of a collectable data type, then an error is thrown.
- pg\_get\_serial\_sequence(tablename, colname)**

Description: Obtains the sequence of the corresponding table name and column name.

Return type: text

Example:

```
openGauss=# select * from pg_get_serial_sequence('t1', 'c1');
pg_get_serial_sequence

public.serial
(1 row)
```

- `pg_sequence_parameters(sequence_oid)`

Description: Obtains the parameters of a specified sequence, including the start value, minimum value, maximum value, and incremental value.

Return type: int16, int16, int16, int16, Boolean

Example:

```
openGauss=# select * from pg_sequence_parameters(16420);
start_value | minimum_value | maximum_value | increment | cycle_option
-----+-----+-----+-----+-----
101 | 1 | 9223372036854775807 | 1 | f
(1 row)
```

## Comment Information Functions

- `col_description(table_oid, column_number)`

Description: Obtains the comment for a table column.

Return type: text

Note: **col\_description** returns the comment for a table column, which is specified by the OID of its table and its column number.

- `obj_description(object_oid, catalog_name)`

Description: Obtains the comment for a database object.

Return type: text

Note: The two-parameter form of **obj\_description** returns the comment for a database object specified by its OID and the name of the containing system catalog. For example, **obj\_description(123456,'pg\_class')** would retrieve the comment for the table with OID 123456. The one-parameter form of **obj\_description** requires only the object OID.

**obj\_description** cannot be used for table columns since columns do not have OIDs of their own.

- `obj_description(object_oid)`

Description: Obtains the comment for a database object.

Return type: text

- `shobj_description(object_oid, catalog_name)`

Description: Obtains the comment for a shared database object.

Return type: text

Note: **shobj\_description** is used just like **obj\_description** except the former is used for retrieving comments on shared objects. Some system catalogs are global to all databases in GaussDB, and the comments for objects in them are stored globally as well.

## Transaction IDs and Snapshots

Internal transaction IDs (XIDs) are 64 bits. **txid\_snapshot**, the data type used by these functions, stores information about transaction ID visibility at a particular moment. [Table 7-44](#) describes its components.

**Table 7-44** Snapshot components

| Name     | Description                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| xmin     | Earliest transaction ID (txid) that is still active. All earlier transactions will either be committed and visible, or rolled back.                                                                                                                                                                                                                                                                                              |
| xmax     | First as-yet-unassigned txid. All txids greater than or equal to this are not yet started as of the time of the snapshot, so they are invisible.                                                                                                                                                                                                                                                                                 |
| xip_list | Active txids at the time of the snapshot. The list includes only those active txids between <b>xmin</b> and <b>xmax</b> ; there might be active txids higher than <b>xmax</b> . A txid that is <b>xmin</b> ≤ <b>txid</b> < <b>xmax</b> and not in this list was already completed at the time of the snapshot, and is either visible or dead according to its commit status. The list does not include txids of subtransactions. |

**txid\_snapshot**'s textual representation is **xmin:xmax:xip\_list**.

For example, **10:20:10,14,15** means **xmin=10**, **xmax=20**, **xip\_list=10, 14, 15**.

The following functions provide server transaction information in an exportable form. The main use of these functions is to determine which transactions were committed between two snapshots.

- **txid\_current()**  
Description: Obtains the current transaction ID.  
Return type: bigint
- **gs\_txid\_oldestxmin()**  
Description: Obtains the minimum transaction ID (specified by **oldesxmin**).  
Return type: bigint
- **txid\_current\_snapshot()**  
Description: Obtains the current snapshot.  
Return type: txid\_snapshot
- **txid\_snapshot\_xip(txid\_snapshot)**  
Description: Obtains in-progress transaction IDs in a snapshot.  
Return type: SETOF bigint
- **txid\_snapshot\_xmax(txid\_snapshot)**  
Description: Obtains **xmax** of snapshots.  
Return type: bigint
- **txid\_snapshot\_xmin(txid\_snapshot)**  
Description: Obtains **xmin** of snapshots.

- Return type: bigint
- `txid_visible_in_snapshot(bigint, txid_snapshot)`  
Description: Specifies whether the transaction ID is visible in a snapshot (do not use subtransaction IDs).  
Return type: Boolean
  - `get_local_prepared_xact()`  
Description: Obtains the two-phase residual transaction information of the current node, including the transaction ID, GID of the two-phase transaction, prepared time, owner OID, database OID, and node name of the current node.  
Return type: xid, text, timestamptz, oid, oid, text
  - `get_remote_prepared_xacts()`  
Description: Obtains the two-phase residual transaction information of all remote nodes, including the transaction ID, GID of the two-phase transaction, prepared time, owner name, database name, and node name.  
Return type: xid, text, timestamptz, name, name, text
  - `global_clean_prepared_xacts(text, text)`  
Description: Concurrently cleans two-phase residual transactions. Only the **gs\_clean** tool can call this function for the cleaning. In other situations, **false** is returned. This function is not supported in the centralized mode.  
Return type: Boolean
  - `gs_get_next_xid_csn()`  
Description: Returns the values of **next\_xid** and **next\_csn** on all nodes globally.  
The return values are as follows:

**Table 7-45** `gs_get_next_xid_csn` parameters

| Column   | Description                              |
|----------|------------------------------------------|
| nodename | Node name.                               |
| next_xid | Next transaction ID of the current node. |
| next_csn | Next CSN of the current node.            |

- `pg_control_system()`  
Description: Returns the status of the system control file.  
Return type: SETOF record
- `pg_control_checkpoint()`  
Description: Returns the system checkpoint status.  
Return type: SETOF record
- `pv_built_in_functions`  
Description: Displays information about all built-in system functions.  
Parameter: nan  
Return type: pronom name, pronomspace oid, proowner oid, proleng oid, procost real, prorows real, provariadic oid, protransform regproc, proisagg



Boolean, proiswindow Boolean, prosecdef Boolean, proleakproof Boolean, proisstrict Boolean, prorerset Boolean, provolatile "char", pronargs smallint, pronargdefaults smallint, prorettytype oid, proargtypes oidvector, proallargtypes integer[], proargmodes "char"[], proargnames text[], proargdefaults pg\_node\_tree, prosrc text, probin text, proconfig text[], proacl aclitem[], prodefaultargpos int2vector, fencedmode Boolean, proshippable Boolean, propackage Boolean, oid oid

- `pv_thread_memory_detail`

Description: Returns the memory information of each thread.

Parameter: nan

Return type: threadid text, tid bigint, thrdtype text, contextname text, level smallint, parent text, totalsize bigint, freesize bigint, usedsize bigint

- `pg_stat_file_recursive`

Description: Lists all files in a path.

Parameter: location text

- `pg_shared_memory_detail`

Description: Returns usage information about all generated shared memory contexts. For details about each column, see [GS\\_SHARED\\_MEMORY\\_DETAIL](#).

Parameter: nan

Return type: contextname text, level smallint, parent text, totalsize bigint, freesize bigint, usedsize bigint

- `get_gtm_lite_status`

**Description:** Returns the backup XID and CSN on the GTM for fault locating. This system function is not supported in GTM-FREE mode or centralized deployment.

- `gs_stat_get_wlm_plan_operator_info`

Description: Obtains operator plan information from the internal hash table.

Parameter: oid

Return type: datname text, queryid int8, plan\_node\_id int4, startup\_time int8, total\_time int8, actual\_rows int8, max\_peak\_memory int4, query\_dop int4, parent\_node\_id int4, left\_child\_id int4, right\_child\_id int4, operation text, orientation text, strategy text, options text, condition text, projection text

- `pg_stat_get_partition_tuples_hot_updated(oid)`

Description: Returns statistics on the number of hot-updated tuples in a partition with a specified partition ID.

Parameter: oid

Return type: bigint

- `gs_session_memory_detail_tp`

Description: Returns the memory usage of the session. For details, see [gs\\_session\\_memory\\_detail](#).

Parameter: nan

Return type: sessid text, sesstype text, contextname text, level smallint, parent text, totalsize bigint, freesize bigint, usedsize bigint

- `gs_thread_memory_detail`

Description: Returns the memory information of each thread.

Parameter: nan

Return type: threadid text, tid bigint, thrdtype text, contextname text, level smallint, parent text, totalsize bigint, freesize bigint, usedsize bigint

- pg\_stat\_get\_wlm\_realtime\_operator\_info

Description: Obtains the operator information of the real-time execution plan from the internal hash table.

Parameter: nan

Return type: queryid bigint, pid bigint, plan\_node\_id integer, plan\_node\_name text, start\_time timestamp with time zone, duration bigint, status text, query\_dop integer, estimated\_rows bigint, tuple\_processed bigint, min\_peak\_memory integer, max\_peak\_memory integer, average\_peak\_memory integer, memory\_skew\_percent integer, min\_spill\_size integer, max\_spill\_size integer, average\_spill\_size integer, spill\_skew\_percent integer, min\_cpu\_time bigint, max\_cpu\_time bigint, total\_cpu\_time bigint, cpu\_skew\_percent integer, warning text

- pg\_stat\_get\_wlm\_operator\_info

Description: Obtains the operator information of the execution plan from the internal hash table.

Parameter: nan

Return type: queryid bigint, pid bigint, plan\_node\_id integer, plan\_node\_name text, start\_time timestamp with time zone, duration bigint, query\_dop integer, estimated\_rows bigint, tuple\_processed bigint, min\_peak\_memory integer, max\_peak\_memory integer, average\_peak\_memory integer, memory\_skew\_percent integer, min\_spill\_size integer, max\_spill\_size integer, average\_spill\_size integer, spill\_skew\_percent integer, min\_cpu\_time bigint, max\_cpu\_time bigint, total\_cpu\_time bigint, cpu\_skew\_percent integer, warning text

- pg\_stat\_get\_wlm\_node\_resource\_info

Description: Obtains the resource information of the current node.

Parameter: nan

Return type: min\_mem\_util integer, max\_mem\_util integer, min\_cpu\_util integer, max\_cpu\_util integer, min\_io\_util integer, max\_io\_util integer, used\_mem\_rate integer

- pg\_stat\_get\_session\_wlmstat

Description: Returns the load information of the current session.

Parameter: pid integer

Return type: datid oid, threadid bigint, sessionid bigint, threadpid integer, usesysid oid, appname text, query text, priority bigint, block\_time bigint, elapsed\_time bigint, total\_cpu\_time bigint, skew\_percent integer, statement\_mem integer, active\_points integer, dop\_value integer, current\_cgroup text, current\_status text, enqueue\_state text, attribute text, is\_plana Boolean, node\_group text, srespool name

- pg\_stat\_get\_wlm\_instance\_info

Description: Returns the load information of the current instance.

Parameter: nan

Return type: instancename text, timestamp, timestamp with time zone, used\_cpu integer, free\_memory integer, used\_memory integer, io\_await double

precision, io\_util double precision, disk\_read double precision, disk\_write double precision, process\_read bigint, process\_write bigint, logical\_read bigint, logical\_write bigint, read\_counts bigint, write\_counts bigint

- `pg_stat_get_wlm_instance_info_with_cleanup`

Description: Returns the load information of the current instance and saves the information to the system catalog.

Parameter: nan

Return type: instancename text, timestamp, timestamp with time zone, used\_cpu integer, free\_memory integer, used\_memory integer, io\_await double precision, io\_util double precision, disk\_read double precision, disk\_write double precision, process\_read bigint, process\_write bigint, logical\_read bigint, logical\_write bigint, read\_counts bigint, write\_counts bigint

- `pg_stat_get_wlm_realtime_session_info`

Description: Returns the load information of the real-time session.

Parameter: nan

Return type: nodename text, threadid bigint, block\_time bigint, duration bigint, estimate\_total\_time bigint, estimate\_left\_time bigint, schemaname text, query\_band text, spill\_info text, control\_group text, estimate\_memory integer, min\_peak\_memory integer, max\_peak\_memory integer, average\_peak\_memory integer, memory\_skew\_percent integer, min\_spill\_size integer, max\_spill\_size integer, average\_spill\_size integer, spill\_skew\_percent integer, min\_dn\_time bigint, max\_dn\_time bigint, average\_dn\_time bigint, dntime\_skew\_percent integer, min\_cpu\_time bigint, max\_cpu\_time bigint, total\_cpu\_time bigint, cpu\_skew\_percent integer, min\_peak\_iops integer, max\_peak\_iops integer, average\_peak\_iops integer, iops\_skew\_percent integer, warning text, query text, query\_plan text, cpu\_top1\_node\_name text, cpu\_top2\_node\_name text, cpu\_top3\_node\_name text, cpu\_top4\_node\_name text, cpu\_top5\_node\_name text, mem\_top1\_node\_name text, mem\_top2\_node\_name text, mem\_top3\_node\_name text, mem\_top4\_node\_name text, mem\_top5\_node\_name text, cpu\_top1\_value bigint, cpu\_top2\_value bigint, cpu\_top3\_value bigint, cpu\_top4\_value bigint, cpu\_top5\_value bigint, mem\_top1\_value bigint, mem\_top2\_value bigint, mem\_top3\_value bigint, mem\_top4\_value bigint, mem\_top5\_value bigint, top\_mem\_dn text, top\_cpu\_dn text

- `pg_stat_get_wlm_session_iostat_info`

Description: Returns the session load I/O information.

Parameter: nan

Return type: threadid bigint, maxcurr\_iops integer, mincurr\_iops integer, maxpeak\_iops integer, minpeak\_iops integer, iops\_limits integer, io\_priority integer, curr\_io\_limits integer

- `pg_stat_get_wlm_statistics`

Description: Returns session load statistics.

Parameter: nan

Return type: statement text, block\_time bigint, elapsed\_time bigint, total\_cpu\_time bigint, qualification\_time bigint, skew\_percent integer, control\_group text, status text, action text

## 7.5.23 System Administration Functions

### 7.5.23.1 Configuration Settings Functions

Configuration setting functions are used for querying and modifying configuration parameters during running.

- `current_setting(setting_name)`

Description: Specifies the current setting.

Return type: text

Note: **current\_setting** obtains the current setting of **setting\_name** by query. It is equivalent to the **SHOW** statement.

Example:

```
openGauss=# SELECT current_setting('datestyle');
```

```
current_setting

ISO, MDY
(1 row)
```

- `set_working_grand_version_num_manually(tmp_version)`

Description: Upgrades new features of the database by switching the authorization version.

Return type: void

- `shell_in(type)`

Description: Inputs a route for the shell type that has not yet been filled.

Return type: void

- `shell_out(type)`

Description: Outputs a route for the shell type that has not yet been filled.

Return type: void

- `set_config(setting_name, new_value, is_local)`

Description: Sets the parameter and returns a new value.

Return type: text

Note: **set\_config** sets **setting\_name** to **new\_value**. If **is\_local** is set to **true**, **new\_value** applies only to the current transaction. If you want **new\_value** to apply for the current session, set the value to **false** instead. The function corresponds to the **SET** statement.

Example:

```
openGauss=# SELECT set_config('log_statement_stats', 'off', false);
```

```
set_config

off
(1 row)
```

### 7.5.23.2 Universal File Access Functions

Universal file access functions provide local access APIs for files on a database server. Only files in the database directory and the **log\_directory** directory can be accessed. Use a relative path for files in the database directory, and a path

matching the **log\_directory** configuration setting for log files. Only the database initial user can use these functions.

- **pg\_ls\_dir**(dirname text)

Description: Lists files in a directory.

Return type: setof text

Note: **pg\_ls\_dir** returns all the names in the specified directory, except the special entries "." and "..".

Example:

```
openGauss=# SELECT pg_ls_dir('./');
 pg_ls_dir

.postgresql.conf.swp
postgresql.conf
pg_tblspc
PG_VERSION
pg_ident.conf
core
server.crt
pg_serial
pg_twophase
postgresql.conf.lock
pg_stat_tmp
pg_notify
pg_subtrans
pg_ctl.lock
pg_xlog
pg_clog
base
pg_snapshots
postmaster.opts
postmaster.pid
server.key.rand
server.key.cipher
pg_multixact
pg_errorinfo
server.key
pg_hba.conf
pg_replslot
.pg_hba.conf.swp
cacert.pem
pg_hba.conf.lock
global
gaussdb.state
(32 rows)
```

- **pg\_read\_file**(filename text, offset bigint, length bigint)

Description: Returns the content of a text file.

Return type: text

Note: **pg\_read\_file** returns part of a text file. It can return a maximum of *length* bytes from *offset*. The actual size of fetched data is less than *length* if the end of the file is reached first. If *offset* is negative, it is the length rolled back from the file end. If *offset* and *length* are omitted, the entire file is returned.

Example:

```
openGauss=# SELECT pg_read_file('postmaster.pid',0,100);
 pg_read_file

53078 +
/srv/BigData/hadoop/data1/dbnode+
1500022474 +
```

```
8000 +
/var/run/FusionInsight +
localhost +
2
(1 row)
```

- `pg_read_binary_file(filename text [, offset bigint, length bigint,missing_ok boolean])`

Description: Returns the content of a binary file.

Return type: bytea

Note: **pg\_read\_binary\_file** is similar to **pg\_read\_file**, except that the result is a **bytea** value; accordingly, no encoding checks are performed. In combination with the **convert\_from** function, this function can be used to read a file in a specified encoding.

```
openGauss=# SELECT convert_from(pg_read_binary_file('filename'), 'UTF8');
```

- `pg_stat_file(filename text)`

Description: Returns status information about a file.

Return type: record

Note: **pg\_stat\_file** returns a record containing the file size, last access timestamp, last modification timestamp, last file status change timestamp, and a Boolean value indicating if it is a directory. Typical use cases are as follows:

```
openGauss=# SELECT * FROM pg_stat_file('filename');
openGauss=# SELECT (pg_stat_file('filename')).modification;
```

Example:

```
openGauss=# SELECT convert_from(pg_read_binary_file('postmaster.pid'), 'UTF8');
convert_from
```

```

4881 +
/srv/BigData/gaussdb/data1/dbnode+
1496308688 +
25108 +
/opt/user/Bigdata/gaussdb/gaussdb_tmp +
* +
25108001 43352069 +
(1 row)
```

```
openGauss=# SELECT * FROM pg_stat_file('postmaster.pid');
```

```
size | access | modification | change
| creation | isdir
-----+-----
117 | 2017-06-05 11:06:34+08 | 2017-06-01 17:18:08+08 | 2017-06-01 17:18:08+08
| f
(1 row)
openGauss=# SELECT (pg_stat_file('postmaster.pid')).modification;
modification

2017-06-01 17:18:08+08
(1 row)
```

### 7.5.23.3 Server Signal Functions

Server signal functions send control signals to other server processes. Only a system administrator has the permission to execute the following functions:

- `pg_cancel_backend(pid int)`

Description: Cancels a statement that is being executed by a backend thread.

Return type: Boolean

Note: **pg\_cancel\_backend** sends a query cancellation (SIGINT) signal to the backend process identified by **pid**. The PID of an active backend process can be found in the **pid** column of the **pg\_stat\_activity** view, or can be found by listing the database process using **ps** on the server. A user with the SYSADMIN permission, the owner of the database connected to the backend process, the owner of the backend process, or a user who inherits permissions of the built-in role **gs\_role\_signal\_backend** has the permission to use this function.

- **pg\_cancel\_session**(pid bigint, sessionid bigint)

Description: Cancels the statement that is being executed by an active session in thread pool mode. When the input parameters **pid** and **sessionid** are the same and both are thread IDs, the function of **pg\_cancel\_backend** is the same as that of **pg\_cancel\_backend**.

Return type: Boolean

#### NOTE

The input parameters of **pg\_cancel\_session** can be queried using the **pid** and **sessionid** columns in **pg\_stat\_activity**. The input parameters can be used to clear the statements that are being executed by active sessions in thread pool mode.

- **pg\_reload\_conf**()

Description: Causes all server processes to reload their configuration files.

Return type: Boolean

Note: **pg\_reload\_conf** sends a SIGHUP signal to the server. As a result, all server processes reload their configuration files.

- **pg\_rotate\_logfile**()

Description: Rotates the log files of the server.

Return type: Boolean

Note: **pg\_rotate\_logfile** sends a signal to the log file manager, instructing the manager to immediately switch to a new output file. This function works only when **redirect\_stderr** is used for log output. Otherwise, no log file manager subprocess exists.

- **pg\_terminate\_backend**(pid int)

Description: Terminates a backend thread.

Return type: Boolean

Note: Each of these functions returns **true** if they are successful and **false** otherwise. A user with the SYSADMIN permission, the owner of the database connected to the backend thread, the owner of the backend thread, or a user who inherits permissions of the built-in role **gs\_role\_signal\_backend** can use this function.

---

#### NOTICE

This function can terminate non-thread pool threads and active thread pool threads, but cannot terminate inactive thread pool threads.

---

Example:

```
openGauss=# SELECT pid from pg_stat_activity;
 pid

140657876268816
(1 rows)

openGauss=# SELECT pg_terminate_backend(140657876268816);
pg_terminate_backend

t
(1 row)
```

- `pg_terminate_session(pid int64, sessionid int64)`

Description: Terminates a backend session in thread pool mode.

Return type: Boolean

Note: Each of these functions returns **true** if they are successful and **false** otherwise. A user with the SYSADMIN permission, the owner of the database connected to the session, the owner of the session, or a user who inherits permissions of the built-in role `gs_role_signal_backend` has the permission to use this function.

---

#### NOTICE

When the input parameters **pid** and **sessionid** are the same and both are thread IDs, this function can terminate non-thread pool threads and active thread pool threads.

When the input parameters **pid** and **sessionid** are different, this function can terminate active sessions or close inactive sessions and the socket connection of the client.

- 
- `pg_terminate_active_session_socket(pid int64, sessionid int64)`

Description: Closes the socket connection between an active session and the client.

Return type: Boolean

Note: Each of these functions returns **true** if they are successful and **false** otherwise. This function can be used only by the initial user.

## 7.5.23.4 Backup and Restoration Control Functions

### Backup Control Functions

Backup control functions help with online backup.

- `pg_create_restore_point(name text)`

Description: Creates a named point for performing a restore (requires an administrator role).

Return type: text

Note: **pg\_create\_restore\_point** creates a named transaction log record that can be used as a restoration target, and returns the corresponding transaction log location. The given name can then be used with **recovery\_target\_name** to specify the point up to which restoration will proceed. Avoid creating multiple restoration points with the same name, since restoration will stop at the first one whose name matches the restoration target.



- **pg\_current\_xlog\_location()**  
Description: Obtains the write position of the current transaction log.  
Return type: text  
Note: **pg\_current\_xlog\_location** displays the write position of the current transaction log in the same format as those of the previous functions. Read-only operations do not require SYSADMIN permissions.
- **pg\_current\_xlog\_insert\_location()**  
Description: Obtains the insert position of the current transaction log.  
Return type: text  
Note: **pg\_current\_xlog\_insert\_location** displays the insert position of the current transaction log. The insertion point is the logical end of the transaction log at any instant, while the write location is the end of what has been written out from the server's internal buffers. The write position is the end that can be detected externally from the server. This operation can be performed to archive only some of completed transaction log files. The insert position is mainly used for commissioning the server. Read-only operations do not require SYSADMIN permissions.
- **gs\_current\_xlog\_insert\_end\_location()**  
Description: Obtains the insert position of the current transaction log.  
Return type: text  
Note: **gs\_current\_xlog\_insert\_end\_location** displays the insert position of the current transaction log.
- **pg\_start\_backup(label text [, fast boolean ])**  
Description: Starts to perform online backup. **operation\_mode** can be enabled only by SYSADMIN, replication role, or OPRADMIN.  
Return type: text  
Note: **pg\_start\_backup** receives a user-defined backup label (usually the name of the position where the backup dump file is stored). This function writes a backup label file to the data directory of the database and then returns the start position of backed up transaction logs in text mode.  

```
openGauss=# SELECT pg_start_backup('label_goes_here');
pg_start_backup

0/30000020
(1 row)
```
- **pg\_stop\_backup()**  
Description: Completes online backup. **operation\_mode** can be enabled only by SYSADMIN, replication role, or OPRADMIN.  
Return type: text  
Note: **pg\_stop\_backup** deletes the label file created by **pg\_start\_backup** and creates a backup history file in the transaction log archive area. The history file includes the label given to **pg\_start\_backup**, the start and end transaction log locations for the backup, and the start and end time of the backup. The return value is the backup's ending transaction log location. After the end position is calculated, the insert position of the current transaction log automatically goes ahead to the next transaction log file. In this way, the ended transaction log file can be immediately archived so that backup is complete.

- **pg\_switch\_xlog()**  
Description: Switches to a new transaction log file. **operation\_mode** can be enabled only by SYSADMIN or OPRADMIN.  
Return type: text  
Note: **pg\_switch\_xlog** moves to the next transaction log file so that the current log file can be archived (if continuous archive is used). The return value is the ending transaction log location + 1 within the just-completed transaction log file. If there has been no transaction log activity since the last transaction log switchover, pg\_switch\_xlog does not move but returns the start location of the transaction log file currently in use.
- **pg\_xlogfile\_name(location text)**  
Description: Converts the position string in a transaction log to a file name.  
Return type: text  
Note: **pg\_xlogfile\_name** extracts only the transaction log file name. If the given transaction log position is the transaction log file border, a transaction log file name will be returned for both the two functions. This is usually the desired behavior for managing transaction log archiving, since the preceding file is the last one that currently needs to be archived.
- **pg\_xlogfile\_name\_offset(location text)**  
Description: Converts the position string in a transaction log to a file name and returns the byte offset in the file.  
Return type: text, integer  
Note: **pg\_xlogfile\_name\_offset** can extract transaction log file names and byte offsets from the returned results of the preceding functions. Example:  

```
openGauss=# SELECT * FROM pg_xlogfile_name_offset(pg_stop_backup());
NOTICE: pg_stop_backup cleanup done, waiting for required WAL segments to be archived
NOTICE: pg_stop_backup complete, all required WAL segments have been archived
 file_name | file_offset
-----+-----
000000010000000000000003 | 272
(1 row)
```
- **pg\_xlog\_location\_diff(location text, location text)**  
Description: Calculates the difference in bytes between two transaction log locations.  
Return type: numeric
- **pg\_cbm\_tracked\_location()**  
Description: Queries the LSN location parsed by CBM.  
Return type: text
- **pg\_cbm\_get\_merged\_file(startLSNArg text, endLSNArg text)**  
Description: Combines CBM files within the specified LSN range into one and returns the name of the combined file.  
Return type: text  
Note: Only SYSADMIN or OPRADMIN can obtain the CBM combination file.
- **pg\_cbm\_get\_changed\_block(startLSNArg text, endLSNArg text)**  
Description: Combines CBM files within the specified LSN range into a table and return records of this table.  
Return type: records

Note: The table records returned by **pg\_cbm\_get\_changed\_block** include the start LSN, end LSN, tablespace OID, database OID, table relfilenode, table fork number, whether the table is deleted, whether the table is created, whether the table is truncated, number of pages in the truncated table, number of modified pages, and list of modified page numbers.

- **pg\_cbm\_recycle\_file**(targetLSNArg text)  
Description: Deletes the CBM files that are no longer used and returns the first LSN after the deletion.  
Return type: text
- **pg\_cbm\_force\_track**(targetLSNArg text,timeOut int)  
Description: Forcibly executes the CBM trace to the specified Xlog position and returns the Xlog position of the actual trace end point.  
Return type: text
- **pg\_enable\_delay\_ddl\_recycle**()  
Description: Enables DDL delay and returns the Xlog position of the enabling point. You need to enable **operation\_mode** as an administrator or O&M administrator.  
Return type: text
- **pg\_disable\_delay\_ddl\_recycle**(barrierLSNArg text, isForce bool)  
Description: Disables DDL delay and returns the Xlog range where DDL delay takes effect. You need to enable **operation\_mode** as an administrator or O&M administrator.  
Return type: records
- **pg\_enable\_delay\_xlog\_recycle**()  
Description: Enables Xlog recycle delay. This function is used in primary database node restoration. **operation\_mode** can be enabled only by an administrator or OPRADMIN.  
Return type: void
- **pg\_disable\_delay\_xlog\_recycle**()  
Description: Disables Xlog recycle delay. This function is used in primary database node restoration. **operation\_mode** can be enabled only by an administrator or OPRADMIN.  
Return type: void
- **pg\_cbm\_rotate\_file**(rotate\_lsn text)  
Description: Forcibly switches the file after the CBM parses **rotate\_lsn**. This function is called during the build process.  
Return type: void
- **gs\_roach\_stop\_backup**(backupid text)  
Description: Stops a backup started by the internal backup tool GaussRoach. It is similar to the **pg\_stop\_backup system** function but is more lightweight.  
Return type: text. The content is the insertion position of the current log.
- **gs\_roach\_enable\_delay\_ddl\_recycle**(backupid name)  
Description: Enables DDL delay and returns the log location of the enabling point. It is similar to the **pg\_enable\_delay\_ddl\_recycle** system function but is more lightweight. In addition, different **backupid** values can be used to concurrently enable DDL statements with delay.

Return type: text. The content is the log location of the start point.

- `gs_roach_disable_delay_ddl_recycle(backupid text)`

Description: Disables DDL delay and returns the log range where DDL delay takes effect. It is similar to the **pg\_enable\_delay\_ddl\_recycle** system function but is more lightweight. In addition, different **backupid** values can be used to concurrently disable DDL statements with delay.

Return type: records. The content is the range of logs for which DDL is delayed to take effect.

- `gs_roach_switch_xlog(request_ckpt bool)`

Description: Switches the currently used log segment file and triggers a full checkpoint if **request\_ckpt** is set to **true**.

Return type: text. The content is the location of the segment log.

- `gs_block_dw_io(timeout int, identifier text)`

Description: Blocks dual-write page flushing.

Parameter description:

- `timeout`

Block duration.

Value range: [0,3600] (s). The value **0** indicates that the block duration is 0s.

- `identifier`

ID of the operation.

Value range: a string, supporting only uppercase letters, lowercase letters, digits, and underscores (\_).

Return type: Boolean

Note: To call this function, the user must have the SYSADMIN or OPRADMIN permission. An O&M administrator must enable **operation\_mode**.

- `gs_is_dw_io_blocked()`

Description: Checks whether disk flushing on the current dual-write page is blocked. If disk flushing is blocked, **true** is returned.

Return type: Boolean

Note: To call this function, the user must have the SYSADMIN or OPRADMIN permission. An O&M administrator must enable **operation\_mode**.

## Restoration Control Functions

Restoration control functions provide information about the status of standby nodes. These functions may be executed both during restoration and in normal running.

- `pg_is_in_recovery()`

Description: Returns **true** if restoration is still in progress.

Return type: Boolean

- `pg_last_xlog_receive_location()`

Description: Obtains the last transaction log location received and synchronized to disk by streaming replication. While streaming replication is in progress, this will increase monotonically. If restoration has been

completed, then this value will remain static at the value of the last WAL record received and synchronized to disk during restoration. If streaming replication is disabled or if it has not yet started, the function returns a null value.

Return type: text

- `pg_last_xlog_replay_location()`

Description: Obtains last transaction log location replayed during restoration. If restoration is still in progress, this will increase monotonically. If restoration has been completed, then this value will remain static at the value of the last WAL record received during that restoration. When the server has been started normally without restoration, the function returns a null value.

Return type: text

- `pg_last_xact_replay_timestamp()`

Description: Obtains the timestamp of last transaction replayed during restoration. This is the time to commit a transaction or abort a WAL record on the primary node. If no transactions have been replayed during restoration, this function will return a null value. If restoration is still in progress, this will increase monotonically. If restoration has been completed, then this value will remain static at the value of the last WAL record received during that restoration. If the server normally starts without manual intervention, this function will return a null value.

Return type: timestamp with time zone

Restoration control functions control restoration processes. These functions may be executed only during restoration.

- `pg_is_xlog_replay_paused()`

Description: Returns **true** if restoration is paused.

Return type: Boolean

- `pg_xlog_replay_pause()`

Description: Pauses restoration immediately.

Return type: void

- `pg_xlog_replay_resume()`

Description: Restarts restoration if it was paused.

Return type: void

- `gs_get_active_archiving_standby()`

Description: Queries information about archive standby nodes in the same shard. The standby node name, archive location, and number of archived logs are returned.

Return type: text, text, int

- `gs_pitr_get_warning_for_xlog_force_recycle()`

Description: Checks whether logs are recycled because a large number of logs are stacked in the archive slot after archiving is enabled.

Return type: Boolean

- `gs_pitr_clean_history_global_barriers(stop_barrier_timestamp cstring)`

Description: Clears all barrier records generated before the specified time. The earliest barrier record is returned. The input parameter is of the cstring type

and is a Linux timestamp. You need to perform this operation as an administrator or O&M administrator.

Return type: text

- `gs_pitr_archive_slot_force_advance(stop_barrier_timestamp cstring)`  
Description: Forcibly updates the archive slot and clears unnecessary barrier records. The new archive slot location is returned. The input parameter is of the `cstring` type and is a Linux timestamp. You need to perform this operation as an administrator or O&M administrator.

Return type: text

While restoration is paused, no further database changes are applied. In hot standby mode, all new queries will see the same consistent snapshot of the database, and no further query conflicts will be generated until restoration is resumed.

If streaming replication is disabled, the paused state may continue indefinitely without problem. While streaming replication is in progress, WAL records will continue to be received, which will eventually fill available disk space. This progress depends on the duration of the pause, the rate of WAL generation, and available disk space.

### 7.5.23.5 DR Control Functions for Dual Database Instances

DR control functions for dual-database instances can be used to create an archive slot. The archive slot specifies the OBS information for storing physical logs.

- `pg_create_physical_replication_slot_extern(slotname text, dummy_standby bool, extra_content text, need_recycle_xlog bool)`  
Description: Creates an OBS or a NAS archive slot. **slotname** indicates the slot name of the DR standby node. The primary and standby nodes must use the same slot name. **dummy\_standby** specifies whether one primary and multiple standby nodes are deployed. The value **false** indicates one primary and multiple standby nodes, and the value **true** indicates not. **extra\_content** contains some information about the archive slot. For an OBS archive slot, the format is **OBS;obs\_server\_ip;obs\_bucket\_name;obs\_ak;obs\_sk;archive\_path;is\_recovery;is\_vote\_replicate**, where **OBS** indicates the archive media of the archive slot, **obs\_server\_ip** indicates the IP address of OBS, **obs\_bucket\_name** indicates the bucket name, **obs\_ak** indicates the AK of OBS, **obs\_sk** indicates the SK of OBS, **archive\_path** indicates the archive path **i**, and **is\_recovery** indicates whether the slot is an archive slot or a recovery slot. The value **0** indicates that the slot is an archive slot and is used by the primary database instance. The value **1** indicates that the slot is a recovery slot and is used by the DR database instance. **is\_vote\_replicate** specifies whether the voting copy is archived first. The value **0** indicates that the synchronous standby node is archived first, and the value **1** indicates that the voting copy is archived first. This column is reserved in the current version and is not adapted yet. For a NAS archive slot, the format is **NAS;archive\_path;is\_recovery;is\_vote\_replicate**. Compared with the OBS archive slot, the NAS archive slot does not have the OBS configuration information, while the meanings of other fields are the same.  
If the media is not specified, the OBS archive slot is used by default. The **extra\_content** format is

**obs\_server\_ip;obs\_bucket\_name;obs\_ak;obs\_sk;archive\_path;is\_recovery;is\_vote\_replicate. need\_recycle\_xlog** specifies whether to recycle old archived logs when creating an archive slot. The value **true** indicates that old archived logs are recycled, and the value **false** indicates that old archived logs are not recycled.

Return type: records, including **slotname** and **xlog\_position** of the current standby for DR.

Note: Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit permissions of the built-in role **gs\_role\_replication**. Currently, multiple archive slots cannot be created.

Example:

Create an OBS archive slot.

```
openGauss=# select * from pg_create_physical_replication_slot_exten('uuid', false, 'OBS;obs.cn-north-7.ulanhqab.huawei.com;dyk;19D772JBCACXX3KWS51D;*****;openGauss_uuid/dn1;0;0', false);
 slotname | xlog_position
-----+-----
 uuid |
(1 row)
```

Create a NAS archive slot.

```
openGauss=# select * from pg_create_physical_replication_slot_exten('uuid', false, 'NAS;/data/nas/media/openGauss_uuid/dn1;0;0', false);
 slotname | xlog_position
-----+-----
 uuid |
```

- **gs\_set\_obs\_delete\_location(delete\_location text)**

Description: Sets the location where OBS archived logs can be deleted. The value of **delete\_location** is an LSN. The logs before this location have been replayed and flushed to disks in the DR database instance and can be deleted on OBS.

Return type: **xlog\_file\_name** text, indicating the name of the log file that can be deleted. The value of this parameter is returned regardless of whether OBS is deleted successfully.

```
openGauss=# select gs_set_obs_delete_location('0/54000000');
 gs_set_obs_delete_location

 00000001000000000000000054_00
(1 row)
```

- **gs\_set\_obs\_delete\_location\_with\_slotname(cstring, cstring )**

Description: Sets the location where OBS archived logs can be deleted in a DR relationship. The first parameter indicates the LSN. The logs before this location have been replayed and flushed to disks in the DR database instance and can be deleted on OBS. The second parameter indicates the name of the archive slot.

Return type: **xlog\_file\_name** text, indicating the name of the log file that can be deleted. The value of this parameter is returned regardless of whether OBS is deleted successfully.

- **gs\_hadr\_do\_switchover()**

Description: Truncates services during a planned switchover in the primary database instance in OBS-based remote DR solutions.

Return type: Boolean, indicating whether the service truncation is successful and whether the switchover process can be performed normally.

- **gs\_streaming\_dr\_in\_switchover()**  
Description: Truncates services during a planned switchover in the primary database instance in streaming replication-based remote DR solutions.  
Return type: Boolean, indicating whether the service truncation is successful and whether the switchover process can be performed normally.

### 7.5.23.6 DR Query Functions for Dual-Database Instances

- **gs\_get\_global\_barrier\_status()**  
Description: If two-city 3DC DR is enabled, logs of the primary database instance and standby database instance for DR are synchronized through OBS. The barrier log is flushed to the disk of the primary database instance. The progress of archiving logs of the primary database instance and the progress of replaying logs of the standby database instance for DR are determined by replaying the standby database instance for DR.  
**gs\_get\_global\_barrier\_status** is used to query the latest global barrier that has been archived in OBS for the primary database instance.  
Return type: text  
**global\_barrier\_id**: globally latest barrier ID.  
**global\_achive\_barrier\_id**: globally latest archived barrier ID.
- **gs\_get\_global\_barriers\_status()**  
Description: If two-city 3DC DR solutions based on OBS are enabled, logs of the primary database instance and multiple standby database instances for DR are synchronized through OBS. The barrier log is flushed to the disk of the primary database instance. The progress of archiving logs of the primary database instance and the progress of replaying logs of the standby database instances for DR are determined by replaying the standby database instances for DR. **gs\_get\_global\_barriers\_status** is used to query the latest global barrier that has been archived in OBS for the primary database instance.  
Return type: text  
**slot\_name**: name of the slot used for DR.  
**global\_barrier\_id**: globally latest barrier ID.  
**global\_achive\_barrier\_id**: globally latest archived barrier ID.
- **gs\_get\_local\_barrier\_status()**  
Description: If two-city 3DC DR is enabled, logs of the primary database instance and standby database instance for DR are synchronized through OBS. The barrier log is flushed to the disk of the primary database instance. The progress of archiving logs of the primary database instance and the progress of replaying logs of the standby database instance for DR are determined by replaying the standby database instance for DR. **gs\_get\_local\_barrier\_status** is used to query the current log replay status of each node of the standby database instance for DR.  
Return type: text  
**barrier\_id**: latest barrier ID of a node of the standby database instance for DR.  
**barrier\_lsn**: LSN of the latest barrier ID returned by a node in the standby database instance for DR.



**archive\_lsn:** location of archived logs obtained by a node in the standby database instance for DR. This parameter does not take effect currently.

**flush\_lsn:** location of logs that have been flushed to disks on a node in the standby database instance for DR.

- `gs_upload_obs_file('slot_name', 'src_file', 'dest_file')`

Description: Function used by the primary database instance to upload data to OBS if two-city 3DC DR is enabled.

Return type: void

**slot\_name:** name of the replication slot created for the primary database instance.

**src\_file:** location of the file to be uploaded in the data directory of the primary database instance.

**dest\_file:** location of the file to be uploaded to OBS.

- `gs_download_obs_file('slot_name', 'src_file', 'dest_file')`

Description: Function used by the standby database instance for DR to download data from OBS to the localhost if two-city 3DC DR is enabled.

Return type: void

**slot\_name:** name of the replication slot created for the database instance for DR.

**src\_file:** location of files to be downloaded from OBS.

**dest\_file:** location of the downloaded file in the data directory of the standby database instance for DR.

- `gs_get_obs_file_context('file_name', 'slot_name')`

Description: Queries file content on OBS if two-city 3DC DR is enabled.

Return type: text

**file\_name:** name of the file on OBS.

**slot\_name:** name of the replication slot created for the primary database instance or database instance for DR.

- `gs_set_obs_file_context('file_name', 'file_context', 'slot_name')`

Description: Creates a file on OBS and writes content into the file if two-city 3DC DR is enabled.

Return type: text

**file\_name:** name of the file on OBS.

**file\_context:** content to be written into the file.

**slot\_name:** name of the replication slot created for the primary database instance or database instance for DR.

- `gs_get_hadr_key_cn()`

Description: Creates a file on OBS and writes content into the file if two-city 3DC DR is enabled.

Return type: text

**file\_name:** name of the file on OBS.

**file\_context:** content to be written into the file.

**slot\_name:** name of the replication slot created for the primary database instance or database instance for DR.

- **gs\_hadr\_has\_barrier\_creator()**

Description: Checks whether the **barrier\_creator** thread exists on the current node if two-city 3DC DR is enabled. If the thread exists, **true** is returned (the SYSADMIN permission is required).

Return type: Boolean

Note: This function is used only when a planned switchover is performed on the standby database instance for DR.
- **gs\_hadr\_in\_recovery()**

Description: Checks whether the current node is in barrier-based log restoration if two-city 3DC DR is enabled. If the current node is in barrier-based log restoration, **true** is returned. The standby database instance for DR can be promoted to the production database instance through the switchover process only after the log restoration is complete (the SYSADMIN permission is required).

Return type: Boolean

Note: This function is used only when a planned switchover is performed on the standby database instance for DR.
- **gs\_streaming\_dr\_get\_switchover\_barrier()**

Description: Checks whether the DN instance of the standby database instance for DR has received the switchover barrier logs and replayed the logs in two-city 3DC DR solutions based on streaming replication. If it has, **true** is returned. In the standby database instance for DR, the procedure for promoting the standby database instance for DR to the production database instance in the switchover process can be started only after the switchover barrier logs of all DN instances are replayed (restricted to system administrators).

Return type: Boolean

Note: This function is used only when a planned switchover is performed on the standby database instance for DR in streaming DR solutions.
- **gs\_streaming\_dr\_service\_truncation\_check()**

Description: Checks whether the DN instance of the primary database instance has sent the switchover barrier logs in two-city 3DC DR solutions based on streaming replication. If it has, **true** is returned. The procedure for demoting the production database instance to the standby database instance for DR in the switchover process can be started only after the logs are sent (the SYSADMIN permission is required).

Return type: Boolean

Note: This function is used only when a planned switchover is performed on the standby database instance for DR.
- **gs\_hadr\_local\_rto\_and\_rpo\_stat()**

Description: Displays the log flow control information about the primary and standby database instances in the streaming DR scenario. (This view can be used only on the primary DN of the primary database instance. Statistics cannot be obtained from the standby DN or the standby database instance.)

Return type: record. The types and meanings of the fields are as follows:

| Parameter               | Type | Description                                                                                                                                         |
|-------------------------|------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| hadr_sender_node_name   | text | Node name, including the primary database instance and the main standby node of the standby database instance.                                      |
| hadr_receiver_node_name | text | Name of the main standby node of the standby database instance.                                                                                     |
| source_ip               | text | IP address of the primary DN of the primary database instance.                                                                                      |
| source_port             | int  | Communication port of the primary DN of the primary database instance.                                                                              |
| dest_ip                 | text | IP address of the main standby DN of the standby database instance.                                                                                 |
| dest_port               | int  | Communication port of the main standby DN of the standby database instance.                                                                         |
| current_rto             | int  | Flow control information, that is, log RTO time of the current primary and standby database instances (unit: second).                               |
| target_rto              | int  | Flow control information, that is, the RTO time between the target primary and standby database instances (unit: second).                           |
| current_rpo             | int  | Flow control information, that is, log RPO time of the current primary and standby database instances (unit: second).                               |
| target_rpo              | int  | Flow control information, that is, the RPO time between the target primary and standby database instances (unit: second).                           |
| rto_sleep_time          | int  | RTO flow control information, that is, the expected sleep time (unit: $\mu$ s) required by walsender on the host to reach the specified RTO.        |
| rpo_sleep_time          | int  | RPO flow control information, that is, the expected sleep time (unit: $\mu$ s) required for inserting Xlogs on the host to reach the specified RPO. |

- gs\_hadr\_remote\_rto\_and\_rpo\_stat()

Description: Displays the log flow control information of the primary and standby database instances of other non-local data shards in streaming DR mode. This function is not supported in centralized deployment scenarios.

### 7.5.23.7 Snapshot Synchronization Functions

Snapshot synchronization functions save the current snapshot and return its identifier.

- `pg_export_snapshot()`

Description: Saves the current snapshot and returns its identifier.

Return type: text

Note: **pg\_export\_snapshot** saves the current snapshot and returns a text string identifying the snapshot. This string must be passed to clients that want to import the snapshot. A snapshot can be imported when the **set transaction snapshot snapshot\_id;** command is executed. Doing so is possible only when the transaction is set to the **SERIALIZABLE** or **REPEATABLE READ** isolation level. GaussDB does not support these two isolation levels currently. The output of the function cannot be used as the input of **set transaction snapshot**.

- `pg_export_snapshot_and_csn()`

Description: Saves the current snapshot and returns its identifier. Compared with **pg\_export\_snapshot()**, **pg\_export\_snapshot()** returns a CSN, indicating the CSN of the current snapshot.

Return type: text

### 7.5.23.8 Database Object Functions

#### Database Object Size Functions

Database object size functions calculate the actual disk space used by database objects.

- `pg_column_size(any)`

Description: Specifies the number of bytes used to store a particular value (possibly compressed)

Return type: int

Note: **pg\_column\_size** displays the space for storing an independent data value.

```
openGauss=# SELECT pg_column_size(1);
 pg_column_size

 4
(1 row)
```

- `pg_database_size(oid)`

Description: Specifies the disk space used by the database with the specified OID.

Return type: bigint

- `pg_database_size(name)`

Description: Specifies the disk space used by the database with the specified name.

Return type: bigint

Note: **pg\_database\_size** receives the OID or name of a database and returns the disk space used by the corresponding object.

Example:

```
openGauss=# SELECT pg_database_size('testdb');
pg_database_size

51590112
(1 row)
```

- **pg\_relation\_size(oid)**  
Description: Abbreviation of **pg\_relation\_size** (... , 'main'), specifying the disk space used by the table with a specified OID or index.  
Return type: bigint
- **pg\_relation\_size(relation regclass, fork text)**  
Description: Specifies the disk space used by the specified bifurcating tree ('main', 'fsm', or 'vm') of a certain table or index.  
Return type: bigint
- **pg\_relation\_size(relation regclass)**  
Description: Is an abbreviation of **pg\_relation\_size**(..., 'main').  
Return type: bigint  
Note: **pg\_relation\_size** receives the OID or name of a table or an index, and returns the size.
- **pg\_partition\_size(oid,oid)**  
Description: Specifies the disk space used by the partition with a specified OID. The first **oid** is the OID of the table and the second **oid** is the OID of the partition.  
Return type: bigint
- **pg\_partition\_size(text, text)**  
Description: Specifies the disk space used by the partition with a specified name. The first **text** is the table name and the second **text** is the partition name.  
Return type: bigint
- **pg\_partition\_indexes\_size(oid,oid)**  
Description: Specifies the disk space used by the index of the partition with a specified OID. The first **oid** is the OID of the table and the second **oid** is the OID of the partition.  
Return type: bigint
- **pg\_partition\_indexes\_size(text,text)**  
Description: Specifies the disk space used by the index of the partition with a specified name. The first **text** is the table name and the second **text** is the partition name.  
Return type: bigint
- **pg\_indexes\_size(regclass)**  
Description: Specifies the total disk space used by the index appended to the specified table.  
Return type: bigint
- **pg\_size\_pretty(bigint)**  
Description: Converts a size in bytes expressed as a 64-bit integer into a human-readable format with size units.

Return type: text

- `pg_size_pretty(numeric)`

Description: Converts a size in bytes expressed as a numeric value into a human-readable format with size units.

Return type: text

Note: **pg\_size\_pretty** formats the results of other functions into a human-readable format. KB/MB/GB/TB can be used.

- `pg_table_size(regclass)`

Description: Specifies the disk space used by the specified table, excluding indexes (but including TOAST, free space mapping, and visibility mapping).

Return type: bigint

- `pg_tablespace_size(oid)`

Description: Specifies the disk space used by the tablespace with a specified OID.

Return type: bigint

- `pg_tablespace_size(name)`

Description: Specifies the disk space used by the tablespace with a specified name.

Return type: bigint

Note: `pg_tablespace_size` receives the OID or name of a database and returns the disk space used by the corresponding object.

- `pg_total_relation_size(oid)`

Description: Specifies the disk space used by the table with a specified OID, including the index data.

Return type: bigint

- `pg_total_relation_size(regclass)`

Description: Specifies the total disk space used by the specified table, including all indexes and TOAST data.

Return type: bigint

- `pg_total_relation_size(text)`

Description: Specifies the disk space used by the table with a specified name, including the index data. The table name can be schema-qualified.

Return type: bigint

Note: `pg_total_relation_size` receives the OID or name of a table, and returns the sizes of the data and related indexes in bytes.

- `datalength(any)`

Description: Specifies the number of bytes used by an expression of a specified data type (data management space, data compression, or data type conversion is not considered).

Return type: int

Note: **datalength** is used to calculate the space of an independent data value.

Example:

```
openGauss=# SELECT datalength(1);
datalength
```

-----

4

(1 row)

The following table lists the supported data types and calculation methods.

| Data Type                 |                                         |                  | Storage Space                                                                                                        |
|---------------------------|-----------------------------------------|------------------|----------------------------------------------------------------------------------------------------------------------|
| Numeri<br>c data<br>types | Integ<br>er<br>types                    | TINYINT          | 1                                                                                                                    |
|                           |                                         | SMALLINT         | 2                                                                                                                    |
|                           |                                         | INTEGER          | 4                                                                                                                    |
|                           |                                         | BINARY_INTEGER   | 4                                                                                                                    |
|                           |                                         | BIGINT           | 8                                                                                                                    |
|                           | Arbitr<br>ary<br>precis<br>ion<br>types | DECIMAL          | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                           |                                         | NUMERIC          | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                           |                                         | NUMBER           | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                           | Sequ<br>ence<br>integ<br>er             | SMALLSERIAL      | 2                                                                                                                    |
|                           |                                         | SERIAL           | 4                                                                                                                    |
|                           |                                         | BIGSERIAL        | 8                                                                                                                    |
|                           |                                         | LARGESERIAL      | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                           | Floati<br>ng<br>point<br>types          | FLOAT4           | 4                                                                                                                    |
|                           |                                         | DOUBLE PRECISION | 8                                                                                                                    |
|                           |                                         | FLOAT8           | 8                                                                                                                    |
|                           |                                         | BINARY_DOUBLE    | 8                                                                                                                    |

|                      |                 |                |                                                                                                                      |
|----------------------|-----------------|----------------|----------------------------------------------------------------------------------------------------------------------|
|                      |                 | FLOAT[(p)]     | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                      |                 | DEC[(p,s)]     | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
|                      |                 | INTEGER[(p,s)] | Every four decimal digits occupy two bytes. The digits before and after the decimal point are calculated separately. |
| Boolean data types   | Boolean type    | BOOLEAN        | 1                                                                                                                    |
| Character data types | Character types | CHAR           | n                                                                                                                    |
|                      |                 | CHAR(n)        | n                                                                                                                    |
|                      |                 | CHARACTER(n)   | n                                                                                                                    |
|                      |                 | NCHAR(n)       | n                                                                                                                    |
|                      |                 | VARCHAR(n)     | n                                                                                                                    |
|                      |                 | CHARACTER      | Actual number of bytes of a character                                                                                |
|                      |                 | VARYING(n)     | Actual number of bytes of a character                                                                                |
|                      |                 | VARCHAR2(n)    | Actual number of bytes of a character                                                                                |
|                      |                 | NVARCHAR(n)    | Actual number of bytes of a character                                                                                |
|                      |                 | NVARCHAR2(n)   | Actual number of bytes of a character                                                                                |
|                      |                 | TEXT           | Actual number of bytes of a character                                                                                |
|                      |                 | CLOB           | Actual number of bytes of a character                                                                                |
| Time data types      | Time types      | DATE           | 8                                                                                                                    |
|                      |                 | TIME           | 8                                                                                                                    |
|                      |                 | TIMEZ          | 12                                                                                                                   |



|  |  |                        |    |
|--|--|------------------------|----|
|  |  | TIMESTAMP              | 8  |
|  |  | TIMESTAMPZ             | 8  |
|  |  | SMALLDATETIME          | 8  |
|  |  | INTERVAL DAY TO SECOND | 16 |
|  |  | INTERVAL               | 16 |
|  |  | RELTIME                | 4  |
|  |  | ABSTIME                | 4  |
|  |  | TINTERVAL              | 12 |

## Database Object Position Functions

- `pg_relation_filenode(relation regclass)`  
Description: Specifies the ID of a filenode with the specified relationship.  
Return type: oid  
Description: `pg_relation_filenode` receives the OID or name of a table, index, or sequence, and returns the **filenode** number allocated to it. **filenode** is the basic component of the file name used by the relationship. For most tables, the result is the same as that of **pg\_class.relfilenode**. For a specified system directory, **relfilenode** is set to **0** and this function must be used to obtain the correct value. If a relationship that is not stored is transmitted, such as a view, this function returns **NULL**.
- `pg_relation_filepath(relation regclass)`  
Description: Specifies the name of a file path with the specified relationship.  
Return type: text  
Description: **pg\_relation\_filepath** is similar to **pg\_relation\_filenode**, except that **pg\_relation\_filepath** returns the whole file path name for the relationship (relative to the data directory **PGDATA** of the database).
- `pg_filenode_relation(tablespace oid, filenode oid)`  
Description: Obtains the table names corresponding to the tablespace and relfilenode.  
Return type: regclass
- `pg_partition_filenode(partition_oid)`  
Description: Obtains **filenode** corresponding to the OID lock of a specified partitioned table.  
Return type: oid
- `pg_partition_filepath(partition_oid)`  
Description: Specifies the file path name of a partition.  
Return type: text

## Recycle Bin Object Functions

- `gs_is_recycle_object(classid, objid, objname)`  
Description: Determines whether an object is in the recycle bin.  
Return type: Boolean

### 7.5.23.9 Advisory Lock Functions

Advisory lock functions manage advisory locks.

- `pg_advisory_lock(key bigint)`  
Description: Obtains an exclusive session-level advisory lock.  
Return type: void  
Note: **pg\_advisory\_lock** locks resources defined by an application. The resources can be identified using a 64-bit or two nonoverlapped 32-bit key values. If another session locks the resources, the function blocks the resources until they can be used. The lock is exclusive. Multiple locking requests are pushed into the stack. Therefore, if the same resource is locked three times, it must be unlocked three times so that it is released to another session.
- `pg_advisory_lock(key1 int, key2 int)`  
Description: Obtains an exclusive session-level advisory lock.  
Return type: void  
Note: Only users with the SYSADMIN permission can add session-level exclusive advisory locks to the key-value pair (65535, 65535).
- `pg_advisory_lock(lock_id int4, lock_id int4, database_name Name)`  
Description: Obtains the exclusive advisory lock of a specified database by inputting the lock ID and database name.  
Return type: void
- `pg_advisory_lock_shared(key bigint)`  
Description: Obtains a shared session-level advisory lock.  
Return type: void
- `pg_advisory_lock_shared(key1 int, key2 int)`  
Description: Obtains a shared session-level advisory lock.  
Return type: void  
Note: **pg\_advisory\_lock\_shared** works in the same way as **pg\_advisory\_lock**, except that **pg\_advisory\_lock\_shared** obtains an advisory lock shared with other sessions requesting the lock, while **pg\_advisory\_lock** obtains an exclusive advisory lock.
- `pg_advisory_unlock(key bigint)`  
Description: Releases an exclusive session-level advisory lock.  
Return type: Boolean
- `pg_advisory_unlock(key1 int, key2 int)`  
Description: Releases an exclusive session-level advisory lock.  
Return type: Boolean

Note: **pg\_advisory\_unlock** releases the obtained exclusive advisory lock. If the release is successful, the function returns **true**. If the lock was not held, it will return **false**. In addition, an SQL warning will be reported by the server.

- **pg\_advisory\_unlock(lock\_id int4, lock\_id int4, database\_name Name)**

Description: Releases an exclusive advisory lock of a specified database by inputting the lock ID and database name.

Return type: Boolean

Note: If the release is successful, **true** is returned. If no lock is held, **false** is returned.

- **pg\_advisory\_unlock\_shared(key bigint)**

Description: Releases a shared session-level advisory lock.

Return type: Boolean

- **pg\_advisory\_unlock\_shared(key1 int, key2 int)**

Description: Releases a shared session-level advisory lock.

Return type: Boolean

Note: **pg\_advisory\_unlock\_shared** works in the same way as **pg\_advisory\_unlock**, except it releases a shared session-level advisory lock.

- **pg\_advisory\_unlock\_all()**

Description: Releases all advisory locks owned by the current session.

Return type: void

Note: **pg\_advisory\_unlock\_all** releases all advisory locks owned by the current session. The function is implicitly called when the session ends even if the client is abnormally disconnected.

- **pg\_advisory\_xact\_lock(key bigint)**

Description: Obtains an exclusive transaction-level advisory lock.

Return type: void

- **pg\_advisory\_xact\_lock(key1 int, key2 int)**

Description: Obtains an exclusive transaction-level advisory lock.

Return type: void

Note: **pg\_advisory\_xact\_lock** works in the same way as **pg\_advisory\_lock**, except the lock is automatically released at the end of the current transaction and cannot be released explicitly. Only users with the SYSADMIN permission can add transaction-level exclusive advisory locks to the key-value pair (65535, 65535).

- **pg\_advisory\_xact\_lock\_shared(key bigint)**

Description: Obtains a shared transaction-level advisory lock.

Return type: void

- **pg\_advisory\_xact\_lock\_shared(key1 int, key2 int)**

Description: Obtains a shared transaction-level advisory lock.

Return type: void

Note: **pg\_advisory\_xact\_lock\_shared** works in the same way as **pg\_advisory\_lock\_shared**, except the lock is automatically released at the end of the current transaction and cannot be released explicitly.

- `pg_try_advisory_lock(key bigint)`  
Description: Obtains an exclusive session-level advisory lock if available.  
Return type: Boolean  
Note: **pg\_try\_advisory\_lock** is similar to **pg\_advisory\_lock**, except **pg\_try\_advisory\_lock** does not block the resource until the resource is released. It either immediately obtains the lock and returns **true** or returns **false**, which indicates the lock cannot be performed currently.
- `pg_try_advisory_lock(key1 int, key2 int)`  
Description: Obtains an exclusive session-level advisory lock if available.  
Return type: Boolean  
Note: Only users with the SYSADMIN permission can add session-level exclusive advisory locks to the key-value pair (65535, 65535).
- `pg_try_advisory_lock_shared(key bigint)`  
Description: Obtains a shared session-level advisory lock if available.  
Return type: Boolean
- `pg_try_advisory_lock_shared(key1 int, key2 int)`  
Description: Obtains a shared session-level advisory lock if available.  
Return type: Boolean  
Note: **pg\_try\_advisory\_lock\_shared** works in the same way as **pg\_try\_advisory\_lock**, except that **pg\_try\_advisory\_lock\_shared** attempts to obtain a shared lock instead of an exclusive lock.
- `pg_try_advisory_xact_lock(key bigint)`  
Description: Obtains an exclusive transaction-level advisory lock if available.  
Return type: Boolean
- `pg_try_advisory_xact_lock(key1 int, key2 int)`  
Description: Obtains an exclusive transaction-level advisory lock if available.  
Return type: Boolean  
Note: **pg\_try\_advisory\_xact\_lock** works in the same way as **pg\_try\_advisory\_lock**, except the lock, if acquired, is automatically released at the end of the current transaction and cannot be released explicitly. Only users with the SYSADMIN permission can add transaction-level exclusive advisory locks to the key-value pair (65535, 65535).
- `pg_try_advisory_xact_lock_shared(key bigint)`  
Description: Obtains a shared transaction-level advisory lock if available.  
Return type: Boolean
- `pg_try_advisory_xact_lock_shared(key1 int, key2 int)`  
Description: Obtains a shared transaction-level advisory lock if available.  
Return type: Boolean  
Note: **pg\_try\_advisory\_xact\_lock\_shared** works in the same way as **pg\_try\_advisory\_lock\_shared**, except the lock, if acquired, is automatically released at the end of the current transaction and cannot be released explicitly.
- `lock_cluster_ddl()`

Description: Attempts to obtain a session-level exclusive advisory lock for all active primary database nodes in the database.

Return type: Boolean

Note: Only users with the SYSADMIN permission can call this function.

- `unlock_cluster_ddl()`

Description: Attempts to add a session-level exclusive advisory lock on the primary database node.

Return type: Boolean

### 7.5.23.10 Logical Replication Functions

- `pg_create_logical_replication_slot('slot_name', 'plugin_name')`

Description: Creates a logical replication slot.

Parameter description:

- `slot_name`

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (`_`), question marks (`?`), hyphens (`-`), and periods (`.`). One or two periods cannot be used alone as the replication slot name.

- `plugin_name`

Indicates the name of the plug-in.

Value range: a string, supporting "**mppdb\_decoding**".

Return type: name, text

Note: The first return value is the slot name, and the second one is the start LSN for decoding in the logical replication slot. Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit the permissions of the built-in role `gs_role_replication`. Currently, this function can be called only on the primary node.

- `pg_create_physical_replication_slot('slot_name', 'isDummyStandby')`

Description: Creates a physical replication slot.

Parameter description:

- `slot_name`

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (`_`), question marks (`?`), hyphens (`-`), and periods (`.`). One or two periods cannot be used alone as the replication slot name.

- `isDummyStandby`

Not supported in the current version.

Return type: name, text

Note: Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit permissions of the built-in role `gs_role_replication`.

- `pg_drop_replication_slot('slot_name')`

Description: Deletes a streaming replication slot.

Parameter description:

- slot\_name

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.). One or two periods cannot be used alone as the replication slot name.

Return type: void

Note: Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit permissions of the built-in role `gs_role_replication`. Currently, this function can be called only on the primary node.

- `pg_logical_slot_peek_changes('slot_name', 'LSN', upto_nchanges, 'options_name', 'options_value')`

Description: Performs decoding but does not go to the next streaming replication slot. (The decoded result will be returned again during the next decoding.)

Parameter description:

- slot\_name

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.). One or two periods cannot be used alone as the replication slot name.

- LSN

Indicates a target LSN. Decoding is performed only when an LSN is less than or equal to this value.

Value range: a string, in the format of *xlogid/xrecoff*, for example, **'1/2AAFC60'**. (If this parameter is set to **NULL**, the target LSN indicating the end position of decoding is not specified.)

- upto\_nchanges

Indicates the number of decoded records (including the **begin** and **commit** timestamps). Assume that there are three transactions, which involve 3, 5, and 7 records, respectively. If **upto\_nchanges** is set to **4**, 8 records of the first two transactions will be decoded. Specifically, decoding is stopped when the number of decoded records exceeds the value of **upto\_nchanges** after decoding in the first two transactions is finished.

Value range: a non-negative integer

 **NOTE**

If any of the **LSN** and **upto\_nchanges** values are reached, decoding ends.

- **options**: Specifies optional parameters, consisting of **options\_name** and **options\_value**.

- include-xids

Specifies whether the decoded **data** column contains XID information.

Valid value: **0** and **1**. The default value is **1**.

- **0**: The decoded **data** column does not contain XID information.

- **1**: The decoded **data** column contains XID information.
- **skip-empty-xacts**  
Specifies whether to ignore empty transaction information during decoding.  
Valid value: **0** and **1**. The default value is **0**.
  - **0**: The empty transaction information is not ignored during decoding.
  - **1**: The empty transaction information is ignored during decoding.
- **include-timestamp**  
Specifies whether decoded information contains the **commit** timestamp.  
Valid value: **0** and **1**. The default value is **0**.
  - **0**: The decoded information does not contain the **commit** timestamp.
  - **1**: The decoded information contains the **commit** timestamp.
- **only-local**  
Specifies whether to decode only local logs.  
Value range: **0** and **1**. The default value is **1**.
  - **0**: Non-local logs and local logs are decoded.
  - **1**: Only local logs are decoded.
- **force-binary**  
Specifies whether to output the decoding result in binary format.  
Value range: **0** and **1**. The default value is **0**.
  - **0**: The decoding result is output in binary format.
- **white-table-list**  
Whitelist parameter, including the schema and table name to be decoded.  
Value range: a string that contains table names in the whitelist. Different tables are separated by commas (,). An asterisk (\*) is used to fuzzily match all tables. Schema names and table names are separated by periods (.). No space character is allowed. The following is an example:

```
select * from pg_logical_slot_peek_changes('slot1', NULL, 4096, 'white-table-list', 'public.t1,public.t2');
```
- **max-txn-in-memory**  
Memory control parameter. The unit is MB. If the memory occupied by a single transaction is greater than the value of this parameter, data is flushed to disks.  
Value range: an integer ranging from 0 to 100. The default value is **0**, indicating that memory control is disabled.

- **max-reorderbuffer-in-memory**

Memory control parameter. The unit is GB. If the total memory (including the cache) of transactions being concatenated in the sender thread is greater than the value of this parameter, the current decoding transaction is flushed to disks.

Value range: an integer ranging from 0 to 100. The default value is **0**, indicating that memory control is disabled.

Return type: text, xid, text

Note: The function returns the decoding result. Each decoding result contains three columns, corresponding to the above return types and indicating the LSN position, XID, and decoded content, respectively.

Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit the permissions of the built-in role `gs_role_replication`.

- **pg\_logical\_slot\_get\_changes**('slot\_name', 'LSN', upto\_nchanges, 'options\_name', 'options\_value')

Description: Performs decoding and goes to the next streaming replication slot.

Parameter: This function has the same parameters as **pg\_logical\_slot\_peek\_changes**. For details, see [pg\\_logical\\_slot\\_peek\\_ch...](#)

Note: Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit permissions of the built-in role `gs_role_replication`. Currently, this function can be called only on the primary node.

- **pg\_logical\_slot\_peek\_binary\_changes**('slot\_name', 'LSN', upto\_nchanges, 'options\_name', 'options\_value')

Description: Performs decoding in binary format and does not go to the next streaming replication slot. (The decoded data can be obtained again during the next decoding.)

Parameter description:

- **slot\_name**

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (`_`), question marks (`?`), hyphens (`-`), and periods (`.`). One or two periods cannot be used alone as the replication slot name.

- **LSN**

Indicates a target LSN. Decoding is performed only when an LSN is less than or equal to this value.

Value range: a string, in the format of *xlogid/xrecoff*, for example, **'1/2AAFC60'**. (If this parameter is set to **NULL**, the target LSN indicating the end position of decoding is not specified.)

- **upto\_nchanges**

Indicates the number of decoded records (including the **begin** and **commit** timestamps). Assume that there are three transactions, which involve 3, 5, and 7 records, respectively. If **upto\_nchanges** is set to **4**, 8 records of the first two transactions will be decoded. Specifically, decoding is stopped when the number of decoded records exceeds the



value of **upto\_nchanges** after decoding in the first two transactions is finished.

Value range: a non-negative integer

 **NOTE**

If any of the **LSN** and **upto\_nchanges** values are reached, decoding ends.

- **options**: Specifies optional parameters, consisting of **options\_name** and **options\_value**.

- **include-xids**

Specifies whether the decoded **data** column contains XID information.

Valid value: **0** and **1**. The default value is **1**.

- **0**: The decoded **data** column does not contain XID information.
- **1**: The decoded **data** column contains XID information.

- **skip-empty-xacts**

Specifies whether to ignore empty transaction information during decoding.

Valid value: **0** and **1**. The default value is **0**.

- **0**: The empty transaction information is not ignored during decoding.
- **1**: The empty transaction information is ignored during decoding.

- **include-timestamp**

Specifies whether decoded information contains the **commit** timestamp.

Valid value: **0** and **1**. The default value is **0**.

- **0**: The decoded information does not contain the **commit** timestamp.
- **1**: The decoded information contains the **commit** timestamp.

- **only-local**

Specifies whether to decode only local logs.

Value range: **0** and **1**. The default value is **1**.

- **0**: Non-local logs and local logs are decoded.
- **1**: Only local logs are decoded.

- **force-binary**

Specifies whether to output the decoding result in binary format.

Value range: **0** or **1**. The default value is **0**. The result is output in binary format.

- **white-table-list**

Whitelist parameter, including the schema and table name to be decoded.

Value range: a string that contains table names in the whitelist. Different tables are separated by commas (,). An asterisk (\*) is used to fuzzily match all tables. Schema names and table names are separated by periods (.). No space character is allowed. The following is an example:

```
select * from pg_logical_slot_peek_binary_changes('slot1', NULL, 4096, 'white-table-list', 'public.t1,public.t2');
```

Return type: text, xid, bytea

Note: The function returns the decoding result. Each decoding result contains three columns, corresponding to the above return types and indicating the LSN position, XID, and decoded content in binary format, respectively. Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit the permissions of the built-in role `gs_role_replication`.

- `pg_logical_slot_get_binary_changes('slot_name', 'LSN', upto_nchanges, 'options_name', 'options_value')`

Description: Performs decoding in binary format and does not go to the next streaming replication slot.

Parameter: This function has the same parameters as **`pg_logical_slot_peek_binary_changes`**. For details, see [•pg\\_logical\\_slot\\_peek\\_bi...](#)

Note: Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit permissions of the built-in role `gs_role_replication`.

- `pg_replication_slot_advance ('slot_name', 'LSN')`

Description: Directly goes to the streaming replication slot for a specified LSN, without outputting any decoded result.

Parameter description:

- `slot_name`

Indicates the name of the streaming replication slot.

Value range: a string, supporting only lowercase letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.). One or two periods cannot be used alone as the replication slot name.

- `LSN`

Indicates a target LSN. Next decoding will be performed only in transactions whose commit position is greater than this value. If an input LSN is smaller than the position recorded in the current streaming replication slot, an error is reported. If the input LSN is greater than the LSN of the current physical log, the latter LSN will be directly used for decoding.

Value range: a string, in the format of *xlogid/xrecoff*

Return type: name, text

Note: A return result contains the slot name and LSN that is actually used for decoding. Users who call this function must have the SYSADMIN permission or the REPLICATION permission, or inherit the permissions of the built-in role `gs_role_replication`. Currently, this function can be called only on the primary node.

- `pg_logical_get_area_changes('LSN_start', 'LSN_end', upto_nchanges, 'decoding_plugin', 'xlog_path', 'options_name', 'options_value')`

Description: Specifies an LSN range or an Xlog file for decoding when no DDL operation is performed.

The constraints are as follows:

- When the API is called, only when **wal\_level** is set to **logical**, the generated log files can be parsed. If the used Xlog file is not of the logical level, the decoded content does not have the corresponding value and type, and there is no other impact.
- The Xlog file can be parsed only by a copy of a fully homogeneous DN to ensure that the metadata corresponding to the data can be found and no DDL or VACUUM FULL operation is performed.
- You can find the Xlog to be parsed.
- Do not read too many Xlog files at a time. You are advised to read one Xlog file at a time. It is estimated that the memory occupied by one Xlog file is two to three times the size of the Xlog file.
- The Xlog file before scale-out cannot be decoded.

Parameter description:

- **LSN\_start**

Specifies the LSN at the start of decoding.

Value range: a string, in the format of *xlogid/xrecoff*, for example, **'1/2AAFC60'**. If this parameter is set to NULL, the start position of decoding logs is not limited.

- **LSN\_end**

Specifies the LSN at the end of decoding.

Value range: a string, in the format of *xlogid/xrecoff*, for example, **'1/2AAFC60'**. (If this parameter is set to **NULL**, the target LSN indicating the end position of decoding is not specified.)

- **upto\_nchanges**

Indicates the number of decoded records (including the **begin** and **commit** timestamps). Assume that there are three transactions, which involve 3, 5, and 7 records, respectively. If **upto\_nchanges** is set to **4**, 8 records of the first two transactions will be decoded. Specifically, decoding is stopped when the number of decoded records exceeds the value of **upto\_nchanges** after decoding in the first two transactions is finished.

Value range: a non-negative integer

#### NOTE

If any of the **LSN** and **upto\_nchanges** values are reached, decoding ends.

- **decoding\_plugin**

Decoding plug-in, which is a .so plug-in that specifies the output format of the decoded content.

Value range: **mppdb\_decoding** and **sql\_decoding**.

- **xlog\_path**

Decoding plug-in, which specifies the Xlog absolute path and file level of the decoding file.

Value range: **NULL** or a character string of the absolute path of the Xlog file.

- **options**: This parameter is optional and consists of multiple pairs of **options\_name** and **options\_value**. You can retain the default values. For details, see [pg\\_logical\\_slot\\_peek\\_changes](#).

Example:

```
openGauss=# SELECT pg_current_xlog_location();
pg_current_xlog_location

0/E62E238
(1 row)

openGauss=# create table t1 (a int primary key,b int,c int);
NOTICE: CREATE TABLE / PRIMARY KEY will create implicit index "t1_pkey" for table "t1"
CREATE TABLE
openGauss=# insert into t1 values(1,1,1);
INSERT 0 1
openGauss=# insert into t1 values(2,2,2);
INSERT 0 1

openGauss=# select data from pg_logical_get_area_changes('0/
E62E238',NULL,NULL,'sql_decoding',NULL);
 location | xid | data
-----+-----+-----
0/E62E8D0 | 27213 | COMMIT (at 2022-01-26 15:08:03.349057+08) 3020226
0/E6325F0 | 27214 | COMMIT (at 2022-01-26 15:08:07.309869+08) 3020234
.....
```

- **pg\_get\_replication\_slots()**

Description: Obtains the replication slot list.

Return type: text, text, text, oid, Boolean, xid, xid, text, Boolean, text

Example:

```
openGauss=# select * from pg_get_replication_slots();
 slot_name | plugin | slot_type | datoid | active | xmin | catalog_xmin | restart_lsn |
 dummy_standby | confirmed_flush
-----+-----+-----+-----+-----+-----+-----+-----+-----
dn_s1 | | physical | 0 | t | | | 0/23DB14E0 | f |
slot1 | mppdb_decoding | logical | 16304 | f | | | 60966 | 0/1AFA1BB0 | f |
0/23DA5700
(2 rows)
```

- **gs\_get\_parallel\_decode\_status()**

Description: Monitors the length of the read log queue and decoding result queue of each decoder thread to locate the concurrent decoding performance bottleneck.

Return type: text, int, text, text, text, int64, int64

Example:

```
openGauss=# select * from gs_get_parallel_decode_status();
 slot_name | parallel_decode_num | read_change_queue_length | decode_change_queue_length |
 reader_lsn | working_txn_cnt | working_txn_memory
-----+-----+-----+-----+-----+-----+-----
slot1 | 2 | queue0: 1005, queue1: 320 | queue0: 63, queue1: 748 | 0/1DCE2578
| 42 | 192927504
(1 row)
```

Note: In the return values, **slot\_name** indicates the replication slot name, **parallel\_decode\_num** indicates the number of parallel decoder threads in the replication slot, **read\_change\_queue\_length** indicates the current length of

the log queue read by each decoder thread, **decode\_change\_queue\_length** indicates the current length of the decoding result queue of each decoder thread, **reader\_lsn** indicates the log location read by the reader thread, **working\_txn\_cnt** indicates the number of transactions being concatenated in the current sender thread, and **working\_txn\_memory** indicates the total memory (in bytes) occupied by the concatenation transactions in the sender thread.

- **pg\_replication\_origin\_create** (node\_name)

Description: Creates a replication source with a given external name and returns the internal ID assigned to it.

Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name  
Specifies the name of the replication source to be created.  
Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

Return type: oid

- **pg\_replication\_origin\_drop** (node\_name)

Description: Deletes a previously created replication source, including any associated replay progress.

Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name  
Specifies the name of the replication source to be deleted.  
Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

- **pg\_replication\_origin\_oid** (node\_name)

Description: Searches for a replication source by name and returns the internal ID. If no such replication source is found, an error is thrown.

Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name  
Specifies the name of the replication source to be queried.  
Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

Return type: oid

- **pg\_replication\_origin\_session\_setup** (node\_name)

Description: Marks the current session for replaying from a given origin, allowing you to track replay progress. This parameter can be used only when no origin is selected. Run the **pg\_replication\_origin\_session\_reset** command to cancel the configuration.

Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name

Specifies the name of the replication source.

Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

- `pg_replication_origin_session_reset ()`  
Description: Cancels the **pg\_replication\_origin\_session\_setup()** effect.  
Note: The user who calls this function must have the SYSADMIN permission.
- `pg_replication_origin_session_is_setup ()`  
Description: Returns a true value if a replication source is selected in the current session.  
Note: The user who calls this function must have the SYSADMIN permission.  
Return type: Boolean
- `pg_replication_origin_session_progress (flush)`  
Description: Returns the replay position of the replication source selected in the current session.  
Note: The user who calls this function must have the SYSADMIN permission.  
Parameter description:
  - `flush`  
Determines whether the corresponding local transaction has been flushed to disk.  
Value range: BooleanReturn type: LSN
- `pg_replication_origin_xact_setup (origin_lsn, origin_timestamp)`  
Description: Marks the current transaction as recommitted at a given LSN and timestamp. This function can be called only when **pg\_replication\_origin\_session\_setup** is used to select a replication source.  
Note: The user who calls this function must have the SYSADMIN permission.  
Parameter description:
  - `origin_lsn`  
Position for replaying the replication source.  
Value range: LSN
  - `origin_timestamp`  
Time point when a transaction is committed  
Value range: timestamp with time zone
- `pg_replication_origin_xact_reset ()`  
Description: Cancels the **pg\_replication\_origin\_xact\_setup()** effect.  
Note: The user who calls this function must have the SYSADMIN permission.
- `pg_replication_origin_advance (node_name, lsn)`  
Description:  
Sets the replication progress of a given node to a given position. This is primarily used to set the initial position, or to set a new position after a configuration change or similar change.  
Note: Improper use of this function may cause inconsistent replication data.  
Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name

Specifies the name of an existing replication source.

Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

- lsn

Position for replaying the replication source.

Value range: LSN

- pg\_replication\_origin\_progress (node\_name, flush)

Description: Returns the position for replaying the given replication source.

Note: The user who calls this function must have the SYSADMIN permission.

Parameter description:

- node\_name

Specifies the name of the replication source.

Value range: a string, supporting only letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.).

- flush

Determines whether the corresponding local transaction has been flushed to disk.

Value range: Boolean

- pg\_show\_replication\_origin\_status()

Description: Displays the replication status of the replication source.

Note: The user who calls this function must have the SYSADMIN permission.

Return type:

- **local\_id**: OID, which specifies the ID of the replication source.
- **external\_id**: text, which specifies the name of the replication source.
- **remote\_lsn**: LSN of the replication source.
- **local\_lsn**: local LSN.

### 7.5.23.11 Segment-Page Storage Functions

- local\_segment\_space\_info(tablespace\_name TEXT, database\_name TEXT)

Description: Generates usage information about all extent groups in the tablespace.

Return type:

|              |                                                                             |
|--------------|-----------------------------------------------------------------------------|
| node_name    | Node name                                                                   |
| extent_size  | Extent specifications of an extent group. The unit is the number of blocks. |
| forknum      | Fork number                                                                 |
| total_blocks | Total number of extents in a physical file                                  |

|                  |                                                                                                                                                                                                |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| meta_data_blocks | Number of blocks occupied by the metadata managed in a tablespace, including the space header and map page but excluding the segment head                                                      |
| used_data_blocks | Number of extents used for storing data, including the segment head                                                                                                                            |
| utilization      | Percentage of the number of used blocks to the total number of blocks, that is, (the value of <b>used_data_blocks</b> + the value of <b>meta_data_block</b> )/the value of <b>total_blocks</b> |
| high_water_mark  | High-water mark, indicating the number of allocated extents and maximum physical page number. Blocks that exceed the high-water mark are not used and can be directly recycled.                |

Example:

```
select * from local_segment_space_info('pg_default', 'testdb');
 node_name | extent_size | forknum | total_blocks | meta_data_blocks | used_data_blocks |
utilization | high_water_mark
-----+-----+-----+-----+-----+-----+-----
dn_6001_6002_6003 | 1 | 0 | 16384 | 4157 | 1 | .253784 |
4158
dn_6001_6002_6003 | 8 | 0 | 16384 | 4157 | 8 | .254211 |
4165
(2 rows)
```

- pg\_stat\_segment\_extent\_usage(int4 tablespace oid, int4 database oid, int4 extent\_type, int4 forknum)

Description: Specifies the usage information of each allocated extent in an extent group returned each time. **extent\_type** indicates the type of the extent group. The value is an integer ranging from 1 to 5. If the value is not within the range, an error is reported. **forknum** indicates the fork number. The value is an integer ranging from 0 to 4. Currently, only the following values are valid: **0** for data files, **1** for FSM files, and **2** for visibility map files.

Return type:

| Name           | Description                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| start_block    | Start physical page number of an extent                                                                                                                      |
| extent_size    | Size of an extent                                                                                                                                            |
| usage_type     | Usage type of an extent, for example, <b>segment head</b> and <b>data extent</b>                                                                             |
| owner_location | Object location of an extent to which a pointer points. For example, the owner of a data extent is the head of the segment to which the data extent belongs. |



| Name         | Description                                                                                                                                                                                            |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| special_data | Position of an extent in its owner. The value of this field is related to the usage type. For example, special data of a data extent is the extent ID in the segment to which the data extent belongs. |

The value of **usage\_type** is enumerated. The meaning of each value is as follows:

- **Non-bucket table segment head**: data segment head of a non-hash bucket table
- **Non-bucket table fork head**: fork segment header of a non-segment-page table
- **Data extent**: data block

Example:

```
select * from pg_stat_segment_extent_usage((select oid::int4 from pg_tablespace where
spcname='pg_default'), (select oid::int4 from pg_database where datname='testdb'), 1, 0);
```

| start_block | extent_size | usage_type  | owner_location | special_data |
|-------------|-------------|-------------|----------------|--------------|
| 4157        | 1           | Data extent | 4294967295     | 0            |
| 4158        | 1           | Data extent | 4157           | 0            |

- **local\_space\_shrink**(tablespacename TEXT, databasename TEXT)

Description: Shrinks specified physical segment-page space on the current node. Only the currently connected database can be shrunk.

Return value: empty

- **gs\_space\_shrink**(int4 tablespace, int4 database, int4 extent\_type, int4 forknum)

Description: Works similar to **local\_space\_shrink**, that is, shrinks specified physical segment-page space. However, the parameters are different. The input parameters are the OIDs of the tablespace and database, and the value of **extent\_type** is an integer ranging from 2 to 5. Note: The value 1 of **extent\_type** indicates segment-page metadata. Currently, the physical file that contains the metadata cannot be shrunk. This function is used only by tools. You are advised not to use it directly.

Return value: empty

- **pg\_stat\_remain\_segment\_info**()

Description: Displays residual extents on the current node due to faults. Residual extents are classified into two types: segments that are allocated but not used and extents that are allocated but not used. The main difference is that a segment contains multiple extents. During reclamation, all extents in the segment need to be recycled.

Return type:

| Name     | Description   |
|----------|---------------|
| space_id | Tablespace ID |

|          |                                                                                                                     |
|----------|---------------------------------------------------------------------------------------------------------------------|
| db_id    | Database ID                                                                                                         |
| block_id | Extent ID                                                                                                           |
| type     | Extent type. The options are as follows:<br><b>ALLOC_SEGMENT</b> , <b>DROP_SEGMENT</b> , and <b>SHRINK_EXTENT</b> . |

The values of **type** are described as follows:

- **ALLOC\_SEGMENT**: When a user creates a segment-page table and the segment is just allocated but the transaction of creating a table is not committed, the node is faulty. As a result, the segment is not used after being allocated.
- **DROP\_SEGMENT**: When a user deletes a segment-page table and the transaction is successfully committed, the bit corresponding to the segment page of the table is not reset and a fault, such as power failure, occurs. As a result, the segment is not used or released.
- **SHRINK\_EXTENT**: When a user shrinks a segment-page table and does not release the idle extent, a fault, such as power failure, occurs. As a result, the extent remains and cannot be reused.

Example:

```
select * from pg_stat_remain_segment_info();
space_id | db_id | block_id | type
-----+-----+-----+-----
1663 | 16385 | 4156 | ALLOC_SEGMENT
```

- `pg_free_remain_segment(int4 spaceId, int4 dbId, int4 segmentId)`

Description: Releases a specified residual extent. The value must be obtained from the **pg\_stat\_remain\_segment\_info** function. The function verifies input values. If the specified extent is not among the recorded residual extents, an error message is returned. If the specified extent is a single extent, the extent is released independently. If it is a segment, the segment and all extents in the segment are released.

Return value: empty

### 7.5.23.12 Other Functions

- `plan_seed()`

Description: Obtains the seed value of the previous query statement (internal use).

Return type: int

- `pg_stat_get_env()`

Description: Obtains the environment variable information of the current node. Only users with the SYSADMIN or MONADMIN permission can access the environment variable information.

Return type: record

Example:

```
openGauss=# select pg_stat_get_env();
pg_stat_get_env
```

```

(sgnode,"localhost,XXX.XXX.XXX.XXX",28589,26000,/home/omm,/home/omm/data/
single_node,pg_log)
(1 row)
```

- `pg_catalog.plancache_clean()`  
Description: Clears the global plan cache that is not used on nodes.  
Return type: Boolean
- `pg_catalog.plancache_status()`  
Description: Displays information about the global plan cache on nodes. The information returned by the function is the same as that in [GLOBAL\\_PLANCACHE\\_STATUS](#).  
Return type: record
- `textlen(text)`  
Description: Provides the method of querying the logical length of text.  
Return type: int
- `threadpool_status()`  
Description: Displays the status of worker threads and sessions in the thread pool.  
Return type: record
- `get_local_active_session()`  
Description: Provides sampling records of the historical active sessions stored in the memory of the current node.  
Return type: record
- `pg_stat_get_thread()`  
Description: Provides status information about all threads on the current node. Users with the SYSADMIN or MONADMIN permission can view information about all threads, and common users can view only their own thread information.  
Return type: record
- `pg_stat_get_sql_count()`  
Description: Provides the counts of the SELECT, UPDATE, INSERT, DELETE, and MERGE INTO statements executed on the current node. Users with the SYSADMIN or MONADMIN permission can view information about all users, and common users can view only their own statistics.  
Return type: record
- `pg_stat_get_data_senders()`  
Description: Provides detailed information about the data-copy sender thread active at the moment.  
Return type: record
- `get_wait_event_info()`  
Description: Provides detailed information about the wait event.  
Return type: record
- `generate_wdr_report(begin_snap_id bigint, end_snap_id bigint, report_type cstring, report_scope cstring, node_name cstring)`

Description: Generates system diagnosis reports based on two snapshots. You need to run the command in the system database. By default, the initial user or users with the MONADMIN permission can access the database. The result can be queried only in the system database but cannot be queried in the user database.

Return type: record

**Table 7-46** generate\_wdr\_report parameter description

| Parameter     | Description                                                                                                                                                                                                                                                                                                                | Range                                                                                                                                                                    |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| begin_snap_id | Snapshot ID that starts the diagnosis report period.                                                                                                                                                                                                                                                                       | N/A                                                                                                                                                                      |
| end_snap_id   | Snapshot ID that ends the diagnosis report period. By default, the value of <b>end_snap_id</b> is greater than that of <b>begin_snap_id</b> .                                                                                                                                                                              | N/A                                                                                                                                                                      |
| report_type   | Specifies the type of the generated report.                                                                                                                                                                                                                                                                                | <ul style="list-style-type: none"><li>• <b>summary</b></li><li>• <b>detail</b></li><li>• <b>all</b>: Both <b>summary</b> and <b>detail</b> are included.</li></ul>       |
| report_scope  | Specifies the scope for a report to be generated.                                                                                                                                                                                                                                                                          | <ul style="list-style-type: none"><li>• <b>cluster</b>: database-level information</li><li>• <b>node</b>: node-level information</li></ul>                               |
| node_name     | When <b>report_scope</b> is set to <b>node</b> , set this parameter to the name of the corresponding node. (You can run the <b>select * from pg_node_env;</b> command to query the node name.)<br><br>If <b>report_scope</b> is set to <b>cluster</b> , this parameter can be omitted, left blank, or set to <b>NULL</b> . | <ul style="list-style-type: none"><li>• <b>cluster</b>: This value is omitted, left blank or set to <b>NULL</b>.</li><li>• <b>node</b>: a node name in GaussDB</li></ul> |

- create\_wdr\_snapshot()

Description: Manually generates system diagnosis snapshots. This function requires the SYSADMIN permission.

Return type: text

- kill\_snapshot()

Description: Kills the WDR snapshot backend thread. Users who call this function must have the SYSADMIN permission, the REPLICATION permission, or inherit permissions of the built-in role gs\_role\_replication.

Return type: void

- `capture_view_to_json(text,integer)`

Description: Saves the view result to the directory specified by GUC: **perf\_directory**. If **is\_crossdb** is set to **1**, the view is accessed once for all databases. If the value of **is\_crossdb** is **0**, the current database is accessed only once. Only users with the SYSADMIN or MONADMIN permission can execute this function.

Return type: int

- `reset_unique_sql`

Description: Clears the unique SQL statements in the memory of the database node. (The SYSADMIN permission is required.)

Return type: Boolean

**Table 7-47** reset\_unique\_sql parameter description

| Parameter   | Type | Description                                                                                                                                                                                                                                                                  |
|-------------|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scope       | text | Clearance scope type. The options are as follows: <ul style="list-style-type: none"><li>• <b>GLOBAL</b>: Clears all nodes. If the value is <b>GLOBAL</b>, this function can be executed only on the primary node.</li><li>• <b>LOCAL</b>: Clears the current node.</li></ul> |
| clean_type  | text | <ul style="list-style-type: none"><li>• <b>BY_USERID</b>: Unique SQL statements are cleared based on user IDs.</li><li>• <b>BY_CNID</b>: Unique SQL statements are cleared based on primary node IDs.</li><li>• <b>ALL</b>: All data is cleared.</li></ul>                   |
| clean_value | int8 | Clearance value corresponding to the clearance type. If the second parameter is set to <b>ALL</b> , the third parameter does not take effect and can be set to any value.                                                                                                    |

- `wdr_xdb_query(db_name_str text, query text)`

Description: Provides the capability of executing local cross-database queries. For example, when connecting to the **testdb** database, access tables in the **test** database.

```
select col1 from wdr_xdb_query('dbname=test','select col1 from t1') as dd(col1 int);
```

Return type: record

- **pg\_wlm\_jump\_queue(pid int)**

Description: Moves a task to the top of the queue of the primary node of the database.

Return type: Boolean

- **true**: success
- **false**: failure

- **gs\_wlm\_switch\_cgroup(pid int, cgroup text)**

Description: Moves a job to another Cgroup to change the job priority.

Return type: Boolean

- **true**: success
- **false**: failure

- **pv\_session\_memctx\_detail(threadid tid, MemoryContextName text)**

Description: Records information about the memory context

**MemoryContextName** of the thread **tid** into the *threadid\_timestamp.log* file in the *\$GAUSSLOG/pg\_log/\${node\_name}/dumpmem* directory. *threadid* can be obtained from *sessid* in the **GS\_SESSION\_MEMORY\_DETAIL** view. In the officially released version, only the **MemoryContextName** that is an empty string (two single quotation marks indicate that the input is an empty string) is accepted. In this case, all memory context information is recorded. Otherwise, no operation is performed. For the DEBUG version for internal development and test personnel to debug, you can specify the **MemoryContextName** to be counted. In this case, all the memory usage of the context is recorded in the specified file. Only the administrator can execute this function.

Return type: Boolean

- **true**: success
- **false**: failure

- **pg\_shared\_memctx\_detail(MemoryContextName text)**

Description: Records information about the memory context

**MemoryContextName** into the *threadid\_timestamp.log* file in the *\$GAUSSLOG/pg\_log/\${node\_name}/dumpmem* directory. This function is provided only for internal development and test personnel to debug in the DEBUG version. Calling this function in the officially released version does not involve any operation. Only the administrator can execute this function.

Return type: Boolean

- **true**: success
- **false**: failure

- **local\_bgwriter\_stat()**

Description: Displays the information about pages flushed by the bgwriter thread of this instance, number of pages in the candidate buffer chain, and buffer elimination information.

Return type: record

- `local_candidate_stat()`

Description: Displays the number of pages in the candidate buffer chain of this instance and buffer elimination information, including the normal buffer pool and segment buffer pool.

Return type: record

- `local_ckpt_stat()`

Description: Displays the information about checkpoints and flushing pages of the current instance.

Return type: record

- `local_double_write_stat()`

Description: Displays the doublewrite file status of the current instance.

Return type: record

**Table 7-48** `local_double_write_stat` parameters

| Parameter                          | Type | Description                                                                                                                     |
|------------------------------------|------|---------------------------------------------------------------------------------------------------------------------------------|
| <code>node_name</code>             | text | Instance name                                                                                                                   |
| <code>curr_dwn</code>              | int8 | Sequence number of the doublewrite file                                                                                         |
| <code>curr_start_page</code>       | int8 | Start page for restoring the doublewrite file                                                                                   |
| <code>file_trunc_num</code>        | int8 | Number of times that the doublewrite file is reused                                                                             |
| <code>file_reset_num</code>        | int8 | Number of reset times after the doublewrite file is full                                                                        |
| <code>total_writes</code>          | int8 | Total number of I/Os of the doublewrite file                                                                                    |
| <code>low_threshold_writes</code>  | int8 | Number of I/Os for writing doublewrite files with low efficiency (the number of I/O flushing pages at a time is less than 16)   |
| <code>high_threshold_writes</code> | int8 | Number of I/Os for writing doublewrite files with high efficiency (the number of I/O flushing pages at a time is more than 421) |
| <code>total_pages</code>           | int8 | Total number of pages that are flushed to the doublewrite file area                                                             |
| <code>low_threshold_pages</code>   | int8 | Number of pages that are flushed with low efficiency                                                                            |
| <code>high_threshold_pages</code>  | int8 | Number of pages that are flushed with high efficiency                                                                           |
| <code>file_id</code>               | int8 | ID of the current doublewrite file                                                                                              |

- `local_single_flush_dw_stat()`

Description: Displays the elimination of doublewrite files on a single page in the instance.

Return type: record

- `local_pagewriter_stat()`

Description: Displays the page flushing information and checkpoint information of the current instance.

Return type: record

- `local_redo_stat()`

Description: Displays the replay status of the current standby instance.

Return type: record

Note: The returned replay status includes the current replay position and the replay position of the minimum restoration point.

- `local_recovery_status()`

Description: Displays log flow control information about the primary and standby nodes.

Return type: record

- `gs_wlm_node_recover(boolean isForce)`

Description: Obtains top SQL query statement-level statistics recorded in the current memory. If the input parameter is not **0**, the information is cleared from the memory.

Return type: record

- `gs_cgroup_map_ng_conf(group name)`

Description: Reads the Cgroup configuration file of a specified logical database.

Return type: record

- `gs_wlm_switch_cgroup(sess_id int8, cgroup name)`

Description: Switches the Cgroup of a specified session.

Return type: record

- `comm_client_info()`

Description: Queries information about active client connections of a single node.

Return type: SETOF record

- `pg_get_flush_lsn()`

Description: Returns the location of the Xlog flushed from the current node.

Return type: text

- `pg_get_sync_flush_lsn()`

Description: Returns the location of the Xlog flushed by the majority on the current node.

Return type: text

- `dbe_perf.get_global_full_sql_by_timestamp(start_timestamp timestamp with time zone, end_timestamp timestamp with time zone)`

Description: Obtains full SQL information at the database level. The result can be queried only in the system database but cannot be queried in the user database.



Return type: record

**Table 7-49** db\_perf.get\_global\_full\_sql\_by\_timestamp parameter description

| Parameter       | Type                     | Description                              |
|-----------------|--------------------------|------------------------------------------|
| start_timestamp | timestamp with time zone | Start point of the SQL start time range. |
| end_timestamp   | timestamp with time zone | End point of the SQL start time range.   |

- db\_perf.get\_global\_slow\_sql\_by\_timestamp(start\_timestamp timestamp with time zone, end\_timestamp timestamp with time zone)

Description: Obtains slow SQL information at the database level. The result can be queried only in the system database but cannot be queried in the user database.

Return type: record

**Table 7-50** db\_perf.get\_global\_slow\_sql\_by\_timestamp parameter description

| Parameter       | Type                     | Description                              |
|-----------------|--------------------------|------------------------------------------|
| start_timestamp | timestamp with time zone | Start point of the SQL start time range. |
| end_timestamp   | timestamp with time zone | End point of the SQL start time range.   |

- statement\_detail\_decode(detail text, format text, pretty boolean)

Description: Parses the **details** column in a full or slow SQL statement. The result can be queried only in the system database but cannot be queried in the user database.

Return type: text

**Table 7-51** statement\_detail\_decode parameter description

| Parameter     | Type | Description                                                |
|---------------|------|------------------------------------------------------------|
| <b>detail</b> | text | Set of events generated by the SQL statement (unreadable). |
| format        | text | Parsing output format. The value is <b>plaintext</b> .     |

| Parameter | Type    | Description                                                                                                                                                                                                                                                                                                                                     |
|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pretty    | boolean | Specifies whether to display the text in pretty format when <b>format</b> is set to <b>plaintext</b> . The options are as follows: <ul style="list-style-type: none"> <li>The value <b>true</b> indicates that events are separated by <b>\n</b>.</li> <li>The value <b>false</b> indicates that events are separated by commas (,).</li> </ul> |

- `pg_control_system()`  
Description: Returns the status of the system control file.  
Return type: SETOF record
- `pg_control_checkpoint()`  
Description: Returns the system checkpoint status.  
Return type: SETOF record
- `get_prepared_pending_xid`  
Description: Returns nextxid when restoration is complete.  
Parameter: nan  
Return type: text
- `pg_clean_region_info`  
Description: Clears the region map.  
Parameter: nan  
Return type: character varying
- `pg_get_delta_info`  
Description: Obtains delta information from a single DN.  
Parameter: rel text, schema\_name text  
Return type: part\_name text, live\_tuple bigint, data\_size bigint, and blocknum bigint
- `pg_get_replication_slot_name`  
Description: Obtains the slot name.  
Parameter: nan  
Return type: text
- `pg_get_running_xacts`  
Description: Obtains running xact.  
Parameter: **nan**  
Return type: handle integer, gxid xid, state tinyint, node text, xmin xid, vacuum Boolean, timeline bigint, prepare\_xid xid, pid bigint, and next\_xid xid
- `pg_get_variable_info`  
Description: Obtains the shared memory variable *cache*.  
Parameter: nan  
Return type: node\_name text, nextOid oid, nextXid xid, oldestXid xid, xidVacLimit xid, oldestXidDB oid, lastExtendCSNLogpage xid,

- startExtendCSNLogpage xid, nextCommitSeqNo xid, latestCompletedXid xid, and startupMaxXid xid
- **pg\_get\_xidlimit**  
Description: Obtains transaction ID information from the shared memory.  
Parameter: nan  
Return type: nextXid xid, oldestXid xid, xidVacLimit xid, xidWarnLimit xid, xidStopLimit xid, xidWrapLimit xid, and oldestXidDB oid
- **pg\_stat\_file\_recursive**  
Description: Lists all files in the path.  
Parameter: location text  
Return type: path text, filename text, size bigint, and isdir Boolean
- **pg\_stat\_get\_activity\_for temptable**  
Description: Returns records of backend processes related to the temporary table.  
Parameter: nan  
Return type: datid oid, timelineid integer, tempid integer, and sessionid bigint
- **pg\_stat\_get\_activity\_ng**  
Description: Returns records of backend processes related to nodegroup.  
Parameter: pid bigint  
Return type: datid oid, pid bigint, sessionid bigint, and node\_group text
- **pg\_stat\_get\_cgroup\_info**  
Description: Returns Cgroup information.  
Parameter: **nan**  
Return type: cgroup\_name text, percent integer, usage\_percent integer, shares bigint, usage bigint, cpuset text, relpath text, valid text, and node\_group text
- **pg\_stat\_get\_realtime\_info\_internal**  
Description: Returns real-time information. Currently, this API is unavailable. **FailedToGetSessionInfo** is returned.  
Parameter: oid, oid, bigint, cstring, oid  
Return type: text
- **pg\_test\_err\_contain\_err**  
Description: Tests the error type and return information.  
Parameter: integer  
Return type: void
- **get\_global\_user\_transaction()**  
Description: Returns transaction information about each user on all nodes.  
Return type: node\_name name, username name, commit\_counter bigint, rollback\_counter bigint, resp\_min bigint, resp\_max bigint, resp\_avg bigint, resp\_total bigint, bg\_commit\_counter bigint, bg\_rollback\_counter bigint, bg\_resp\_min bigint, bg\_resp\_max bigint, bg\_resp\_avg bigint, and bg\_resp\_total bigint
- **pg\_collation\_for**

Description: Returns the collation rule corresponding to the input parameter string.

Parameter: any (Explicit type conversion is required for constants.)

Return type: text

- `pgxc_unlock_for_sp_database(name Name)`

Description: Releases a specified database lock.

Parameter: database name

Return type: Boolean

- `pgxc_lock_for_sp_database(name Name)`

Description: Locks a specified database.

Parameter: database name

Return type: Boolean

- `copy_error_log_create()`

Description: Creates the error table (**public.pgxc\_copy\_error\_log**) required for creating the **COPY FROM** error tolerance mechanism.

Return type: Boolean

#### NOTE

- This function attempts to create the **public.pgxc\_copy\_error\_log** table. For details about the table, see [Table 7-52](#).
- In addition, it creates a B-tree index on the **relname** column and executes **REVOKE ALL on public.pgxc\_copy\_error\_log FROM public** to manage permissions on the error table (the permissions are the same as those of the **COPY** statement).
- **public.pgxc\_copy\_error\_log** is a row-store table. Therefore, this function can be executed and **COPY** error tolerance is available only when row-store tables can be created in the database instance. Note that after the GUC parameter **enable\_hadoop\_env** is enabled, row-store tables cannot be created in the database instance (the default value is **off** for GaussDB).
- Same as the error table and the **COPY** statement, the function requires **SYSADMIN** or higher permissions.
- If the **public.pgxc\_copy\_error\_log** table or the **copy\_error\_log\_relname\_idx** index exists before the function creates it, the function will report an error and roll back.

**Table 7-52** Error table `public.pgxc_copy_error_log`

| Column    | Type                     | Description                                                              |
|-----------|--------------------------|--------------------------------------------------------------------------|
| relname   | character varying        | Table name, in the form of <i>Schema name.Table name</i> .               |
| begintime | timestamp with time zone | Time when a data format error was reported                               |
| filename  | character varying        | Name of the source data file where a data format error occurs            |
| lineno    | bigint                   | Number of the row where a data format error occurs in a source data file |

| Column    | Type | Description                                               |
|-----------|------|-----------------------------------------------------------|
| rawrecord | text | Raw record of a data format error in the source data file |
| detail    | text | Error details                                             |

- dynamic\_func\_control(scope text, function\_name text, action text, "{params}" text[])

Description: Dynamically enables built-in functions. Currently, only full SQL statements can be dynamically enabled.

Return type: record

**Table 7-53** Parameter description of dynamic\_func\_control

| Parameter     | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scope         | text   | Scope where the function is to be dynamically enabled. Currently, only <b>LOCAL</b> is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| function_name | text   | Function name. Currently, only <b>STMT</b> is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| action        | text   | When <b>function_name</b> is set to <b>STMT</b> , the value of <b>action</b> can only be <b>TRACK</b> , <b>UNTRACK</b> , <b>LIST</b> , or <b>CLEAN</b> . <ul style="list-style-type: none"><li>• <b>TRACK</b>: records the full SQL information of normalized SQL statements.</li><li>• <b>UNTRACK</b>: cancels the recording of full SQL information of normalized SQL statements.</li><li>• <b>LIST</b>: lists normalized SQL information that is recorded in the current track.</li><li>• <b>CLEAN</b>: cleans normalized SQL information that is recorded in the current track.</li></ul> |
| params        | text[] | When <b>function_name</b> is set to <b>STMT</b> , the parameters corresponding to different actions are set as follows: <ul style="list-style-type: none"><li>• <b>TRACK</b>: '{"Normalized SQLID", "L0/L1/L2"}'</li><li>• <b>UNTRACK</b>: '{"Normalized SQLID"}'</li><li>• <b>LIST</b> - '{}'</li><li>• <b>CLEAN</b> - '{}'</li></ul>                                                                                                                                                                                                                                                        |

- gs\_parse\_page\_bypath(path text, blocknum bigint, relation\_type text, read\_memory boolean)

Description: Parses a specified table page and returns the path for storing the parsed content.

Return type: text

Note: Only SYSADMIN or OPRADMIN can execute this function.

**Table 7-54** gs\_parse\_page\_bypath parameters

| Parameter     | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| path          | text    | <ul style="list-style-type: none"><li>For an ordinary table or segment-page table, the relative path is <i>Tablespace name/Database OID/Relfilenode of the table (physical file name)</i>. For example, <b>base/16603/16394</b>.</li><li>You can run the <b>pg_relation_filepath(table_name text)</b> command to query the relative path of the table file. To obtain the path of the partitioned table, view the <b>pg_partition</b> system catalog and call <b>pg_partition_filepath(partition_oid)</b>.</li><li>Valid path formats are as follows:<ul style="list-style-type: none"><li>global/relNode</li><li>base/dbNode/relNode</li><li>pg_tblspc/spcNode/version_dir/dbNode/relNode</li></ul></li></ul> |
| blocknum      | bigint  | <ul style="list-style-type: none"><li><b>-1</b>: Information about all blocks (forcibly parsed from disks)</li><li><b>0–MaxBlockNumber</b>: Information about the corresponding block</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| relation_type | text    | <ul style="list-style-type: none"><li><b>heap</b>: Astore table</li><li><b>uheap</b>: Ustore table</li><li><b>btree</b>: B-tree index</li><li><b>ubtree</b>: UB-tree index</li><li><b>segment</b>: Segment-page table</li><li><b>indexurq</b>: Ustore recycling queue</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                |
| read_memory   | boolean | <ul style="list-style-type: none"><li><b>false</b>: The system parses the page from the disk file.</li><li><b>true</b>: The system attempts to parse the page from the shared buffer. If the page does not exist in the shared buffer, the system parses the page from the disk file.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                |

- gs\_xlogdump\_lsn(start\_lsn text, end\_lsn text)**  
Description: Parses Xlogs within the specified LSN range and returns the path for storing the parsed content. You can use **pg\_current\_xlog\_location()** to obtain the current Xlog location.  
Return type: text  
Parameters: LSN start position and LSN end position

Note: Only SYSADMIN or OPRADMIN can execute this function.

- `gs_xlogdump_xid(c_xid xid)`

Description: Parses Xlogs of a specified XID and returns the path for storing the parsed content. You can use `txid_current()` to obtain the current XID.

Parameter: XID

Return type: text

Note: Only SYSADMIN or OPRADMIN can execute this function.

- `gs_xlogdump_tablepath(path text, blocknum bigint, relation_type text)`

Description: Parses logs corresponding to a specified table page and returns the path for storing the parsed content.

Return type: text

Note: Only SYSADMIN or OPRADMIN can execute this function.

**Table 7-55** `gs_xlogdump_tablepath` parameters

| Parameter     | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| path          | text   | <ul style="list-style-type: none"><li>• For an ordinary table or segment-page table, the relative path is <i>Tablespace name/Database OID/Relfilenode of the table (physical file name)</i>. For example, <b>base/16603/16394</b>.</li><li>• You can run the <b><code>pg_relation_filepath(table_name text)</code></b> command to query the relative path of the table file. To obtain the path of the partitioned table, view the <b><code>pg_partition</code></b> system catalog and call <b><code>pg_partition_filepath(partition_oid)</code></b>.</li><li>• Valid path formats are as follows:<ul style="list-style-type: none"><li>– global/relNode</li><li>– base/dbNode/relNode</li><li>– pg_tblspc/spcNode/version_dir/dbNode/relNode</li></ul></li></ul> |
| blocknum      | bigint | <ul style="list-style-type: none"><li>• <b>-1</b>: Information about all blocks (forcibly parsed from disks)</li><li>• <b>0–MaxBlockNumber</b>: Information about the corresponding block</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| relation_type | text   | <ul style="list-style-type: none"><li>• <b>heap</b>: Astore table</li><li>• <b>uheap</b>: Ustore table</li><li>• <b>btree</b>: B-tree index</li><li>• <b>ubtree</b>: UB-tree index</li><li>• <b>segment</b>: Segment-page table</li><li>• <b>indexurq</b>: Ustore recycling queue</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

- `gs_xlogdump_parsepage_tablepath(path text, blocknum bigint, relation_type text, read_memory boolean)`

Description: Parses the specified table page and logs corresponding to the table page and returns the path for storing the parsed content. It can be regarded as one execution of **gs\_parse\_page\_bypath** and **gs\_xlogdump\_tablepath**. The prerequisite for executing this function is that the table file exists. To view logs of deleted tables, call **gs\_xlogdump\_tablepath**.

Return type: text

Note: Only SYSADMIN or OPRADMIN can execute this function.

**Table 7-56** `gs_xlogdump_parsepage_tablepath` parameters

| Parameter     | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| path          | text   | <ul style="list-style-type: none"><li>• For an ordinary table or segment-page table, the relative path is <i>Tablespace name/ Database OID/ Relfilenode of the table (physical file name)</i>. For example, <b>base/16603/16394</b>.</li><li>• You can run the <b>pg_relation_filepath(table_name text)</b> command to query the relative path of the table file. To obtain the path of the partitioned table, view the <b>pg_partition</b> system catalog and call <b>pg_partition_filepath(partition_oid)</b>.</li><li>• Valid path formats are as follows:<ul style="list-style-type: none"><li>– global/relNode</li><li>– base/dbNode/relNode</li><li>– pg_tblspc/spcNode/version_dir/dbNode/relNode</li></ul></li></ul> |
| blocknum      | bigint | <ul style="list-style-type: none"><li>• <b>-1</b>: Information about all blocks (forcibly parsed from disks)</li><li>• <b>0–MaxBlockNumber</b>: Information about the corresponding block</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| relation_type | text   | <ul style="list-style-type: none"><li>• <b>heap</b>: Astore table</li><li>• <b>uheap</b>: Ustore table</li><li>• <b>btree</b>: B-tree index</li><li>• <b>ubtree</b>: UB-tree index</li><li>• <b>segment</b>: Segment-page table</li><li>• <b>indexurq</b>: Ustore recycling queue</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  |



| Parameter   | Type    | Description                                                                                                                                                                                                                                                                                         |
|-------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| read_memory | boolean | <ul style="list-style-type: none"><li>• <b>false</b>: The system parses the page from the disk file.</li><li>• <b>true</b>: The system attempts to parse the page from the shared buffer. If the page does not exist in the shared buffer, the system parses the page from the disk file.</li></ul> |

- `gs_index_verify(Oid oid, uint32 blkno)`

Description: Checks whether the sequence of keys on the UB-tree index page or index tree is correct.

Return type: record

**Table 7-57** `gs_index_verify` parameters

| Parameter | Type   | Description                                                                                                                                                                                                                                    |
|-----------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid       | Oid    | <ul style="list-style-type: none"><li>• Index file relfilenode, which can be queried using <b><code>select relfilenode from pg_class where relname='Index file name'</code></b>.</li></ul>                                                     |
| blkno     | uint32 | <ul style="list-style-type: none"><li>• <b>0</b>: indicates that all pages in the index tree are checked.</li><li>• If the value is greater than 0, the index page whose page code is equal to the value of <b>blkno</b> is checked.</li></ul> |

- `gs_index_recycle_queue(Oid oid, int type, uint32 blkno)`

Description: Parses the UB-tree index recycling queue information.

Return type: record

**Table 7-58** `gs_index_recycle_queue` parameters

| Parameter | Type   | Description                                                                                                                                                                                                                                                 |
|-----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid       | Oid    | <ul style="list-style-type: none"><li>• Index file relfilenode, which can be queried using <b><code>select relfilenode from pg_class where relname='Index file name'</code></b>.</li></ul>                                                                  |
| type      | int    | <ul style="list-style-type: none"><li>• <b>0</b>: indicates that the entire queue to be recycled is parsed.</li><li>• <b>1</b>: indicates that the entire empty page queue is parsed.</li><li>• <b>2</b>: indicates that a single page is parsed.</li></ul> |
| blkno     | uint32 | <ul style="list-style-type: none"><li>• ID of the recycling queue page. This parameter is valid only when <b>type</b> is set to <b>2</b>. The value of <b>blkno</b> ranges from 1 to 4294967294.</li></ul>                                                  |

- `gs_stat_wal_entrytable(int64 idx)`  
Description: Exports the content of the WAL insertion status table in the Xlog.  
Return type: record

**Table 7-59** `gs_stat_wal_entrytable` parameters

| Category         | Parameter | Type   | Description                                                                                                                                                                                                    |
|------------------|-----------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | idx       | int64  | <ul style="list-style-type: none"><li>• <b>-1</b>: queries all elements in an array.</li><li>• <b>0–Maximum value</b>: content of a specific array element.</li></ul>                                          |
| Output parameter | idx       | uint64 | Records the indexes in the corresponding array.                                                                                                                                                                |
| Output parameter | endsln    | uint64 | Records the LSN label.                                                                                                                                                                                         |
| Output parameter | lrc       | int32  | Records the corresponding LRC.                                                                                                                                                                                 |
| Output parameter | status    | uint32 | Specifies whether the Xlog corresponding to the current entry has been completely copied to the WAL buffer. <ul style="list-style-type: none"><li>• <b>0</b>: Not copied.</li><li>• <b>1</b>: Copied</li></ul> |

- `gs_walwriter_flush_position()`  
Description: Outputs the refresh position of write-ahead logs.  
Return type: record

**Table 7-60** gs\_walwriter\_flush\_position parameters

| Category         | Parameter               | Type   | Description                                                                                                                                                                |
|------------------|-------------------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Output parameter | last_flush_status_entry | int32  | Index obtained after the Xlog flushes the tblEntry of the last flushed disk.                                                                                               |
| Output parameter | last_scanned_lrc        | int32  | LRC obtained after the Xlog flushes the last tblEntry scanned last time.                                                                                                   |
| Output parameter | curr_lrc                | int32  | Latest LRC usage in the WALInsertStatusEntry status table. The LRC indicates the LRC value corresponding to the WALInsertStatusEntry when the next Xlog record is written. |
| Output parameter | curr_byte_pos           | uint64 | The latest Xlog location after the Xlog is written to the WAL file, which is also the next Xlog insertion point.                                                           |
| Output parameter | prev_byte_size          | uint32 | Length of the previous Xlog record.                                                                                                                                        |
| Output parameter | flush_result            | uint64 | Location of the current global Xlog flush.                                                                                                                                 |
| Output parameter | send_result             | uint64 | Xlog sending location on the current host.                                                                                                                                 |
| Output parameter | shm_rqst_write_pos      | uint64 | The write position of the LogwrtRqst request in the XLogCtl recorded in the shared memory.                                                                                 |
| Output parameter | shm_rqst_flush_pos      | uint64 | The flush position of the LogwrtRqst request in the XLogCtl recorded in the shared memory.                                                                                 |
| Output parameter | shm_result_write_pos    | uint64 | The write position of the LogwrtResult request in the XLogCtl recorded in the shared memory.                                                                               |

| Category         | Parameter            | Type   | Description                                                                                  |
|------------------|----------------------|--------|----------------------------------------------------------------------------------------------|
| Output parameter | shm_result_flush_pos | uint64 | The flush position of the LogwrtResult request in the XLogCtl recorded in the shared memory. |
| Output parameter | curr_time            | text   | Current time.                                                                                |

- gs\_walwriter\_flush\_stat(int operation)

Description: Collects statistics on the frequency of writing and synchronizing WALs, data volume, and Xlog file information.

Return type: record

**Table 7-61** gs\_walwriter\_flush\_stat parameters

| Category         | Parameter                    | Type   | Description                                                                                                                                                                                                                                    |
|------------------|------------------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | operation                    | int    | <ul style="list-style-type: none"><li>• <b>-1</b>: Disables the statistics function. (Default value)</li><li>• <b>0</b>: Enable the statistics function.</li><li>• <b>1</b>: Query statistics.</li><li>• <b>2</b>: Reset statistics.</li></ul> |
| Output parameter | write_times                  | uint64 | Number of times that the Xlog calls the <b>write</b> API.                                                                                                                                                                                      |
| Output parameter | sync_times                   | uint64 | Number of times that the Xlog calls the <b>sync</b> API.                                                                                                                                                                                       |
| Output parameter | total_xlog_sync_bytes        | uint64 | Total number of backend thread requests for writing data to Xlogs.                                                                                                                                                                             |
| Output parameter | total_actual_xlog_sync_bytes | uint64 | Total number of Xlogs that call the <b>sync</b> API for disk flushing.                                                                                                                                                                         |

| Category         | Parameter              | Type   | Description                                                                        |
|------------------|------------------------|--------|------------------------------------------------------------------------------------|
| Output parameter | avg_write_bytes        | uint32 | Number of Xlogs written each time the <b>XLogWrite</b> API is called.              |
| Output parameter | avg_actual_write_bytes | uint32 | Number of Xlogs written each time the <b>write</b> API is called.                  |
| Output parameter | avg_sync_bytes         | uint32 | Average number of Xlogs for requesting the <b>sync</b> API each time.              |
| Output parameter | avg_actual_sync_bytes  | uint32 | Actual number of Xlogs for disk flushing by calling the <b>sync</b> API each time. |
| Output parameter | total_write_time       | uint64 | Total time of calling the write operation (unit: $\mu$ s).                         |
| Output parameter | total_sync_time        | uint64 | Total time for calling the <b>sync</b> API (unit: $\mu$ s).                        |
| Output parameter | avg_write_time         | uint32 | Average time for calling the <b>write</b> API each time (unit: $\mu$ s).           |
| Output parameter | avg_sync_time          | uint32 | Average time for calling the <b>sync</b> API each time (unit: $\mu$ s).            |
| Output parameter | curr_init_xlog_segno   | uint64 | ID of the latest Xlog segment file.                                                |
| Output parameter | curr_open_xlog_segno   | uint64 | ID of the Xlog segment file that is being written.                                 |

| Category         | Parameter       | Type | Description                               |
|------------------|-----------------|------|-------------------------------------------|
| Output parameter | last_reset_time | text | Time when statistics were last collected. |
| Output parameter | curr_time       | text | Current time.                             |

- gs\_catalog\_attribute\_records()

Description: Returns the definition of each column in a system catalog specified by OID. Only common system catalogs whose OIDs are less than 10000 are supported. Indexes and TOAST tables are not supported.

Parameter: OID of the system catalog

Return type: record

- gs\_comm\_proxy\_thread\_status()

Description: Collects statistics on data packets sent and received by the proxy communications library **comm\_proxy** when a user-mode network is configured for the database instance.

Parameter: nan

Return type: record

#### NOTE

The query result of this function is displayed only when the user-mode network is deployed in a centralized environment and **enable\_dfx in comm\_proxy\_attr** is set to **true**. In other scenarios, an error message is displayed, indicating that queries are not supported.

- pg\_ls\_tmpdir()

Description: Returns the name, size, and last modification time of each file in the temporary directory (**pgsql\_tmp**) of the default tablespace.

Parameter: nan

Return type: record

Note: Only SYSADMIN or MONADMIN can execute this function.

| Category         | Parameter    | Type       | Description                 |
|------------------|--------------|------------|-----------------------------|
| Output parameter | name         | text       | File name                   |
| Output parameter | size         | int8       | File size (unit: byte)      |
| Output parameter | modification | timestampz | Last file modification time |

- pg\_ls\_tmpdir(oid)

Description: Returns the name, size, and last modification time of each file in the temporary directory (**pgsql\_tmp**) of the specified tablespace.

Parameter: oid

Return type: record

Note: Only SYSADMIN or MONADMIN can execute this function.

| Category         | Parameter    | Type       | Description                 |
|------------------|--------------|------------|-----------------------------|
| Input parameter  | oid          | oid        | Tablespace ID               |
| Output parameter | name         | text       | File name                   |
| Output parameter | size         | int8       | File size (unit: byte)      |
| Output parameter | modification | timestampz | Last file modification time |

- pg\_ls\_waldir()

Description: Returns the name, size, and last modification time of each file in the WAL directory.

Parameter: nan

Return type: record

Note: Only SYSADMIN or MONADMIN can execute this function.

| Category         | Parameter    | Type       | Description                 |
|------------------|--------------|------------|-----------------------------|
| Output parameter | name         | text       | File name                   |
| Output parameter | size         | int8       | File size (unit: byte)      |
| Output parameter | modification | timestampz | Last file modification time |

- gs\_write\_term\_log(void)

Description: Writes a log to record the current **term** value of a DN. The standby DN returns **false**. After the data is successfully written to the primary DN, **true** is returned.

Return type: Boolean

### 7.5.23.13 Undo System Functions

- gs\_undo\_meta(type, zoneld, location)

Description: Specifies metadata of each module in the undo system.

Parameter description:

- **type** (metadata type)  
The value **0** indicates the metadata corresponding to **Undo Zone(Record)**.  
The value **1** indicates the metadata corresponding to **Undo Zone(Transaction Slot)**.  
The value **2** indicates the metadata corresponding to **Undo Space(Record)**.  
The value **3** indicates the metadata corresponding to **Undo Space(Transaction Slot)**.
- **zoneId** (undo zone ID)  
The value **-1** indicates the metadata of all undo zones.  
The value range 0–1024 x 1024 indicates the metadata of the corresponding zone ID.
- **location** (read location)  
The value **0** indicates that data is read from the current memory.  
The value **1** indicates that data is read from a physical file.

Return type: record

**Table 7-62** Output example of `gs_undo_meta(0,-1,0)`

| Category         | Parameter   | Type | Description                                                                                                                                        |
|------------------|-------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Output parameter | zoneId      | oid  | ID of the undo zone                                                                                                                                |
| Output parameter | persistType | oid  | Persistence level                                                                                                                                  |
| Output parameter | insert      | text | Position of the next undo record to be inserted                                                                                                    |
| Output parameter | discard     | text | Position of the undo record that is recycled in common mode                                                                                        |
| Output parameter | end         | text | Position of the undo record that is forcibly recycled. Values smaller than the value of this parameter indicate that the record has been recycled. |
| Output parameter | used        | text | Used undo space                                                                                                                                    |
| Output parameter | lsn         | text | LSN to modify the zone                                                                                                                             |



| Category         | Parameter | Type | Description                         |
|------------------|-----------|------|-------------------------------------|
| Output parameter | pid       | oid  | ID of the process bound to the zone |

- `gs_undo_translot(location, zoneld)`

Description: Specifies transaction slot information of the undo system.

Parameter description:

- **location** (read location)  
The value **0** indicates that data is read from the current memory.  
The value **1** indicates that data is read from a physical file.
- **zoneld** (undo zone ID)  
The value **-1** indicates the metadata of all undo zones.  
The value range **0–1024 x 1024** indicates the metadata of the corresponding zone ID.

Return type: record

**Table 7-63** Output example of `gs_undo_translot(0,-1)`

| Category         | Parameter    | Type | Description                                                                                                                                          |
|------------------|--------------|------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Output parameter | groupid      | oid  | ID of the used undo zone                                                                                                                             |
| Output parameter | xactId       | text | Transaction ID                                                                                                                                       |
| Output parameter | startUndoPtr | text | Position where the undo record is inserted at the start of the transaction corresponding to a slot                                                   |
| Output parameter | endUndoPtr   | text | Position where the undo record is inserted at the end of the transaction corresponding to a slot                                                     |
| Output parameter | lsn          | text | Pointer to the corresponding slot                                                                                                                    |
| Output parameter | slot_states  | oid  | Transaction state. The options are as follows: <b>0</b> : committed; <b>1</b> : being executed; <b>2</b> : being rolled back; <b>3</b> : rolled back |

- `gs_stat_undo()`

Description: Undo statistics.

Return type: record

**Table 7-64** gs\_stat\_undo parameters

| Category         | Parameter              | Type   | Description                                                                                                                                                                                                                                                                                |
|------------------|------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Output parameter | curr_used_zone_count   | uint32 | Number of used undo zones.                                                                                                                                                                                                                                                                 |
| Output parameter | top_used_zones         | text   | Information about the first three undo zones with the maximum usage. The output format is as follows: <ul style="list-style-type: none"><li>● <b>zoneld1</b>: used size of zone 1.</li><li>● <b>zoneld2</b>: used size of zone 2.</li><li>● <b>zoneld3</b>: used size of zone 3.</li></ul> |
| Output parameter | curr_used_undo_size    | uint32 | Total size of the undo tablespace that is being used. The unit is MB.                                                                                                                                                                                                                      |
| Output parameter | undo_threshold         | uint32 | Calculation result of the value of the GUC parameter <b>undo_space_limit_size</b> x 80%. The unit is MB.                                                                                                                                                                                   |
| Output parameter | oldest_xid_in_undo     | uint64 | XID of the transaction recycled to the undo space. The undo records generated by the transaction whose XID is smaller than the value of XID are recycled.                                                                                                                                  |
| Output parameter | oldest_xmin            | uint64 | Oldest active transaction.                                                                                                                                                                                                                                                                 |
| Output parameter | total_undo_chain_len   | int64  | Total length of all accessed undo chains.                                                                                                                                                                                                                                                  |
| Output parameter | max_undo_chain_len     | int64  | Maximum length of the accessed undo chain.                                                                                                                                                                                                                                                 |
| Output parameter | create_undo_file_count | uint32 | Number of created undo files.                                                                                                                                                                                                                                                              |

| Category         | Parameter               | Type   | Description                   |
|------------------|-------------------------|--------|-------------------------------|
| Output parameter | discard_undo_file_count | uint32 | Number of deleted undo files. |

- `gs_undo_record(undoptr)`  
Description: Undo record resolution.  
Parameter description:
  - **undoptr** (undo record pointer)Return type: record
- `gs_undo_dump_parsepage_mv(relpath text, blkno bigint, reltype text, rmem boolean)`  
Description: Parses the page header information of the disk page in the Ustore table, header information of each tuple, flag bit information, and all historical undo version information that can be queried.  
Return type: text  
Note: Only SYSADMIN or OPRADMIN can execute this function.

**NOTE**

Currently, this API supports only Ustore tables.

**Table 7-65** `gs_undo_dump_parsepage_mv` parameters

| Category        | Parameter | Type   | Description                                                                                                                                                                                                                                         |
|-----------------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter | relpath   | text   | Relative path of the Ustore table data file, in the format of <i>Tablespace name/Database OID/relfilenode</i> . For example, <b>base/16603/16384</b> . You can run the <b>pg_relation_filepath('tablename')</b> command to query the relative path. |
| Input parameter | blkno     | bigint | <ul style="list-style-type: none"><li>• <b>-1</b>: All block pages are parsed.</li><li>• <b>0–MaxBlocNumber</b>: A specified block page is parsed.</li></ul>                                                                                        |
| Input parameter | reltype   | text   | Table type. Currently, only the Ustore table is supported. The value is <b>uheap</b> .                                                                                                                                                              |

| Category         | Parameter | Type    | Description                                                                                                                                                                             |
|------------------|-----------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | rmem      | boolean | <ul style="list-style-type: none"><li>false</li><li>true</li></ul> Currently, the value can only be <b>false</b> , indicating that the corresponding page is parsed from the disk file. |
| Output parameter | output    | text    | Absolute path of the parsing result file.                                                                                                                                               |

- `gs_undo_meta_dump_zone(zone_id int, read_memory boolean)`  
Description: Parses zone metadata in the undo module.  
Return type: record

**Table 7-66** gs\_undo\_meta\_dump\_zone parameters

| Category         | Parameter    | Type    | Description                                                                                                                                                                    |
|------------------|--------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | zone_id      | int     | Undo zone ID <ul style="list-style-type: none"><li>-1: All undo zones are queried.</li><li><b>0-1048575</b>: The undo zone metadata is queried based on the zone ID.</li></ul> |
| Input parameter  | read_memory  | boolean | <ul style="list-style-type: none"><li><b>true</b>: Data is read from the current memory.</li><li><b>false</b>: Data is read from the physical file.</li></ul>                  |
| Output parameter | zone_id      | int     | Undo zone ID                                                                                                                                                                   |
| Output parameter | persist_type | int     | Persistence level <ul style="list-style-type: none"><li><b>0</b>: ordinary table</li><li><b>1</b>: unlogged table</li><li><b>2</b>: temporary table</li></ul>                  |
| Output parameter | insert       | text    | Position of the next undo record to be inserted                                                                                                                                |

| Category         | Parameter    | Type | Description                                                                                                                                        |
|------------------|--------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Output parameter | discard      | text | Position of the undo record that is recycled in common mode                                                                                        |
| Output parameter | forcediscard | text | Position of the undo record that is forcibly recycled. Values smaller than the value of this parameter indicate that the record has been recycled. |
| Output parameter | lsn          | text | LSN to modify the zone                                                                                                                             |

- `gs_undo_meta_dump_spaces(zone_id int, read_memory boolean)`  
Description: Parses the metadata of the undo record space and undo slot space in the undo module.  
Return type: record

**Table 7-67** `gs_undo_meta_dump_spaces` parameters

| Category         | Parameter             | Type    | Description                                                                                                                                                                               |
|------------------|-----------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | zone_id               | int     | Undo zone ID <ul style="list-style-type: none"><li>• <b>-1</b>: All undo zones are queried.</li><li>• <b>0-1048575</b>: The undo zone metadata is queried based on the zone ID.</li></ul> |
| Input parameter  | read_memory           | boolean | <ul style="list-style-type: none"><li>• <b>true</b>: Data is read from the current memory.</li><li>• <b>false</b>: Data is read from the physical file.</li></ul>                         |
| Output parameter | zone_id               | int     | Undo zone ID                                                                                                                                                                              |
| Output parameter | undorecord_space_tail | text    | End position of the undo record space                                                                                                                                                     |

| Category         | Parameter             | Type | Description                             |
|------------------|-----------------------|------|-----------------------------------------|
| Output parameter | undorecord_space_head | text | Start position of the undo record space |
| Output parameter | undorecord_space_lsn  | text | LSN of the modified undo record space   |
| Output parameter | undoslot_space_tail   | text | End position of the undo slot space     |
| Output parameter | undoslot_space_head   | text | Start position of the undo slot space   |
| Output parameter | undo_slot_space_lsn   | text | LSN of the modified undo slot space     |

- `gs_undo_meta_dump_slot(zone_id int, read_memory boolean)`  
Description: Parses slot metadata in the undo module.  
Return type: record

**Table 7-68** `gs_undo_meta_dump_slot` parameters

| Category         | Parameter   | Type    | Description                                                                                                                                                                               |
|------------------|-------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | zone_id     | int     | Undo zone ID <ul style="list-style-type: none"><li>• <b>-1</b>: All undo zones are queried.</li><li>• <b>0-1048575</b>: The undo zone metadata is queried based on the zone ID.</li></ul> |
| Input parameter  | read_memory | boolean | <ul style="list-style-type: none"><li>• <b>true</b>: Data is read from the current memory.</li><li>• <b>false</b>: Data is read from the physical file.</li></ul>                         |
| Output parameter | zone_id     | int     | Undo zone ID                                                                                                                                                                              |

| Category         | Parameter          | Type | Description                                                                                                          |
|------------------|--------------------|------|----------------------------------------------------------------------------------------------------------------------|
| Output parameter | allocate           | text | Allocation position of the undo transaction slot                                                                     |
| Output parameter | recycle            | text | Recycling position of the undo transaction slot                                                                      |
| Output parameter | frozen_xid         | text | Frozen XID, which is used to determine the visibility                                                                |
| Output parameter | global_frozen_xid  | text | Minimum frozen XID in the system. Transactions whose XID is smaller than the value of this parameter are visible.    |
| Output parameter | recycle_xid        | text | Recycled XID. Transactions whose XID is smaller than the value of this parameter are recycled.                       |
| Output parameter | global_recycle_xid | text | Minimum recycled XID in the system. Transactions whose XID is smaller than the value of this parameter are recycled. |

- gs\_undo\_translot\_dump\_slot(zone\_id int, read\_memory boolean)  
Description: Parses slots in a zone.  
Return type: record

**Table 7-69** gs\_undo\_translot\_dump\_slot parameters

| Category        | Parameter | Type | Description                                                                                                                                                             |
|-----------------|-----------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter | zone_id   | oid  | Undo zone ID <ul style="list-style-type: none"><li>-1: All undo zones are queried.</li><li>0-1048575: The undo zone metadata is queried based on the zone ID.</li></ul> |

| Category         | Parameter      | Type    | Description                                                                                                                                                                                      |
|------------------|----------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | read_memory    | boolean | <ul style="list-style-type: none"><li>• <b>true</b>: Data is read from the current memory.</li><li>• <b>false</b>: Data is read from the physical file.</li></ul>                                |
| Output parameter | zone_id        | text    | Undo zone ID                                                                                                                                                                                     |
| Output parameter | slot_xid       | text    | Transaction ID                                                                                                                                                                                   |
| Output parameter | start_undo_ptr | text    | Position where the undo record is inserted at the start of the transaction corresponding to a slot                                                                                               |
| Output parameter | end_undo_ptr   | text    | Position where the undo record is inserted at the end of the transaction corresponding to a slot                                                                                                 |
| Output parameter | slot_ptr       | text    | Position of a transaction slot.                                                                                                                                                                  |
| Output parameter | slot_states    | oid     | Transaction state <ul style="list-style-type: none"><li>• <b>0</b>: committed</li><li>• <b>1</b>: being executed</li><li>• <b>2</b>: being rolled back</li><li>• <b>3</b>: rolled back</li></ul> |

- `gs_undo_translot_dump_xid(slot_xid xid, read_memory boolean)`  
Description: Parses the slot in a zone based on the XID.  
Return type: record

**Table 7-70** `gs_undo_translot_dump_xid` parameters

| Category        | Parameter | Type | Description       |
|-----------------|-----------|------|-------------------|
| Input parameter | slot_xid  | xid  | XID to be queried |



| Category         | Parameter      | Type    | Description                                                                                                                                                                                      |
|------------------|----------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | read_memory    | boolean | <ul style="list-style-type: none"><li>• <b>true</b>: Data is read from the current memory.</li><li>• <b>false</b>: Data is read from the physical file.</li></ul>                                |
| Output parameter | zone_id        | text    | Undo zone ID                                                                                                                                                                                     |
| Output parameter | slot_xid       | text    | Transaction ID                                                                                                                                                                                   |
| Output parameter | start_undo_ptr | text    | Position where the undo record is inserted at the start of the transaction corresponding to a slot                                                                                               |
| Output parameter | end_undo_ptr   | text    | Position where the undo record is inserted at the end of the transaction corresponding to a slot                                                                                                 |
| Output parameter | lsn            | text    | LSN of the modified slot                                                                                                                                                                         |
| Output parameter | slot_states    | oid     | Transaction state <ul style="list-style-type: none"><li>• <b>0</b>: committed</li><li>• <b>1</b>: being executed</li><li>• <b>2</b>: being rolled back</li><li>• <b>3</b>: rolled back</li></ul> |

- **gs\_undo\_dump\_record(undoptr bigint)**  
Description: Parses undo records.  
Return type: record

**Table 7-71** gs\_undo\_dump\_record parameters

| Category         | Parameter | Type   | Description                                     |
|------------------|-----------|--------|-------------------------------------------------|
| Input parameter  | undoptr   | bigint | Start position of the undo record to be parsed. |
| Output parameter | undoptr   | bigint | Start position of the undo record to be parsed. |

| Category         | Parameter    | Type | Description                                             |
|------------------|--------------|------|---------------------------------------------------------|
| Output parameter | xactid       | text | Transaction ID.                                         |
| Output parameter | cid          | text | Command ID.                                             |
| Output parameter | relid        | text | Relation OID.                                           |
| Output parameter | relfilenode  | text | Relfinode of the file.                                  |
| Output parameter | utype        | text | Undo record type                                        |
| Output parameter | blkprev      | text | Position of the previous undo record in the same block. |
| Output parameter | blockno      | text | Block number.                                           |
| Output parameter | uoffset      | text | Undo record offset.                                     |
| Output parameter | prevurp      | text | Position of the previous undo record.                   |
| Output parameter | payloadlen   | text | Length of the undo record data.                         |
| Output parameter | oldxactid    | text | Previous XID.                                           |
| Output parameter | partitionoid | text | Partition OID.                                          |
| Output parameter | tablespace   | text | Tablespace.                                             |

| Category         | Parameter         | Type | Description                              |
|------------------|-------------------|------|------------------------------------------|
| Output parameter | alreadyread_bytes | text | Length of the read undo record.          |
| Output parameter | prev_undorec_len  | text | Length of the previous undo record.      |
| Output parameter | td_id             | text | ID of the transaction directory.         |
| Output parameter | reserved          | text | Specifies whether to reserve the record. |
| Output parameter | flag              | text | Flag 1.                                  |
| Output parameter | flag2             | text | Flag 2.                                  |
| Output parameter | t_hoff            | text | Length of the undo record data header.   |

- gs\_undo\_dump\_xid(undo\_xid xid)  
Description: Parses undo records based on the XID.  
Return type: record

**Table 7-72** gs\_undo\_dump\_xid parameters

| Category         | Parameter | Type | Description                                    |
|------------------|-----------|------|------------------------------------------------|
| Input parameter  | undo_xid  | xid  | XID                                            |
| Output parameter | undoptr   | xid  | Start position of the undo record to be parsed |
| Output parameter | xactid    | text | XID                                            |

| Category         | Parameter         | Type | Description                                            |
|------------------|-------------------|------|--------------------------------------------------------|
| Output parameter | cid               | text | command id                                             |
| Output parameter | reloid            | text | relation oid                                           |
| Output parameter | relfilenode       | text | Relfinode of the file                                  |
| Output parameter | utype             | text | Undo record type                                       |
| Output parameter | blkprev           | text | Position of the previous undo record in the same block |
| Output parameter | blockno           | text | Block number                                           |
| Output parameter | uoffset           | text | Undo record offset                                     |
| Output parameter | prevurp           | text | Position of the previous undo record                   |
| Output parameter | payloadlen        | text | Length of the undo record data                         |
| Output parameter | oldxactid         | text | Previous XID                                           |
| Output parameter | partitionoid      | text | Partition OID                                          |
| Output parameter | tablespace        | text | Tablespace                                             |
| Output parameter | alreadyread_bytes | text | Length of the read undo record                         |

| Category         | Parameter         | Type | Description                              |
|------------------|-------------------|------|------------------------------------------|
| Output parameter | prev_undo_rec_len | text | Length of the previous undo record       |
| Output parameter | td_id             | text | ID of the transaction directory          |
| Output parameter | reserved          | text | Specifies whether to reserve the record. |
| Output parameter | flag              | text | Flag 1                                   |
| Output parameter | flag2             | text | Flag 2                                   |
| Output parameter | t_hoff            | text | Length of the undo record data header    |

- `gs_verify_undo_record(type, startIdx, endIdx, location)`

Description: Verifies the Undo record.

Return type: record

 **NOTE**

This API is reserved for extension and cannot be used.

- `gs_verify_undo_translot(type, startIdx, endIdx, location)`

Description: Verifies the Undo transaction slot.

Return type: record

 **NOTE**

This API is reserved for extension and cannot be used.

- `gs_verify_undo_meta(type, startIdx, endIdx, location)`

Description: Verifies the Undo meta information.

Return type: record

 **NOTE**

This API is reserved for extension and cannot be used.

## 7.5.24 Statistics Information Functions

Statistics information functions are divided into the following two categories: functions that access databases, using the OID of each table or index in a database to mark the database for which statistics are generated; functions that

access servers, identified by the server process ID, whose value ranges from 1 to the number of currently active servers.

- `pg_stat_get_db_conflict_tablespace(oid)`  
Description: Specifies the number of queries canceled due to a conflict between the restored tablespace and the deleted tablespace in the database.  
Return type: bigint
- `pg_control_group_config()`  
Description: Prints Cgroup configurations on the current node. Only users with the SYSADMIN permission can execute this function.  
Return type: record
- `pg_stat_get_db_stat_reset_time(oid)`  
Description: Specifies the most recent time when database statistics were reset. It is initialized to the system time during the first connection to each database. The reset time is updated when you call **pg\_stat\_reset** on the database and execute **pg\_stat\_reset\_single\_table\_counters** against any table or index in it.  
Return type: timestampz
- `pg_stat_get_function_total_time(oid)`  
Description: Specifies the total wall clock time spent in the function, in microseconds. The time spent on this function calling other functions is included.  
Return type: bigint
- `pg_stat_get_xact_tuples_returned(oid)`  
Description: Specifies the number of rows read through sequential scans when the parameter is a table in the current transaction or the number of index entries returned when the parameter is an index.  
Return type: bigint
- `pg_lock_status()`  
Description: Queries information about locks held by open transactions. All users can execute this function.  
Return type: For details, see the return result of **PG\_LOCKS**, which is obtained by querying this function.
- `pg_stat_get_xact_numscans(oid)`  
Description: Specifies the number of sequential scans performed when the parameter is a table in the current transaction or the number of index scans performed when the parameter is an index.  
Return type: bigint
- `pg_stat_get_xact_blocks_fetched(oid)`  
Description: Specifies the number of disk block fetch requests for a table or an index in the current transaction.  
Return type: bigint
- `pg_stat_get_xact_blocks_hit(oid)`  
Description: Specifies the number of disk block fetch requests for tables or indexes found in cache in the current transaction.  
Return type: bigint

- `pg_stat_get_xact_function_calls(oid)`  
Description: Specifies the number of times the function is called in the current transaction.  
Return type: `bigint`
- `pg_stat_get_xact_function_self_time(oid)`  
Description: Specifies the time spent on this function in the current transaction, excluding the time spent on this function internally calling other functions.  
Return type: `bigint`
- `pg_stat_get_xact_function_total_time(oid)`  
Description: Specifies the total wall clock time (in microseconds) spent on the function in the current transaction, including the time spent on this function internally calling other functions.  
Return type: `bigint`
- `pg_stat_get_wal_senders()`  
Description: Queries walsender information on the primary node.  
Return type: `setofrecord`  
The following table describes return columns.

**Table 7-73** Return column description

| Column                             | Type                     | Description                             |
|------------------------------------|--------------------------|-----------------------------------------|
| <code>pid</code>                   | <code>bigint</code>      | Thread ID of the WAL sender             |
| <code>sender_pid</code>            | <code>integer</code>     | Lightweight thread ID of the WAL sender |
| <code>local_role</code>            | <code>text</code>        | Type of the primary node                |
| <code>peer_role</code>             | <code>text</code>        | Type of the standby node                |
| <code>peer_state</code>            | <code>text</code>        | Status of the standby node              |
| <code>state</code>                 | <code>text</code>        | Status of the WAL sender                |
| <code>catchup_start</code>         | timestamp with time zone | Startup time of a catchup task          |
| <code>catchup_end</code>           | timestamp with time zone | End time of a catchup task              |
| <code>sender_sent_location</code>  | <code>text</code>        | Sending position of the primary node    |
| <code>sender_write_location</code> | <code>text</code>        | Writing position of the primary node    |

| Column                     | Type | Description                                        |
|----------------------------|------|----------------------------------------------------|
| sender_flush_location      | text | Flushing position of the primary node              |
| sender_replay_location     | text | Redo position of the primary node                  |
| receiver_received_location | text | Receiving position of the standby node             |
| receiver_write_location    | text | Writing position of the standby node               |
| receiver_flush_location    | text | Flushing position of the standby node              |
| receiver_replay_location   | text | Redo position of the standby node                  |
| sync_percent               | text | Synchronization percentage                         |
| sync_state                 | text | Synchronization status                             |
| sync_group                 | text | Group to which the synchronous replication belongs |
| sync_priority              | text | Priority of synchronous replication                |
| sync_most_available        | text | Maximum availability mode                          |
| channel                    | text | Channel information of the WAL sender              |

- `get_paxos_replication_info()`  
Description: Queries the primary/standby replication status in Paxos mode.  
Return type: setofrecord  
The following table describes return columns.

**Table 7-74** Return column description

| Column                | Type | Description                                                                            |
|-----------------------|------|----------------------------------------------------------------------------------------|
| paxos_write_location  | text | Position of the Xlog that has been written to the Distribute Consensus Framework (DCF) |
| paxos_commit_location | text | Position of the Xlog agreed in the DCF                                                 |



| Column                | Type | Description                      |
|-----------------------|------|----------------------------------|
| local_write_location  | text | Writing position of a node       |
| local_flush_location  | text | Flushing position of a node      |
| local_replay_location | text | Redo position of a node          |
| dcf_replication_info  | text | DCF module information of a node |

- `pg_stat_get_stream_replications()`  
Description: Queries the primary/standby replication status.  
Return type: setofrecord  
The following table describes return values.

**Table 7-75** Return values

| Return Parameter   | Type    | Description           |
|--------------------|---------|-----------------------|
| local_role         | text    | Local role            |
| static_connections | integer | Connection statistics |
| db_state           | text    | Database status       |
| detail_information | text    | Detailed information  |

- `pg_stat_get_db_numbackends(oid)`  
Description: Specifies the number of active server processes for a database.  
Return type: integer
- `pg_stat_get_db_xact_commit(oid)`  
Description: Specifies the number of transactions committed in a database.  
Return type: bigint
- `pg_stat_get_db_xact_rollback(oid)`  
Description: Specifies the number of transactions rolled back in a database.  
Return type: bigint
- `pg_stat_get_db_blocks_fetched(oid)`  
Description: Specifies the number of disk blocks fetch requests for a database.  
Return type: bigint
- `pg_stat_get_db_blocks_hit(oid)`  
Description: Specifies the number of disk block fetch requests found in cache for a database.  
Return type: bigint

- `pg_stat_get_db_tuples_returned(oid)`  
Description: Specifies the number of tuples returned for a database.  
Return type: bigint
- `pg_stat_get_db_tuples_fetched(oid)`  
Description: Specifies the number of tuples fetched for a database.  
Return type: bigint
- `pg_stat_get_db_tuples_inserted(oid)`  
Description: Specifies the number of tuples inserted in a database.  
Return type: bigint
- `pg_stat_get_db_tuples_updated(oid)`  
Description: Specifies the number of tuples updated in a database.  
Return type: bigint
- `pg_stat_get_db_tuples_deleted(oid)`  
Description: Specifies the number of tuples deleted in a database.  
Return type: bigint
- `pg_stat_get_db_conflict_lock(oid)`  
Description: Specifies the number of lock conflicts in a database.  
Return type: bigint
- `pg_stat_get_db_deadlocks(oid)`  
Description: Specifies the number of deadlocks in a database.  
Return type: bigint
- `pg_stat_get_numscans(oid)`  
Description: Specifies the number of sequential row scans done if parameters are in a table or the number of index scans done if parameters are in an index.  
Return type: bigint
- `pg_stat_get_role_name(oid)`  
Description: Obtains the username based on the user OID. Only users with the SYSADMIN or MONADMIN permission can access the information.  
Return type: text  
Example:

```
openGauss=# select pg_stat_get_role_name(10);
 pg_stat_get_role_name

 aabbcc
(1 row)
```
- `pg_stat_get_tuples_returned(oid)`  
Description: Specifies the number of sequential row scans done if parameters are in a table or the number of index scans done if parameters are in an index.  
Return type: bigint
- `pg_stat_get_tuples_fetched(oid)`  
Description: Specifies the number of table rows fetched by bitmap scans if parameters are in a table or the number of table rows fetched by simple index scans using the index if parameters are in an index.

- Return type: bigint
- `pg_stat_get_tuples_inserted(oid)`  
Description: Specifies the number of rows inserted into a table.  
Return type: bigint
  - `pg_stat_get_tuples_updated(oid)`  
Description: Specifies the number of rows updated in a table.  
Return type: bigint
  - `pg_stat_get_tuples_deleted(oid)`  
Description: Specifies the number of rows deleted from a table.  
Return type: bigint
  - `pg_stat_get_tuples_changed(oid)`  
Description: Specifies the total number of inserted, updated, and deleted rows after a table was last analyzed or autoanalyzed.  
Return type: bigint
  - `pg_stat_get_tuples_hot_updated(oid)`  
Description: Specifies the number of rows hot updated in a table.  
Return type: bigint
  - `pg_stat_get_live_tuples(oid)`  
Description: Specifies the number of live rows in a table.  
Return type: bigint
  - `pg_stat_get_dead_tuples(oid)`  
Description: Specifies the number of dead rows in a table.  
Return type: bigint
  - `pg_stat_get_blocks_fetched(oid)`  
Description: Specifies the number of disk block fetch requests for a table or an index.  
Return type: bigint
  - `pg_stat_get_blocks_hit(oid)`  
Description: Specifies the number of disk block requests found in cache for a table or an index.  
Return type: bigint
  - `pg_stat_get_partition_tuples_inserted(oid)`  
Description: Specifies the number of rows in the corresponding table partition.  
Return type: bigint
  - `pg_stat_get_partition_tuples_updated(oid)`  
Description: Specifies the number of rows that have been updated in the corresponding table partition.  
Return type: bigint
  - `pg_stat_get_partition_tuples_deleted(oid)`  
Description: Specifies the number of rows deleted from the corresponding table partition.  
Return type: bigint

- `pg_stat_get_partition_tuples_changed(oid)`  
Description: Specifies the total number of inserted, updated, and deleted rows after a table partition was last analyzed or autoanalyzed.  
Return type: `bigint`
- `pg_stat_get_partition_live_tuples(oid)`  
Description: Specifies the number of live rows in a partitioned table.  
Return type: `bigint`
- `pg_stat_get_partition_dead_tuples(oid)`  
Description: Specifies the number of dead rows in a partitioned table.  
Return type: `bigint`
- `pg_stat_get_xact_tuples_fetched(oid)`  
Description: Specifies the number of tuple rows scanned in a transaction.  
Return type: `bigint`
- `pg_stat_get_xact_tuples_inserted(oid)`  
Description: Specifies the number of tuple inserted into the active subtransactions related to a table.  
Return type: `bigint`
- `pg_stat_get_xact_tuples_deleted(oid)`  
Description: Specifies the number of deleted tuples in the active subtransactions related to a table.  
Return type: `bigint`
- `pg_stat_get_xact_tuples_hot_updated(oid)`  
Description: Specifies the number of hot updated tuples in the active subtransactions related to a table.  
Return type: `bigint`
- `pg_stat_get_xact_tuples_updated(oid)`  
Description: Specifies the number of updated tuples in the active subtransactions related to a table.  
Return type: `bigint`
- `pg_stat_get_xact_partition_tuples_inserted(oid)`  
Description: Specifies the number of inserted tuples in the active subtransactions related to a table partition.  
Return type: `bigint`
- `pg_stat_get_xact_partition_tuples_deleted(oid)`  
Description: Specifies the number of deleted tuples in the active subtransactions related to a table partition.  
Return type: `bigint`
- `pg_stat_get_xact_partition_tuples_hot_updated(oid)`  
Description: Specifies the number of hot updated tuples in the active subtransactions related to a table partition.  
Return type: `bigint`
- `pg_stat_get_xact_partition_tuples_updated(oid)`

Description: Specifies the number of updated tuples in the active subtransactions related to a table partition.

Return type: bigint

- `pg_stat_get_last_vacuum_time(oid)`

Description: Specifies the most recent time when the autovacuum thread is manually started to clear a table.

Return type: timestampz

- `pg_stat_get_last_autovacuum_time(oid)`

Description: Specifies the time of the last vacuum initiated by the autovacuum daemon on a table.

Return type: timestampz

- `pg_stat_get_vacuum_count(oid)`

Description: Specifies the number of times a table is manually cleared.

Return type: bigint

- `pg_stat_get_autovacuum_count(oid)`

Description: Specifies the number of times the autovacuum daemon is started to clear a table.

Return type: bigint

- `pg_stat_get_last_analyze_time(oid)`

Description: Specifies the last time when a table starts to be analyzed manually or by the autovacuum thread.

Return type: timestampz

- `pg_stat_get_last_autoanalyze_time(oid)`

Description: Specifies the time when the last analysis initiated by the autovacuum daemon on a table.

Return type: timestampz

- `pg_stat_get_analyze_count(oid)`

Description: Specifies the number of times a table is manually analyzed.

Return type: bigint

- `pg_stat_get_autoanalyze_count(oid)`

Description: Specifies the number of times the autovacuum daemon analyzes a table.

Return type: bigint

- `pg_total_autovac_tuples(bool)`

Description: Returns tuple records related to the total autovac, such as **nodename**, **nspname**, **relname**, and tuple IUDs. The input parameter specifies whether to query the **relation** information.

Return type: setofrecord

The following table describes return parameters.

**Table 7-76** Description

| Return Parameter      | Type   | Description                                             |
|-----------------------|--------|---------------------------------------------------------|
| nodename              | name   | Node name                                               |
| nspname               | name   | Name of a namespace                                     |
| relname               | name   | Name of an object, such as a table, an index, or a view |
| partname              | name   | Partition name                                          |
| n_dead_tuples         | bigint | Number of dead rows in a table partition                |
| n_live_tuples         | bigint | Number of live rows in a table partition                |
| changes_since_analyze | bigint | Number of changes generated by ANALYZE                  |

- pg\_autovac\_status(oid)

Description: Returns autovac information, such as **nodename**, **nspname**, **relname**, **analyze**, **vacuum**, thresholds of **analyze** and **vacuum**, and the number of analyzed or vacuumed tuples. Only users with the SYSADMIN permission can use this function.

Return type: setofrecord

The following table describes return parameters.

**Table 7-77** Return parameter description

| Return Parameter | Type    | Description                                             |
|------------------|---------|---------------------------------------------------------|
| nspname          | text    | Name of a namespace                                     |
| relname          | text    | Name of an object, such as a table, an index, or a view |
| nodename         | text    | Node name                                               |
| doanalyze        | Boolean | Specifies whether to execute ANALYZE.                   |
| anltuples        | bigint  | Number of ANALYZE tuples                                |
| anlthresh        | bigint  | ANALYZE threshold                                       |
| dovacuum         | Boolean | Specifies whether to execute VACUUM.                    |
| vactuples        | bigint  | Number of VACUUM tuples                                 |

| Return Parameter | Type   | Description      |
|------------------|--------|------------------|
| vacthresh        | bigint | VACUUM threshold |

- pg\_autovac\_timeout(oid)

Description: Returns the number of consecutive timeouts during the autovac operation on a table. If the table information is invalid or the node information is abnormal, **NULL** will be returned.

Return type: bigint

- pg\_stat\_get\_last\_data\_changed\_time(oid)

Description: Returns the time when **INSERT**, **UPDATE**, **DELETE**, or **EXCHANGE/TRUNCATE/DROP PARTITION** was last performed on a table. The data in the **last\_data\_changed** column of the [PG\\_STAT\\_ALL\\_TABLES](#) view is calculated by using this function. The performance of obtaining the last modification time by using the view is poor when the table has a large amount of data. In this case, you are advised to use the function.

Return type: timestamptz

- pg\_stat\_set\_last\_data\_changed\_time(oid)

Description: Manually changes the time when **INSERT**, **UPDATE**, **DELETE**, or **EXCHANGE/TRUNCATE/DROP PARTITION** was last performed.

Return type: void

- pg\_backend\_pid()

Description: Specifies the thread ID of the server thread attached to the current session.

Return type: integer

- pg\_stat\_get\_activity(integer)

Description: Returns a record about the backend with the specified PID. A record for each active backend in the system is returned if **NULL** is specified. The returned result does not contain the **connection\_info** column. The initial user, system administrators and users with the MONADMIN permission can view all data. Common users can only query their own results.

Example:

```
openGauss=# select * from pg_stat_get_activity(139881386280704);
 datid | pid | sessionid | usesysid | application_name | state |
 query | waiting | xact_start | | query_start |
 backend_start | state_change | client_addr | client_hostname | client_port | enqueue
 | query_id | srespool | global_sessionid | unique_sql_id | trace_id
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
16545 | 139881386280704 | 69 | 10 | gsql | active | select * from
pg_stat_get_activity(139881386280704); | f | 2022-01-18 19:43:05.167718+08 | 2022-01-18
19:43:05.167718+08 | 2022
-01-18 19:42:33.513507+08 | 2022-01-18 19:43:05.16773+08 | | | -1 | |
72620543991624410 | default_pool | 1938253334#69#0 | 3751941862 |
(1 row)
```

Return type: setofrecord

The following table describes return parameters.

**Table 7-78** Description

| Return Parameter | Type                     | Description                                                                                                                                                                                                               |
|------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid            | oid                      | OID of the database that the user session connects to in the backend                                                                                                                                                      |
| pid              | bigint                   | Backend thread ID                                                                                                                                                                                                         |
| sessionid        | bigint                   | Session ID                                                                                                                                                                                                                |
| usesysid         | oid                      | OID of the user logged in to the backend                                                                                                                                                                                  |
| application_name | text                     | Name of the application connected to the backend                                                                                                                                                                          |
| state            | text                     | Overall status of the backend                                                                                                                                                                                             |
| query            | text                     | Latest query at the backend. If <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed.                                                      |
| waiting          | Boolean                  | Specifies whether the backend is currently waiting for a lock. If the backend is currently waiting for a lock, the value is <b>true</b> .                                                                                 |
| xact_start       | timestamp with time zone | Time when current transaction was started (null if no transaction is active).<br><br>If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column. |
| query_start      | timestamp with time zone | Time when the currently active query was started, or time when the last query was started if <b>state</b> is not <b>active</b>                                                                                            |



| Return Parameter | Type                     | Description                                                                                                                                                                                                                |
|------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| backend_start    | timestamp with time zone | Time when this process was started, that is, when the client connected to the server                                                                                                                                       |
| state_change     | timestamp with time zone | Time when <b>state</b> was last modified                                                                                                                                                                                   |
| client_addr      | inet                     | IP address of the client connected to the backend. If this column is <b>NULL</b> , it indicates either the client is connected via a Unix socket on the server or this is an internal process, such as <b>AUTOVACUUM</b> . |
| client_hostname  | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                          |
| client_port      | integer                  | TCP port number that the client uses for communication with the backend ( <b>-1</b> if a Unix socket is used)                                                                                                              |
| enqueue          | text                     | Unsupported currently                                                                                                                                                                                                      |
| query_id         | bigint                   | ID of a query                                                                                                                                                                                                              |
| srespool         | name                     | Name of the resource pool                                                                                                                                                                                                  |
| global_sessionid | text                     | Global session ID                                                                                                                                                                                                          |
| unique_sql_id    | bigint                   | Unique SQL statement ID                                                                                                                                                                                                    |
| trace_id         | text                     | Driver-specific trace ID, which is associated with an application request                                                                                                                                                  |

- pg\_stat\_get\_activity\_with\_conninfo(integer)

Description: Returns a record about the background process with the specified PID. A record for each active background process in the system is returned if **NULL** is specified. The initial user, system administrators and users with the MONADMIN permission can view all data. Common users can only query their own results.

Return type: setofrecord

The following table describes return values.

**Table 7-79** Return values

| Return Value     | Return Type | Description                                                                                                                                                          |
|------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid            | oid         | OID of the database that the user session connects to in the backend                                                                                                 |
| pid              | bigint      | Backend thread ID                                                                                                                                                    |
| sessionid        | bigint      | Session ID                                                                                                                                                           |
| usesysid         | oid         | OID of the user logged in to the backend                                                                                                                             |
| application_name | text        | Name of the application connected to the backend                                                                                                                     |
| state            | text        | Overall status of the backend.                                                                                                                                       |
| query            | text        | Latest query at the backend. If <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed. |
| waiting          | Boolean     | Specifies whether the backend is currently waiting for a lock. If the backend is currently waiting for a lock, the value is <b>true</b> .                            |

| Return Value    | Return Type              | Description                                                                                                                                                                                                        |
|-----------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| xact_start      | timestamp with time zone | Time when current transaction was started (null if no transaction is active). If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column. |
| query_start     | timestamp with time zone | Time when the currently active query was started, or time when the last query was started if <b>state</b> is not <b>active</b>                                                                                     |
| backend_start   | timestamp with time zone | Time when this process was started, that is, when the client connected to the server                                                                                                                               |
| state_change    | timestamp with time zone | Time when <b>state</b> was last modified                                                                                                                                                                           |
| client_addr     | inet                     | IP address of the client connected to the backend. If this column is <b>null</b> , it indicates either the client is connected via a UDS on the server or this is an internal process, such as autovacuum.         |
| client_hostname | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                  |
| client_port     | integer                  | TCP port number that the client uses for communication with the backend (–1 if a Unix socket is used)                                                                                                              |
| enqueue         | text                     | Unsupported currently                                                                                                                                                                                              |

| Return Value     | Return Type | Description                                                                                                                            |
|------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------|
| query_id         | bigint      | ID of a query                                                                                                                          |
| connection_info  | text        | A string in JSON format recording the driver type, driver version, driver deployment path, and process owner of the connected database |
| srespool         | name        | Name of the resource pool                                                                                                              |
| global_sessionid | text        | Global session ID                                                                                                                      |
| unique_sql_id    | bigint      | Unique SQL statement ID                                                                                                                |
| trace_id         | text        | Driver-specific trace ID, which is associated with an application request.                                                             |

- `pg_stat_get_function_calls(oid)`  
Description: Specifies the number of times the function has been called.  
Return type: `bigint`
- `pg_stat_get_function_self_time(oid)`  
Description: Specifies the time spent in only this function. The time spent on this function calling other functions is excluded.  
Return type: `bigint`
- `pg_stat_get_backend_idset()`  
Description: Sets the number of currently active server processes (from 1 to the number of active server processes).  
Return type: `setofinteger`
- `pg_stat_get_backend_pid(integer)`  
Description: Specifies the ID of a specified server thread.  
Return type: `bigint`
- `pg_stat_get_backend_dbid(integer)`  
Description: Specifies the ID of a database connected to a specified server process.  
Return type: `oid`
- `pg_stat_get_backend_userid(integer)`  
Description: Specifies the user ID of a specified server process. This function can be called only by system administrators.  
Return type: `oid`
- `pg_stat_get_backend_activity(integer)`

Description: Queries the current activity of a specified server process. The query result can be obtained only when the user who calls this function is a system administrator or the same user as that of the session being queried and **track\_activities** is enabled.

Return type: text

- `pg_stat_get_backend_waiting(integer)`

Description: Returns a true value if a specified server process is waiting for a lock, but only when the user who calls this function is a system administrator or the same user as that of the session being queried and **track\_activities** is enabled.

Return type: Boolean

- `pg_stat_get_backend_activity_start(integer)`

Description: Specifies the time when a specified server process's currently executing query is started only when the user who calls this function is a system administrator or the same user as that of the session being queried and **track\_activities** is enabled.

Return type: timestamp with time zone

- `pg_stat_get_backend_xact_start(integer)`

Description: Specifies the time when a specified server process's currently executing transaction is started only when the user who calls this function is a system administrator or the same user as that of the session being queried and **track\_activities** is enabled.

Return type: timestamp with time zone

- `pg_stat_get_backend_start(integer)`

Description: Specifies the time when a specified server process is started. If the current user is neither a system administrator nor the same user as that of the session being queried, **NULL** is returned.

Return type: timestamp with time zone

- `pg_stat_get_backend_client_addr(integer)`

Description: Specifies the IP address of a backend connected to a specified client. If the connection is over a UDS, or if the current user is neither a system administrator nor the same user as that of the session being queried, **NULL** will be returned.

Return type: inet

- `pg_stat_get_backend_client_port(integer)`

Description: Specifies the TCP port of a backend connected to a specified client. If the connection is over a UDS, **-1** will be returned. If the current user is neither a system administrator nor the same user as that of the session being queried, **NULL** will be returned.

Return type: integer

- `pg_stat_get_bgwriter_timed_checkpoints()`

Description: Specifies the time when the background writer starts scheduled checkpoints (because the **checkpoint\_timeout** time has expired).

Return type: bigint

- `pg_stat_get_bgwriter_requested_checkpoints()`

Description: Specifies the time when the background writer starts checkpoints based on requests from the backend because **checkpoint\_segments** has been exceeded or the **CHECKPOINT** command has been executed.

Return type: bigint

- `pg_stat_get_bgwriter_buf_written_checkpoints()`

Description: Specifies the number of buffers written by the background writer during checkpoints.

Return type: bigint

- `pg_stat_get_bgwriter_buf_written_clean()`

Description: Specifies the number of buffers written by the background writer for routine cleaning of dirty pages.

Return type: bigint

- `pg_stat_get_bgwriter_maxwritten_clean()`

Description: Specifies the time when the background writer stops its cleaning scan because it has written more buffers than specified in the **bgwriter\_lru\_maxpages** parameter.

Return type: bigint

- `pg_stat_get_buf_written_backend()`

Description: Specifies the number of buffers written by the backend because they need to allocate a new buffer.

Return type: bigint

- `pg_stat_get_buf_alloc()`

Description: Specifies the total number of the allocated buffers.

Return type: bigint

- `pg_stat_clear_snapshot()`

Description: Clears the current statistics snapshot. This function can be executed only by SYSADMIN and MONADMIN.

Return type: void

- `pg_stat_reset()`

Description: Resets all statistics counters for the current database to zero (the SYSADMIN permission is required).

Return type: void

- `pg_stat_reset_shared(text)`

Description: Resets all statistics counters for the current database in each node in a shared cluster to zero (the SYSADMIN permission is required).

Return type: void

- `pg_stat_reset_single_table_counters(oid)`

Description: Resets statistics for a single table or index in the current database to zero (the SYSADMIN permission is required).

Return type: void

- `pg_stat_reset_single_function_counters(oid)`

Description: Resets statistics for a single function in the current database to zero (the SYSADMIN permission is required).

Return type: void

- **fenced\_udf\_process(integer)**  
Description: Shows the number of local UDF Master and Work processes. If the input parameter is set to **1**, the number of master processes is queried. If the input parameter is set to **2**, the number of worker processes is queried. If the input parameter is set to **3**, all worker processes are terminated.  
Return type: text
- **total\_cpu()**  
Description: Obtains the CPU time used by the current node, in jiffies.  
Return type: bigint
- **total\_memory()**  
Description: Obtains the size of the virtual memory used by the current node, in KB.  
Return type: bigint
- **pg\_stat\_get\_db\_cu\_hdd\_asyn(oid)**  
Description: Obtains the times CU is asynchronously read from a disk by a database of the current node.  
Return type: bigint
- **pg\_stat\_bad\_block(text, int, int, int, int, int, timestamp with time zone, timestamp with time zone)**  
Description: Obtains damage information about pages or CUs after the current node is started.  
Example:  

```
select * from pg_stat_bad_block();
```

  
Return type: record
- **pg\_stat\_bad\_block\_clear()**  
Description: Deletes the page and CU damage information that is read and recorded on the node (the SYSADMIN permission is required).  
Return type: void
- **gs\_respool\_exception\_info(pool text)**  
Description: Queries the query rule of a specified resource pool.  
Return type: record
- **gs\_control\_group\_info(pool text)**  
Description: Queries information about Cgroups associated with a resource pool.  
Return type: record

The command output is as follows:

| Attribute | Value               | Description                   |
|-----------|---------------------|-------------------------------|
| name      | class_a:workload_a1 | Class name and workload name. |
| class     | class_a             | Class Cgroup name.            |
| workload  | workload_a1         | Workload Cgroup name.         |

| Attribute | Value | Description                                                                               |
|-----------|-------|-------------------------------------------------------------------------------------------|
| type      | DEFWD | Cgroup type ( <b>Top</b> , <b>CLASS</b> , <b>BAKWD</b> , <b>DEFWD</b> , or <b>TSWD</b> ). |
| gid       | 87    | Cgroup ID.                                                                                |
| shares    | 30    | Percentage of CPU resources to those on the parent node.                                  |
| limits    | 0     | Percentage of CPU cores to those on the parent node.                                      |
| rate      | 0     | Allocation ratio in Timeshare.                                                            |
| cpucores  | 0-3   | Number of CPU cores.                                                                      |

- `gs_all_control_group_info()`  
Description: Collects information about all Cgroups in the database.  
Return type: record
- `gs_get_control_group_info()`  
Description: Collects information about all Cgroups.  
Return type: record
- `get_instr_workload_info(integer)`  
Description: Obtains the transaction volume and time information on the primary database node.  
Return type: record

| Attribute         | Value  | Description                                                     |
|-------------------|--------|-----------------------------------------------------------------|
| resourcepool_oid  | 10     | OID of the resource pool (the logic is equivalent to the load). |
| commit_counter    | 4      | Number of frontend transactions that were committed.            |
| rollback_counter  | 1      | Number of frontend transactions that were rolled back.          |
| resp_min          | 949    | Minimum response time of frontend transactions (unit: $\mu$ s). |
| resp_max          | 201891 | Maximum response time of frontend transactions (unit: $\mu$ s). |
| resp_avg          | 43564  | Average response time of frontend transactions (unit: $\mu$ s). |
| resp_total        | 217822 | Total response time of frontend transactions (unit: $\mu$ s).   |
| bg_commit_counter | 910    | Number of backend transactions that were committed.             |



| Attribute               | Value        | Description                                                    |
|-------------------------|--------------|----------------------------------------------------------------|
| bg_rollback_count<br>er | 0            | Number of backend transactions that were rolled back.          |
| bg_resp_min             | 97           | Minimum response time of backend transactions (unit: $\mu$ s). |
| bg_resp_max             | 678080687    | Maximum response time of backend transactions (unit: $\mu$ s). |
| bg_resp_avg             | 327847884    | Average response time of backend transactions (unit: $\mu$ s). |
| bg_resp_total           | 298341575300 | Total response time of backend transactions (unit: $\mu$ s).   |

- `pv_instance_time()`

Description: Obtains the time consumed in each execution phase on the current node.

Return type: record

| Stat_name<br>Attribute  | Value   | Description                                                                |
|-------------------------|---------|----------------------------------------------------------------------------|
| DB_TIME                 | 1062385 | Total end-to-end wall time consumed by all threads (unit: $\mu$ s).        |
| CPU_TIME                | 311777  | Total CPU time consumed by all threads (unit: $\mu$ s).                    |
| EXECUTION_TIME          | 380037  | Total time consumed on the executor (unit: $\mu$ s).                       |
| PARSE_TIME              | 6033    | Total time consumed for parsing SQL statements (unit: $\mu$ s).            |
| PLAN_TIME               | 173356  | Total time consumed for generating an execution plan (unit: $\mu$ s).      |
| REWRITE_TIME            | 2274    | Total time consumed on query rewriting (unit: $\mu$ s).                    |
| PL_EXECUTION_TI<br>ME   | 0       | Total time consumed for executing PL/<br>pgSQL statements (unit: $\mu$ s). |
| PL_COMPILATION<br>_TIME | 557     | Total time consumed for compiling SQL<br>statements (unit: $\mu$ s).       |
| NET_SEND_TIME           | 1673    | Total time consumed for sending data<br>over network (unit: $\mu$ s).      |
| DATA_IO_TIME            | 426622  | Total time consumed for data read and<br>write (unit: $\mu$ s).            |

- **DBE\_PERF.get\_global\_instance\_time()**  
Description: Provides the time consumed in each key phase in the entire database. To query this function, you must have the MONADMIN permission.  
Return type: record
- **get\_instr\_unique\_sql()**  
Description: Obtains information about execution statements (normalized SQL statements) on the current node as a user with the SYSADMIN permission.  
Return type: record
- **reset\_unique\_sql(text, text, bigint)**  
Description: Resets information about system execution statements (normalized SQL statements) information as a user with the SYSADMIN permission. The value of the first parameter can be **global** or **local**. **global** indicates that information on all nodes is cleared, and **local** indicates that only information on the current node is cleared. The value of the second parameter can be **ALL**, **BY\_USERID**, or **BY\_CNID**. **ALL** indicates that all information is cleared. **BY\_USERID** indicates that the SQL information of the user specified by **USERID** is cleared. **BY\_CNID** indicates that the SQL information related to the primary node of the database in the system is cleared. The third parameter indicates **CNID** and **USERID**. If the second parameter is set to **ALL**, the third parameter does not take effect and can be set to any value.  
Return type: Boolean

 **NOTE**

In the standalone system, the function of the value **global** is the same as that of the value **local** and the second parameter cannot set to be **BY\_CNID**.

- **get\_instr\_wait\_event(NULL)**  
Description: Obtains the statistics on wait events of the current node.  
Return type: record
- **get\_instr\_user\_login()**  
Description: Obtains the number of user login and logout times on the current node. Only users with the SYSADMIN or MONADMIN permission can execute this function.  
Return type: record
- **get\_instr\_rt\_percentile(integer)**  
Description: Obtains the distribution information about the response time of 80% and 95% SQL statements of the database.  
Return type: record
- **get\_node\_stat\_reset\_time()**  
Description: Obtains statistics about reset (restart, primary/standby switchover, and database deletion) time of the current node.  
Return type: record
- **DBE\_PERF.get\_global\_os\_runtime()**  
Description: Displays the running status of the current OS. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_os\_threads()

Description: Provides information about the threads under all normal nodes of the entire database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_workload\_sql\_count()

Description: Provides statistics about the number of SELECT, UPDATE, INSERT, DELETE, DDL, DML, and DCL statements of different service loads in the entire database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_workload\_sql\_elapse\_time()

Description: Provides statistics about the number of SELECT, UPDATE, INSERT, and DELETE statements and response time information (TOTAL, AVG, MIN, and MAX) for different loads in the entire database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_workload\_transaction()

Description: Obtains the transaction volume and time information on all nodes of the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_session\_stat()

Description: Obtains the session status information on all nodes of the database. To query this function, you must have the MONADMIN permission.

Return type: record

#### NOTE

There are 17 status information items: **commit**, **rollback**, **sql**, **table\_scan**, **blocks\_fetched**, **physical\_read\_operation**, **shared\_blocks\_dirtied**, **local\_blocks\_dirtied**, **shared\_blocks\_read**, **local\_blocks\_read**, **blocks\_read\_time**, **blocks\_write\_time**, **sort\_imemory**, **sort\_idisk**, **cu\_mem\_hit**, **cu\_hdd\_sync\_read**, and **cu\_hdd\_asyread**.

- DBE\_PERF.get\_global\_session\_time()

Description: Provides the time consumed in each key phase of each node in the entire database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_session\_memory()

Description: Displays statistics about memory usage at the session level on each node in the unit of MB, including all the memory allocated to GaussDB and stream threads on DN for jobs currently executed by users. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_session\_memory\_detail()

Description: Displays statistics about thread memory usage on each node by MemoryContext node. To query this function, you must have the MONADMIN permission.

Return type: record

- `create_wlm_session_info(int flag)`

Description: Clears top SQL query statement-level statistics recorded in the current memory. Only the administrator can execute this function.

Return type: int

- `pg_stat_get_wlm_session_info(int flag)`

Description: Obtains top SQL query statement-level statistics recorded in the current memory. If the input parameter is not 0, the information is cleared from the memory. Only users with the SYSADMIN or MONADMIN permission can execute this function.

Return type: record

- `gs_paxos_stat_replication()`

Description: Queries the standby node information on the primary node. Currently, only the centralized DCF mode is supported.

Return type: setofrecord

The following table describes return columns.

| Column                                | Type | Description                                                               |
|---------------------------------------|------|---------------------------------------------------------------------------|
| <code>local_role</code>               | text | Role of the node that sends logs                                          |
| <code>peer_role</code>                | text | Role of the node that receives logs                                       |
| <code>local_dcf_role</code>           | text | DCF role of the node that sends logs                                      |
| <code>peer_dcf_role</code>            | text | DCF role of the node that receives logs                                   |
| <code>peer_state</code>               | text | Status of the node that receives logs                                     |
| <code>sender_write_location</code>    | text | Location in the Xlog buffer where the node that sends logs is written     |
| <code>sender_commit_location</code>   | text | Consistency point reached for the DCF logs of the node that sends logs    |
| <code>sender_flush_location</code>    | text | Location in the Xlog disk where the node that sends logs is written       |
| <code>sender_replay_location</code>   | text | Location where the node that sends logs replays logs                      |
| <code>receiver_write_location</code>  | text | Location in the Xlog buffer where the node that receives logs is written  |
| <code>receiver_commit_location</code> | text | Consistency point reached for the DCF logs of the node that receives logs |
| <code>receiver_flush_location</code>  | text | Location in the Xlog disk where the node that receives logs is written    |

|                          |      |                                                          |
|--------------------------|------|----------------------------------------------------------|
| receiver_replay_location | text | Location where the node that receives logs replays Xlogs |
| sync_percent             | text | Synchronization percentage                               |
| dcf_run_mode             | int4 | DCF synchronization mode                                 |
| channel                  | text | Channel information                                      |

- `gs_wlm_get_resource_pool_info(int)`  
Description: Obtains resource usage statistics of all users. The input parameter can be any value of the INT type or be **NULL**.  
Return type: record
- `gs_wlm_get_all_user_resource_info()`  
Description: Obtains resource usage statistics of all users.  
Return type: record
- `gs_wlm_get_user_info(int)`  
Description: Obtains information about all users. The input parameter is of the int type and can be any int value or **NULL**. Only users with the SYSADMIN permission can execute this function.  
Return type: record
- `gs_wlm_readjust_user_space()`  
Description: Corrects the storage space usage of all users. Only the administrator can execute this function.  
Return type: record
- `gs_wlm_readjust_user_space_through_username(text name)`  
Description: Corrects the storage space usage of a specified user. Common users can use this function to modify only their own usage. Only the administrator can modify the usage of all users. If the value of **name** is **0000**, the usage of all users needs to be modified.  
Return type: record
- `gs_wlm_readjust_user_space_with_reset_flag(text name, boolean isfirst)`  
Description: Corrects the storage space usage of a specified user. If the input parameter **isfirst** is set to **true**, statistics are collected from 0. Otherwise, statistics are collected from the previous result. Common users can use this function to modify only their own usage. Only the administrator can modify the usage of all users. If the value of **name** is **0000**, the usage of all users needs to be modified.  
Return type: record
- `gs_wlm_session_respool(bigint)`  
Description: Obtains the session resource pool information about all backend threads. The input parameter can be any value of the bigint type or can be null.  
Return type: record
- `gs_io_wait_status()`

Description: This API does not support single-node systems or centralized systems and is unavailable currently.

- `global_stat_get_hotkeys_info()`

Description: Obtains the statistics on hot keys in the entire database instance. This API does not support single-node systems or centralized systems and is unavailable currently.

- `global_stat_clean_hotkeys()`

Description: Clears statistics on hot keys in the entire database instance. This API does not support single-node systems or centralized systems and is unavailable currently.

- `DBE_PERF.get_global_session_stat_activity()`

Description: Displays information about threads that are running on each node in the database. To query this function, you must have the MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_thread_wait_status()`

Description: Displays the blocking status of backend threads and auxiliary threads on all nodes. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_operator_history_table()`

Description: Displays the operator-related records (persistent) generated after jobs are executed on the primary database node of the current user. To query this function, you must have the MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_operator_history()`

Description: Displays the operator-related records generated after jobs are executed on the primary database node of the current user. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_operator_runtime()`

Description: Displays real-time operator-related records of jobs executed on the primary database node of the current user. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_statement_complex_history()`

Description: Displays the historical records of complex queries on the primary database node of the current user. To query this function, you must have the MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_statement_complex_history_table()`

Description: Displays the historical records (persistent) of complex queries on the primary database node of the current user. To query this function, you must have the MONADMIN permission.

Return type: record

- `DBE_PERF.get_global_statement_complex_runtime()`  
Description: Displays the real-time information of complex queries on the primary database node of the current user. To query this function, you must have the SYSADMIN or MONADMIN permission.  
Return type: record
- `DBE_PERF.get_global_memory_node_detail()`  
Description: Displays the memory usage of a certain database on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_global_shared_memory_detail()`  
Description: Displays the usage information about all the shared memory contexts of all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_global_statio_all_indexes()`  
Description: Displays statistics about each index displayed in a row in the current database, showing I/O statistics about accesses to that specific index. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_summary_stat_all_tables()`  
Description: Displays statistics about a row in each table (including the TOAST table) in the aggregated data of each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_global_stat_all_tables()`  
Description: Displays statistics about a row in each table (including the TOAST table) of data on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_local_toastname_and_toastindexname()`  
Description: Provides the mapping between the name and index of the local TOAST table and its associated table. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_summary_statio_all_indexes()`  
Description: Collects statistics about each index displayed in a row in the current databases of all nodes and displays the I/O statistics of a specific index. To query this function, you must have the MONADMIN permission.  
Return type: record
- `DBE_PERF.get_global_statio_all_sequences()`  
Description: Provides I/O status information about all sequences in the namespace. To query this function, you must have the MONADMIN permission.  
Return type: record

- DBE\_PERF.get\_global\_statio\_all\_tables()  
Description: Displays the I/O statistics about each table in databases on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_statio\_all\_tables()  
Description: Collects I/O statistics about each table in databases in the database. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_local\_toast\_relation()  
Description: Provides the mapping between the name of the local TOAST table and its associated table. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_statio\_sys\_indexes()  
Description: Displays the I/O status information about all system catalog indexes in namespaces on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_statio\_sys\_indexes()  
Description: Collects the I/O status information about all system catalog indexes in namespaces on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_statio\_sys\_sequences()  
Description: Provides the I/O status information about all the system sequences in the namespace. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_statio\_sys\_tables()  
Description: Provides I/O status information about all system catalogs in namespaces on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_statio\_sys\_tables()  
Description: Displays the I/O status information of all system catalogs in the namespace in the database. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_statio\_user\_indexes()  
Description: Displays the I/O status information about all user relationship table indexes in namespaces on each node. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_statio\_user\_indexes()



Description: Displays the I/O status information about all user relationship table indexes in namespaces in the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_statio\_user\_sequences()

Description: Displays the I/O status information about all user sequences in the namespace of each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_statio\_user\_tables()

Description: Displays the I/O status information about all user relationship tables in namespaces on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_statio\_user\_tables()

Description: Displays the I/O status information about all user relationship tables in namespaces in the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_stat\_db\_cu()

Description: Queries CU hits in a database and in each node in the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_all\_indexes()

Description: Displays statistics of each index in databases on all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_all\_indexes()

Description: Collects statistics of each index in all databases on all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_sys\_tables()

Description: Displays statistics about the system catalogs of all the namespaces in **pg\_catalog** and **information\_schema** schemas on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_sys\_tables()

Description: Collects statistics about the system catalogs of all the namespaces in **pg\_catalog** and **information\_schema** schemas on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_sys\_indexes()

Description: Displays index status information about all the system catalogs in the **pg\_catalog** and **information\_schema** schemas on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_sys\_indexes()

Description: Collects statistics about index status information about all the system catalogs in the pg\_catalog and information\_schema schemas on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_user\_tables()

Description: Displays the status information about customized ordinary tables in all namespaces. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_user\_tables()

Description: Collects statistics about the status information about customized ordinary tables in all namespaces. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_user\_indexes()

Description: Displays the status information about the index of customized ordinary tables in all databases. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_user\_indexes()

Description: Collects statistics about the status information about the index of customized ordinary tables in all databases. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_database()

Description: Displays database statistics of all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_database\_conflicts()

Description: Collects statistics on the database of all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_xact\_all\_tables()

Description: Displays transaction status information about all ordinary tables and TOAST tables in all namespaces. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_stat\_xact\_all\_tables()

Description: Collects statistics about transaction status information about all ordinary tables and TOAST tables in all namespaces. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_stat\_xact\_sys\_tables()  
Description: Displays transaction status information about all system catalogs in namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_stat\_xact\_sys\_tables()  
Description: Collects statistics about transaction status information about all system catalogs in namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_stat\_xact\_user\_tables()  
Description: Displays the transaction status information of the user tables in the namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_summary\_stat\_xact\_user\_tables()  
Description: Collects statistics about the transaction status information of the user tables in the namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_stat\_user\_functions()  
Description: Displays the transaction status information of user-defined functions in the namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_stat\_xact\_user\_functions()  
Description: Collects statistics about the transaction status information of user-defined functions in the namespaces on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_stat\_bad\_block()  
Description: Displays information about table and index read failures on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_file\_redo\_iostat()  
Description: Collects statistics on information about table and index read failures on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_file\_iostat()  
Description: Displays statistics about data file I/Os on all nodes. To query this function, you must have the MONADMIN permission.  
Return type: record
- DBE\_PERF.get\_global\_locks()

Description: Displays lock information of all nodes. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_replication\_slots()

Description: Displays logical replication information on all nodes. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_bgwriter\_stat()

Description: Displays statistics about the background writer process's activities on all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_replication\_stat()

Description: Displays information about log synchronization status on each node, such as the locations where the sender sends logs and where the receiver receives logs. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_transactions\_running\_xacts()

Description: Displays information about running transactions on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_transactions\_running\_xacts()

Description: Collects statistics of information about running transactions on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_transactions\_prepared\_xacts()

Description: Displays information about transactions that are currently prepared for two-phase commit on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_transactions\_prepared\_xacts()

Description: Collects statistics information about transactions that are currently prepared for two-phase commit on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_statement()

Description: Displays the status information of the historically-executed statements on each node. To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_statement\_count()

Description: Displays the number of SELECT, UPDATE, INSERT, and DELETE statements and response time information (TOTAL, AVG, MIN, and MAX) on

each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_config\_settings()

Description: Displays GUC parameter configuration information on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_wait\_events()

Description: Displays the wait event status information on each node. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_statement\_responsetime\_percentile()

Description: Obtains the response time distribution for 80% and 95% SQL statements of the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_summary\_user\_login()

Description: Collects statistics about number of user login and logout times on each node in the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.get\_global\_record\_reset\_time()

Description: Displays the statistics about reset (restart, primary/standby switchover, and database deletion) time of the database. To query this function, you must have the MONADMIN permission.

Return type: record

- DBE\_PERF.track\_memory\_context(context\_list text)

Description: Sets the memory context whose memory allocation details need to be collected. The input parameter is the memory context names, which are separated by commas (,), for example, **ThreadTopMemoryContext**, **SessionCacheMemoryContext**. Note that the memory context names are context-sensitive. In addition, the length of a single memory context is 63, and the excess part is truncated. The maximum number of memory contexts that can be collected at a time is 16. If the number of memory contexts exceeds 16, the setting fails. Each time this function is called, the previous collection result is cleared. When the input parameter is set to "", the collection function is disabled. Only the initial user or a user with the MONADMIN permission can execute this function.

Return type: Boolean

- DBE\_PERF.track\_memory\_context\_detail()

Description: Obtains the memory allocation details of the memory context specified by the **DBE\_PERF.track\_memory\_context** function. For details, see the **DBE\_PERF.track\_memory\_context\_detail** view. Only the initial user or a user with the MONADMIN permission can execute this function.

Return type: record

- pg\_stat\_get\_mem\_mbytes\_reserved(tid)

Description: Collects statistics on variables related to resource management, which is used only for fault locating.

Parameter: thread ID

Return type: text

- `gs_wlm_user_resource_info(name text)`

Description: Queries a user's resource quota and resource usage.

Return type: record

- `pg_stat_get_file_stat()`

Description: Collects statistics on data file I/Os to indicate I/O performance and detect performance problems such as abnormal I/O operations.

Return type: record

- `pg_stat_get_redo_stat()`

Description: Displays statistics on the replay of session thread logs.

Return type: record

- `pg_stat_get_status(int8)`

Description: Tests the blocking status about the backend thread and auxiliary thread of the current instance.

Return type: record

- `get_local_rel_iostat()`

Description: Queries the accumulated I/O status of data files on the current node.

Return type: record

- `DBE_PERF.get_global_rel_iostat()`

Description: Displays statistics about data file I/Os on all nodes. To query this function, you must have the MONADMIN permission.

Return type: record

- `DBE_PERF.global_threadpool_status()`

Description: Displays the status of worker threads and sessions in thread pools on all nodes. For details about the columns returned by the function, see [GLOBAL\\_THREADPOOL\\_STATUS](#).

Return type: record

- `remote_bgwriter_stat()`

Description: Displays the information about pages flushed by the bgwriter threads of all instances in the database, number of pages in the candidate buffer chain, and buffer elimination information. (The query result does not contain the information about the current node and cannot be used on DNs.) This is not supported in centralized mode.

Return type: record

- `pv_os_run_info()`

Description: Displays the running status of the current OS. For details about the columns, see [GS\\_OS\\_RUN\\_INFO](#).

Parameter: nan

Return type: SETOF record

- `pv_session_stat()`  
Description: Collects session status information by session thread or autovacuum thread. For details about the columns, see [GS\\_SESSION\\_STAT](#).  
Parameter: **nan**  
Return type: SETOF record
- `pv_session_time()`  
Description: Collects statistics on the running time of session threads and the time consumed in each execution phase. For details about the columns, see [GS\\_SESSION\\_TIME](#).  
Parameter: **nan**  
Return type: SETOF record
- `pg_stat_get_db_temp_bytes()`  
Description: Collects statistics on the total amount of data written to temporary files through database query. All temporary files are counted, regardless of why the temporary file was created, and regardless of the **log\_temp\_files** setting.  
Parameter: **oid**  
Return type: bigint
- `pg_stat_get_db_temp_files()`  
Description: Queries the number of temporary files created in the database. All temporary files are counted, regardless of why the temporary file was created (for example, sorting or hashing), and regardless of the **log\_temp\_files** setting.  
Parameter: **oid**  
Return type: bigint
- `create_wlm_instance_statistics_info`  
Description: Saves the historical monitoring data of the current instance persistently.  
Parameter: **nan**  
Return type: integer
- `remote_candidate_stat()`  
Description: Displays the checkpoint information and log flushing information about all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record
- `remote_ckpt_stat()`  
Description: Displays the checkpoint information and log flushing information about all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record
- `remote_single_flush_dw_stat()`  
Description: Displays the single-page doublewrite file status of all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record

- `remote_double_write_stat()`  
Description: Displays doublewrite file status of all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record
- `remote_pagewriter_stat()`  
Description: Displays the page flushing information and checkpoint information about all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record
- `remote_recovery_status()`  
Description: Displays log flow control information about the primary and standby nodes (except the current node). Centralized systems are not supported.  
Return type: record
- `remote_redo_stat()`  
Description: Displays the log replay status of all instances in the database (except the current node). Centralized systems are not supported.  
Return type: record
- `dbe_perf.gs_stat_activity_timeout(int)`  
Description: Obtains information about query jobs whose execution time exceeds the timeout threshold on the current node. The correct result can be returned only when the GUC parameter **track\_activities** is set to **on**. The timeout threshold ranges from 0 to 2147483.  
Return type: SETOF record

| Name             | Type        | Description                                                                                                                                                                                  |
|------------------|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| database         | name        | Name of the database to which a user session is connected                                                                                                                                    |
| pid              | bigint      | Backend thread ID                                                                                                                                                                            |
| sessionid        | bigint      | Session ID                                                                                                                                                                                   |
| usesysid         | oid         | OID of the user logged in to the backend                                                                                                                                                     |
| application_name | text        | Name of the application connected to the backend                                                                                                                                             |
| query            | text        | Query that is being executed on the backend                                                                                                                                                  |
| xact_start       | timestamptz | Time when the current transaction is started                                                                                                                                                 |
| query_start      | timestamptz | Time when the current query starts For a stored procedure, function, or package, the first query time is queried and does not change with the running of statements in the stored procedure. |
| query_id         | bigint      | Query statement ID                                                                                                                                                                           |



- **gs\_wlm\_user\_resource\_info(name text)**  
Description: Queries a user's resource quota and resource usage. Common users can query only their own information. Administrators can query information about all users.  
Return type: record
- **create\_wlm\_instance\_statistics\_info**  
Description: Saves the historical monitoring data of the current instance persistently.  
Parameter: **nan**  
Return type: integer
- **gs\_wlm\_persistent\_user\_resource\_info()**  
Description: Archives all user resource usage statistics to the `gs_wlm_user_resource_history` system catalog. To query this function, you must have the SYSADMIN permission.  
Return type: record
- **create\_wlm\_operator\_info(int flag)**  
Description: Clears top SQL operator-level statistics recorded in the current memory. If the input parameter is greater than 0, the information is archived to `gs_wlm_operator_info`. Otherwise, the information is not archived. Only users with the SYSADMIN permission can execute this function.  
Return type: int
- **GS\_ALL\_NODEGROUP\_CONTROL\_GROUP\_INFO(text)**  
Description: Provides Cgroup information for all logical database instances. Before calling this function, you need to specify the name of the logical database instance to be queried.  
For example, to query the Cgroup information for the '**installation**' logical database instance, run the following command:  

```
SELECT * FROM GS_ALL_NODEGROUP_CONTROL_GROUP_INFO('installation')
```

  
Return type: record

The following table describes return columns.

| Name     | Type   | Description                                             |
|----------|--------|---------------------------------------------------------|
| name     | text   | Cgroup name.                                            |
| type     | text   | Cgroup type.                                            |
| gid      | bigint | Cgroup ID.                                              |
| classgid | bigint | ID of the class Cgroup where a workload Cgroup belongs. |
| class    | text   | Class Cgroup.                                           |
| workload | text   | Workload Cgroup.                                        |
| shares   | bigint | CPU quota allocated to the Cgroup.                      |
| limits   | bigint | Limit of CPU resources allocated to a Cgroup.           |

| Name     | Type   | Description                                       |
|----------|--------|---------------------------------------------------|
| wdlevel  | bigint | Workload Cgroup level.                            |
| cpucores | text   | Information about the CPU cores used by a Cgroup. |

- `gs_total_nodegroup_memory_detail`  
Description: Returns information about the memory used by the current logical database, in MB.  
Return type: SETOF record
- `local_redo_time_count()`  
Description: Returns the time consumption statistics on each process of each replayer thread on the current node (valid data exists only on the standby node).

The return values are as follows:

`local_redo_time_count` parameters

| Column                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>thread_name</code> | Thread name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <code>step1_total</code> | Total duration of step 1. The process of each thread is as follows: <ul style="list-style-type: none"><li>• Ultimate RTO:<ul style="list-style-type: none"><li><b>batch redo</b>: obtains a log from a queue.</li><li><b>redo manager</b>: obtains a log from a queue.</li><li><b>redo worker</b>: obtains a log from a queue.</li><li><b>txn manager</b>: reads a log from a queue.</li><li><b>txn worker</b>: reads a log from a queue.</li><li><b>read worker</b>: reads an Xlog page (overall) from a file.</li><li><b>read page worker</b>: obtains a log from a queue.</li><li><b>startup</b>: obtains a log from a queue.</li></ul></li><li>• Parallel replay:<ul style="list-style-type: none"><li><b>page redo</b>: obtains a log from a queue.</li><li><b>startup</b>: reads a log.</li></ul></li></ul> |
| <code>step1_count</code> | Number of accumulated execution times of step 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Column      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| step2_total | <p>Total duration of step 2. The process of each thread is as follows:</p> <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>batch redo</b>: processes logs (overall).</li> <li><b>redo manager</b>: processes logs (overall).</li> <li><b>redo worker</b>: processes logs (overall).</li> <li><b>txn manager</b>: processes logs (overall).</li> <li><b>txn worker</b>: processes logs (overall).</li> <li><b>redo worker</b>: specifies the time required for reading the Xlog page.</li> <li><b>read page worker</b>: generates and sends LSN forwarders.</li> <li><b>startup</b>: checks whether to replay to the specified position.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: processes logs (overall).</li> <li><b>startup</b>: checks whether to replay to the specified position.</li> </ul> </li> </ul> |
| step2_count | Number of accumulated execution times of step 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| step3_total | <p>Total duration of step 3. The process of each thread is as follows:</p> <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>batch redo</b>: updates the standby state.</li> <li><b>redo manager</b>: processes data logs.</li> <li><b>redo worker</b>: replays page logs (overall).</li> <li><b>txn manager</b>: updates the flushing LSN.</li> <li><b>txn worker</b>: replays logs.</li> <li><b>redo worker</b>: pushes the Xlog segment.</li> <li><b>read page worker</b>: obtains a new item.</li> <li><b>startup</b>: collects statistics on the wait time of delayed replay feature.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: updates the standby state.</li> <li><b>startup</b>: collects statistics on the wait time of delayed replay feature.</li> </ul> </li> </ul>                                   |
| step3_count | Number of accumulated execution times of step 3.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

| Column      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| step4_total | <p>Total duration of step 4. The process of each thread is as follows:</p> <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>batch redo</b>: parses Xlogs.</li> <li><b>redo manager</b>: processes DDL operations.</li> <li><b>redo worker</b>: reads data pages.</li> <li><b>txn manager</b>: synchronizes the wait time.</li> <li><b>txn worker</b>: updates the LSN of the current thread.</li> <li><b>read page worker</b>: stores logs in the distribution thread.</li> <li><b>startup</b>: distributes logs (overall).</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays undo logs.</li> <li><b>startup</b>: distributes logs (overall).</li> </ul> </li> </ul>                    |
| step4_count | Number of accumulated execution times of step 4.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| step5_total | <p>Total duration of step 5. The process of each thread is as follows:</p> <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>batch redo</b>: distributes logs to the redo manager.</li> <li><b>redo manager</b>: distributes logs to redo workers.</li> <li><b>redo worker</b>: replays data page logs.</li> <li><b>txn manager</b>: distributes data to the txn worker.</li> <li><b>txn worker</b>: forcibly synchronizes the wait time.</li> <li><b>read page worker</b>: updates the LSN of the current thread.</li> <li><b>startup</b>: decodes logs.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays sharetxn logs.</li> <li><b>startup</b>: replays logs.</li> </ul> </li> </ul> |
| step5_count | Number of accumulated execution times of step 5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Column      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| step6_total | Total duration of step 6. The process of each thread is as follows: <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>redo worker</b>: replays non-data page logs.</li> <li><b>txn manager</b>: updates global LSNs.</li> <li><b>read page worker</b>: performs CRC on logs.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays synctrxn logs.</li> <li><b>startup</b>: forcibly synchronizes the wait time.</li> </ul> </li> </ul> |
| step6_count | Number of accumulated execution times of step 6.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| step7_total | Total duration of step 7. The process of each thread is as follows: <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>redo manager</b>: creates tablespaces.</li> <li><b>redo worker</b>: updates FSM.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays a single log.</li> </ul> </li> </ul>                                                                                                                                      |
| step7_count | Number of accumulated execution times of step 7.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| step8_total | Total duration of step 8. The process of each thread is as follows: <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li><b>redo worker</b>: forcibly synchronizes the wait time.</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays all workers do logs.</li> </ul> </li> </ul>                                                                                                                                                          |
| step8_count | Number of accumulated execution times of step 8.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| step9_total | Total duration of step 9. The process of each thread is as follows: <ul style="list-style-type: none"> <li>Ultimate RTO: <ul style="list-style-type: none"> <li>None</li> </ul> </li> <li>Parallel replay: <ul style="list-style-type: none"> <li><b>page redo</b>: replays multi-worker do logs.</li> </ul> </li> </ul>                                                                                                                                                                                                             |
| step9_count | Number of accumulated execution times of step 9.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- local\_xlog\_redo\_statics()

Description: Returns the statistics on each type of logs that have been replayed on the current node (valid data exists only on the standby node).

The return values are as follows:

**Table 7-80** local\_xlog\_redo\_stats parameters

| Column    | Description                                                                                                                                                                                                                                          |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| xlog_type | Log types.                                                                                                                                                                                                                                           |
| rmid      | Resource manager ID.                                                                                                                                                                                                                                 |
| info      | Xlog operation.                                                                                                                                                                                                                                      |
| num       | Number of logs.                                                                                                                                                                                                                                      |
| extra     | Valid values are available for page replay logs and xact logs. <ul style="list-style-type: none"><li>• The page replay logs indicate the number of pages read from the disk.</li><li>• The xact logs indicate the number of deleted files.</li></ul> |

- `gs_get_shared_memctx_detail(text)`

Description: Returns the memory allocation details of the specified memory context, including the file, line number, and size of each memory allocation (the size of the same line in the same file is accumulated). Only the memory context queried through the **pg\_shared\_memory\_detail** view can be queried. The input parameter is the memory context name (that is, the **contextname** column in the result returned by the **pg\_shared\_memory\_detail** view). To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: SETOF record

| Name | Type | Description                                                                                                                             |
|------|------|-----------------------------------------------------------------------------------------------------------------------------------------|
| file | text | Name of the file where the memory is applied for.                                                                                       |
| line | int8 | Line number of the code in the file where the requested memory is located.                                                              |
| size | int8 | Size of the allocated memory. The value is accumulated if the memory is allocated for multiple times to the same line of the same file. |

- `gs_get_session_memctx_detail(text)`

Description: Returns the memory allocation details of the specified memory context, including the file, line number, and size of each memory allocation (the size of the same line in the same file is accumulated). This parameter is valid only in thread pool mode. Only the memory context queried through the **gs\_session\_memory\_context** view can be queried. The input parameter is the memory context name (that is, the **contextname** column in the result

returned by the **gs\_session\_memory\_context** view). To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: SETOF record

| Name | Type | Description                                                                                                                                       |
|------|------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| file | text | Name of the file where the memory is applied for.                                                                                                 |
| line | int8 | Line number of the code in the file where the requested memory is located.                                                                        |
| size | int8 | Size of the allocated memory, in bytes. The value is accumulated if the memory is allocated for multiple times to the same line of the same file. |

#### NOTE

This view takes effect only in thread pool mode.

- **gs\_get\_history\_memory\_detail(cstring)**  
Description: Queries historical memory snapshot information. The input parameter type is cstring. The value can be **NULL** or the name of the memory snapshot log file.
  - If the value of the input parameter is **NULL**, the list of all memory snapshot log files on the current node is displayed.
  - If the value of the input parameter is the name of the memory snapshot log file in the list queried in [a](#), the detailed information about the memory snapshot recorded in the log file is displayed.
  - If you enter any other input parameter, the system displays a message indicating that the input parameter is incorrect or the file fails to be opened.

To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: text

| Name        | Type | Description                                                                                                                                                                                                                                 |
|-------------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| memory_info | text | Memory information. If the input parameter of the function is set to <b>NULL</b> , the memory snapshot file list is displayed. If the input parameter is set to the name of the memory snapshot file, the content of the file is displayed. |

- **gs\_stack()**  
Description: Displays the call stack of a thread. To query this function, you must have the SYSADMIN or MONADMIN permission.  
Parameter: tid, which indicates the thread ID. **tid** is an optional parameter. If it is specified, the function returns the call stack of the thread corresponding to **tid**. If it is not specified, the function returns the call stacks of all threads.

Return value: If **tid** is specified, the return value is of the TEXT type. If **tid** is not specified, the return value is a SETOF record.

Example:

Run **select \* from gs\_stack(tid)** to obtain the call stack of a specified thread.

```
openGauss=# select * from gs_stack(139663481165568);
gs_stack
```

```

__poll + 0x2d
WaitLatchOrSocket(Latch volatile*, int, int, long) + 0x29f
WaitLatch(Latch volatile*, int, long) + 0x2e
JobScheduleMain() + 0x90f
int GaussDbThreadMain<(knl_thread_role)9>(knl_thread_arg*) + 0x456+
InternalThreadFunc(void*) + 0x2d
ThreadStarterFunc(void*) + 0xa4
start_thread + 0xc5
clone + 0x6d
(1 row)
```

Run **select \* from gs\_stack()** to obtain the call stacks of all threads.

```
openGauss=# select * from gs_stack();
```

```
-[RECORD
```

```
1]-----
tid | 139670364324352
lwtid | 308
stack | __poll + 0x2d
 | CommWaitPollParam::caller(int (*)(pollfd*, unsigned long, int), unsigned long) + 0x34
 | int comm_socket_call<CommWaitPollParam, int (*)(pollfd*, unsigned long,
 | int)>(CommWaitPollParam*, int (*)(pollfd*, unsigned long
 | , int)) + 0x28
 | comm_poll(pollfd*, unsigned long, int) + 0xb1
 | ServerLoop() + 0x72b
 | PostmasterMain(int, char**) + 0x314e
 | main + 0x617
 | __libc_start_main + 0xf5
 | 0x55d38f8db3a7
```

```
[RECORD 2]-----
```

```
tid | 139664851859200
lwtid | 520
stack | __poll + 0x2d
 | WaitLatchOrSocket(Latch volatile*, int, int, long) + 0x29f
 | SysLoggerMain(int) + 0xc86
 | int GaussDbThreadMain<(knl_thread_role)17>(knl_thread_arg*) + 0x45d
 | InternalThreadFunc(void*) + 0x2d
 | ThreadStarterFunc(void*) + 0xa4
 | start_thread + 0xc5
 | clone + 0x6d
```

- **gs\_get\_thread\_memctx\_detail(tid,text)**

Description: Returns the memory allocation details of the specified memory context, including the file, line number, and size of each memory allocation (the size of the same line in the same file is accumulated). Only the memory context queried through the **gs\_thread\_memory\_context** view is supported. The first input parameter is the thread ID (the **tid** column of the data returned by the **gs\_thread\_memory\_context**), and the second parameter is the memory context name (the **contextname** column of the data returned by **gs\_thread\_memory\_context**). To query this function, you must have the SYSADMIN or MONADMIN permission.

Return type: SETOF record

| Name | Type | Description                                       |
|------|------|---------------------------------------------------|
| file | text | Name of the file where the memory is applied for. |



| Name | Type | Description                                                                                                                                       |
|------|------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| line | int8 | Line number of the code in the file where the requested memory is located.                                                                        |
| size | int8 | Size of the allocated memory, in bytes. The value is accumulated if the memory is allocated for multiple times to the same line of the same file. |

## 7.5.25 Trigger Functions

- pg\_get\_triggerdef(oid)

Description: Obtains the definition information of a trigger.

Parameter: OID of the trigger to be queried

Return type: text

Example:

```
openGauss=# select pg_get_triggerdef(oid) from pg_trigger;
 pg_get_triggerdef

CREATE TRIGGER tg1 BEFORE INSERT ON gtest26 FOR EACH STATEMENT EXECUTE PROCEDURE
gtest_trigger_func()
CREATE TRIGGER tg03 AFTER INSERT ON gtest26 FOR EACH ROW WHEN ((new.a IS NOT NULL))
EXECUTE PROCEDURE gtest_trigger_func()
(2 rows)
```

- pg\_get\_triggerdef(oid, boolean)

Description: Obtains the definition information of a trigger.

Parameter: OID of the trigger to be queried and whether it is displayed in pretty mode

### NOTE

Boolean parameters take effect only when the WHEN condition is specified during trigger creation.

Return type: text

Example:

```
openGauss=# select pg_get_triggerdef(oid,true) from pg_trigger;
 pg_get_triggerdef

CREATE TRIGGER tg1 BEFORE INSERT ON gtest26 FOR EACH STATEMENT EXECUTE PROCEDURE
gtest_trigger_func()
CREATE TRIGGER tg03 AFTER INSERT ON gtest26 FOR EACH ROW WHEN (new.a IS NOT NULL)
EXECUTE PROCEDURE gtest_trigger_func()
(2 rows)

openGauss=# select pg_get_triggerdef(oid,false) from pg_trigger;
 pg_get_triggerdef

CREATE TRIGGER tg1 BEFORE INSERT ON gtest26 FOR EACH STATEMENT EXECUTE PROCEDURE
gtest_trigger_func()
CREATE TRIGGER tg03 AFTER INSERT ON gtest26 FOR EACH ROW WHEN ((new.a IS NOT NULL))
EXECUTE PROCEDURE gtest_trigger_func()
(2 rows)
```

## 7.5.26 Hash Function

- `hash_array(anyarray)`

Description: Hashes an array, obtains the result of an array element using the hash function, and returns the combination result.

Parameter: data of the anyarray type

Return type: integer

Example:

```
openGauss=# SELECT hash_array(ARRAY[[1,2,3],[1,2,3]]);
 hash_array

-382888479
(1 row)
```

- `hash_group(key)`

Description: Calculates the hash value of each column in the Group Clause in the streaming engine. (Due to specification changes, the current version no longer supports this feature. Do not use it.)

Parameter: **key** indicates the value of each column in the Group Clause.

Return type: 32-bit hash value

Example:

```
Perform the following steps in sequence.
openGauss=# CREATE TABLE tt(a int, b int,c int,d int);
NOTICE: The 'DISTRIBUTE BY' clause is not specified. Using 'a' as the distribution column by default.
HINT: Please use 'DISTRIBUTE BY' clause to specify suitable data distribution column.
CREATE TABLE
openGauss=# SELECT * FROM tt;
 a | b | c | d
---+---+---+---
(0 rows)

openGauss=# INSERT INTO tt VALUES(1,2,3,4);
INSERT 0 1
openGauss=# SELECT * FROM tt;
 a | b | c | d
---+---+---+---
 1 | 2 | 3 | 4
(1 row)

openGauss=# INSERT INTO tt VALUES(5,6,7,8);
INSERT 0 1
openGauss=# SELECT * FROM tt;
 a | b | c | d
---+---+---+---
 1 | 2 | 3 | 4
 5 | 6 | 7 | 8
(2 rows)

openGauss=# SELECT hash_group(a,b) FROM tt WHERE a=1 and b=2;
 hash_group

 990882385
(1 row)

openGauss=# DROP TABLE tt;
DROP TABLE
```

- `hash_numeric(numeric)`

Description: Calculates the hash value of numeric data.

Parameter: data of the numeric type.

Return type: integer

Example:

```
openGauss=# SELECT hash_numeric(30);
 hash_numeric

-282860963
(1 row)
```

- **hash\_range(anyrange)**

Description: Calculates the hash value of a range.

Parameter: data of the anyrange type

Return type: integer

Example:

```
openGauss=# SELECT hash_range(numrange(1.1,2.2));
 hash_range

683508754
(1 row)
```

- **hashbpchar(character)**

Description: Calculates the hash value of bpchar.

Parameter: data of the character type

Return type: integer

Example:

```
openGauss=# SELECT hashbpchar('hello');
 hashbpchar

-1870292951
(1 row)
```

- **hashchar(char)**

Description: Converts char and Boolean data into hash values.

Parameter: data of the char or bool type

Return type: integer

Example:

```
openGauss=# SELECT hashbpchar('hello');
 hashbpchar

-1870292951
(1 row)
```

```
openGauss=# SELECT hashchar('true');
 hashchar

1686226652
(1 row)
```

- **hashenum(anyenum)**

Description: Converts enumerated values to hash values.

Parameter: data of the anyenum type

Return type: integer

Example:

```
openGauss=# CREATE TYPE b1 AS ENUM('good', 'bad', 'ugly');
CREATE TYPE
openGauss=# CALL hashenum('good'::b1);
 hashenum
```

```

1821213359
(1 row)
openGauss=# DROP TYPE b1;
DROP TYPE
```

- **hashfloat4(real)**

Description: Converts float4 values to hash values.

Parameter: data of the real type

Return type: integer

Example:

```
openGauss=# SELECT hashfloat4(12.1234);
hashfloat4

1398514061
(1 row)
```

- **hashfloat8(double precision)**

Description: Converts float8 values to hash values.

Parameter: data of the double precision type

Return type: integer

Example:

```
openGauss=# SELECT hashfloat8(123456.1234);
hashfloat8

1673665593
(1 row)
```

- **hashinet(inet)**

Description: Supports hashing indexes on inet or cidr. Returns the hash value of inet.

Parameter: data of the inet type

Return type: integer

Example:

```
openGauss=# SELECT hashinet('127.0.0.1 '::inet);
hashinet

-1435793109
(1 row)
```

- **hashint1(tinyint)**

Description: Converts INT1 values to hash values.

Parameter: data of the tinyint type

Return type: uint32

Example:

```
openGauss=# SELECT hashint1(20);
hashint1

-2014641093
(1 row)
```

- **hashint2(smallint)**

Description: Converts INT2 values to hash values.

Parameter: data of the smallint type

Return type: uint32

Example:

```
openGauss=# SELECT hashint2(20000);
 hashint2

-863179081
(1 row)
```

## 7.5.27 Prompt Message Function

- report\_application\_error

Description: This function can be used to throw errors during PL execution.

Return type: void

**Table 7-81** report\_application\_error parameter description

| Parameter | Type | Description                                                                           | Required (Yes/No) |
|-----------|------|---------------------------------------------------------------------------------------|-------------------|
| log       | text | Content of an error message.                                                          | Yes               |
| code      | int4 | Error code corresponding to an error message. The value ranges from -20999 to -20000. | No                |

Example:

```
openGauss=# CREATE OR REPLACE FUNCTION GET_RESULT_UNKNOWNN(
openGauss(# IN context_id int, /*context_id*/
openGauss(# IN col_type text /*col_type*/
openGauss(#)RETURNS INTEGER
openGauss-# AS $$
openGauss$$ BEGIN
openGauss$$ if col_type is NULL then
openGauss$$ PG_CATALOG.REPORT_APPLICATION_ERROR('invalid input for the third parameter
col_type should not be null');
openGauss$$ end if;
openGauss$$ PG_CATALOG.REPORT_APPLICATION_ERROR('UnSupport data type for
column_value(context: '||context_id||', '||PG_CATALOG.QUOTE_LITERAL(col_type)||')');
openGauss$$ return -1;
openGauss$$ END;
openGauss$$ $$ LANGUAGE plpgsql;
CREATE FUNCTION
openGauss=# CALL GET_RESULT_UNKNOWNN(NULL, NULL);
ERROR: invalid input for the third parameter col_type should not be null
CONTEXT: SQL statement "CALL pg_catalog.report_application_error('invalid input for the third
parameter col_type should not be null')
PL/pgSQL function get_result_unknown(integer,text) line 4 at PERFORM
```

## 7.5.28 Global Temporary Table Functions

- pg\_get\_gtt\_relstats(relOid)

Description: Displays basic information about a global temporary table specified by the current session.

Parameter: OID of the global temporary table

Return type: record

Example:

```
openGauss=# select * from pg_get_gtt_relstats(74069);
 relfilenode | relpages | reltuples | relallvisible | relfrozensxid | relminmxid
-----+-----+-----+-----+-----+-----
 74069 | 58 | 13000 | 0 | 11151 | 0
(1 row)
```

- pg\_get\_gtt\_statistics(relOid, attnum, "::text)

Description: Displays statistics about a single column in a global temporary table specified by the current session.

Parameter: `OID` and the **attnum** attribute of the global temporary table

Return type: record

Example:

[illegible]

- `pg_gtt_attached_pid(relOid)`

Description: Displays PIDs and session IDs of all threads that are using the specified global temporary table. If the thread pool is enabled and the session that is using the global temporary table is in the detach state, then the PID is **0** and **sessionid** indicates the session ID.

Parameter: OID of the global temporary table

Return type: record

Example:

```
openGauss=# select * from pg_gtt_attached_pid(74069);
relid | pid | sessionid
-----+-----+-----
74069 | 139648170456832 | 139648170456832
74069 | 139648123270912 | 139648123270912
(2 rows)
```

- `dbperf.get_global_full_sql_by_timestamp(start_timestamp timestamp, end_timestamp timestamp)`

Description: Obtains full SQL information at the instance level.

Return type: record

**Table 7-82** `dbperf.get_global_full_sql_by_timestamp` parameter description

| Parameter                    | Type      | Description                              |
|------------------------------|-----------|------------------------------------------|
| <code>start_timestamp</code> | timestamp | Start point of the SQL start time range. |
| <code>end_timestamp</code>   | timestamp | End point of the SQL start time range.   |

- `dbperf.get_global_slow_sql_by_timestamp(start_timestamp timestamp, end_timestamp timestamp)`

Description: Obtains slow SQL information at the instance level.

Return type: record

**Table 7-83** `dbperf.get_global_slow_sql_by_timestamp` parameter description

| Parameter                    | Type      | Description                              |
|------------------------------|-----------|------------------------------------------|
| <code>start_timestamp</code> | timestamp | Start point of the SQL start time range. |
| <code>end_timestamp</code>   | timestamp | End point of the SQL start time range.   |

- `statement_detail_decode(detail text, format text, pretty bool)`

Parses the details column in a full or slow SQL statement.

**Table 7-84** `statement_detail_decode` parameter description

| Parameter           | Type | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>detail</code> | text | Set of events generated by the SQL statement (unreadable).                                                                                                                                                                                                                                                                                             |
| <code>format</code> | text | Parsing output format. The value is <b>plaintext</b> or <b>json</b> .                                                                                                                                                                                                                                                                                  |
| <code>pretty</code> | bool | Specifies whether to display the text in pretty format when <b>format</b> is set to <b>plaintext</b> . The options are as follows: <ul style="list-style-type: none"><li>• The value <b>true</b> indicates that events are separated by <code>\n</code>.</li><li>• The value <b>false</b> indicates that events are separated by commas (,).</li></ul> |

- `pg_list_gtt_relfrozenxids()`  
Description: Displays the frozen XID of each session.  
If the value of **pid** is **0**, the earliest frozen XID of all sessions is displayed.  
Parameter: none  
Return type: record  
Example:

```
openGauss=# select * from pg_list_gtt_relfrozenxids();
 pid | relfrozenxid
-----+-----
139648123270912 | 11151
139648170456832 | 11155
 0 | 11151
(3 rows)
```

## 7.5.29 Fault Injection System Function

- `gs_fault_inject(int64, text, text, text, text, text)`  
Description: This function cannot be called. WARNING information "unsupported fault injection" is reported when this function is called, which does not affect or change the database.  
Parameter: fault injection of the int64 type (**0**: CLOG extended page; **1**: CLOG page reading; **2**: forcible deadlock)
  - If the first input parameter of text is set to **2** and the second input parameter of text is set to **1**, the second input parameter deadlock occurs. Other input parameters are not deadlocked. When the first input parameter is **0** or **1**, the second input parameter indicates the number of the start page from which the CLOG starts to be extended or read.
  - The third input parameter of text indicates the number of extended or read pages when the first input parameter is **0** or **1**.
  - The fourth to sixth input parameters of text are reserved.Return type: int64

## 7.5.30 AI Feature Functions

- `gs_index_advise(text)`  
Description: Recommends an index for a single query statement.  
Parameter: SQL statement string  
Return type: record
- `hypopg_create_index(text)`  
Description: Creates a virtual index.  
Parameter: character string of the statement for creating an index  
Return type: record
- `hypopg_display_index()`  
Description: Displays information about all created virtual indexes.  
Parameter: none  
Return type: record
- `hypopg_drop_index(oid)`



Description: Deletes a specified virtual index.

Parameter: OID of the index

Return type: Boolean

- `hypopg_reset_index()`

Description: Clears all virtual indexes.

Parameter: none

Return type: none

- `hypopg_estimate_size(oid)`

Description: Estimates the space required for creating a specified index.

Parameter: OID of the index

Return type: int8

- `check_engine_status(ip text, port text)`

Description: Tests whether a predictor engine provides services on a specified IP address and port.

Parameter: IP address and port number of the predictor engine.

Return type: text

#### NOTE

This function is unavailable in the current version.

- `encode_plan_node(optname text, orientation text, strategy text, options text, dop int8, quals text, projection text)`

Description: Encodes the plan operator information in the input parameters.

Parameter: plan operator information

Return type: text

#### NOTE

This function is an internal function. You are advised not to use it directly.

- `model_train_opt(template text, model text)`

Description: Trains a given query performance prediction model.

Parameters: template name and model name of the performance prediction model

Return type: startup\_time\_accuracy FLOAT8, total\_time\_accuracy FLOAT8, rows\_accuracy FLOAT8, peak\_memory\_accuracy FLOAT8

#### NOTE

This function is unavailable in the current version.

- `track_model_train_opt(ip text, port text)`

Description: Returns the training log address of the specified IP address and port predictor engine.

Parameter: IP address and port number of the predictor engine

Return type: text

#### NOTE

This function is unavailable in the current version.

- `encode_feature_perf_hist(datname text)`  
Description: Encodes historical plan operators collected in the target database.  
Parameter: database name  
Return type: queryid bigint, plan\_node\_id int, parent\_node\_id int, left\_child\_id int, right\_child\_id int, encode text, startup\_time bigint, total\_time bigint, rows bigint, and peak\_memory int
- `gather_encoding_info(datname text)`  
Description: Calls `encode_feature_perf_hist` to save the encoded data persistently.  
Parameter: database name  
Return type: int

## 7.5.31 Dynamic Data Masking Functions

### NOTE

This function is an internal function.

- `creditcardmasking(col text, letter char default 'x')`  
Description: Replaces the digits before the last four bits following the col string with letters.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text
- `basicmailmasking(col text, letter char default 'x')`  
Description: Replaces the characters before the first at sign (@) in the col string with letters.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text
- `fullmailmasking(col text, letter char default 'x')`  
Description: Replaces the characters (except @) before the last period (.) in the col string with letters.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text
- `alldigitsmasking(col text, letter char default '0')`  
Description: Replaces the digits in the col string with letters.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text
- `shufflemasking(col text)`  
Description: Sorts the characters in the col string out of order.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text

- **randommasking(col text)**  
Description: Randomizes the characters in the col string.  
Parameter: Character string to be replaced or character string used for replacement  
Return type: text
- **regexpmasking**  
Description: Specifies the internal function of the masking policy, which is used to replace characters using a regular expression.  
Parameter: col text, reg text, replace\_text text, pos INTEGER default 0, reg\_len INTEGER default -1  
Return type: text

## 7.5.32 Hierarchical Recursion Query Functions

The following functions can be used in a hierarchical recursion query statement to return information about the connection path.

- **sys\_connect\_by\_path(col, separator)**  
Description: Returns the connection path from the root node to the current row. This function applies only to hierarchical recursion queries.  
The **col** parameter indicates the name of the column displayed in the path. Only columns of the CHAR, VARCHAR, NVARCHAR2, and TEXT types are supported. The **separator** parameter indicates the separator between path nodes.  
Return type: text

Example:

```
openGauss=# select *, sys_connect_by_path(name, '-') from connect_table start with id = 1 connect by
prior id = pid;
 id | pid | name | sys_connect_by_path
-----+-----+-----+-----
 1 | 0 | a | -a
 2 | 1 | b | -a-b
 4 | 1 | d | -a-d
 3 | 2 | c | -a-b-c
(4 rows)
```

- **connect\_by\_root(col)**  
Description: Returns the value of a column in the top-most parent row of the current row. This function applies only to hierarchical recursion queries.  
The **col** parameter indicates the name of an output column.  
Return type: data type of the specified column **col**.

Example:

```
openGauss=# select *, connect_by_root(name) from connect_table start with id = 1 connect by prior
id = pid;
 id | pid | name | connect_by_root
-----+-----+-----+-----
 1 | 0 | a | a
 2 | 1 | b | a
 4 | 1 | d | a
 3 | 2 | c | a
(4 rows)
```

## 7.5.33 Internal Functions

The following functions of GaussDB use internal data types, which cannot be directly called by users.

- Selectivity calculation functions

|                          |                         |                          |                        |                       |                          |                  |
|--------------------------|-------------------------|--------------------------|------------------------|-----------------------|--------------------------|------------------|
| areajoin<br>el           | areasel                 | arraycon<br>tjoin<br>sel | arraycon<br>tsel       | contjoin<br>el        | contsel                  | eqjoin<br>sel    |
| eqsel                    | iclikejoin<br>sel       | iclikesel                | icnlikejoi<br>n<br>sel | icnlikese<br>l        | icregexe<br>qjoin<br>sel | icregexe<br>qsel |
| icregexn<br>ejoin<br>sel | icregexn<br>esel        | likejoin<br>el           | likesel                | neqjoin<br>el         | neqsel                   | nlikejoin<br>sel |
| nlikesel                 | positionj<br>oin<br>sel | positions<br>el          | regexej<br>oin<br>sel  | regexeq<br>s<br>el    | regexnej<br>oin<br>sel   | regexnes<br>el   |
| scalargtj<br>oin<br>sel  | scalargts<br>el         | scalartj<br>oin<br>sel   | scalart<br>s<br>el     | tmatchj<br>oin<br>sel | tmatchs<br>el            | -                |

- Statistics collection functions

|                 |                 |              |
|-----------------|-----------------|--------------|
| array_tyanalyze | range_tyanalyze | ts_tyanalyze |
| local_rto_stat  | -               | -            |

- Internal functions for sorting

|                        |                       |                      |                         |                           |
|------------------------|-----------------------|----------------------|-------------------------|---------------------------|
| bpchar_sorts<br>upport | bytea_sortsu<br>pport | date_sortsup<br>port | numeric_sort<br>support | timestamp_s<br>ortsupport |
|------------------------|-----------------------|----------------------|-------------------------|---------------------------|

- Internal type processing functions

|                   |                              |                    |                  |                              |                             |                                |
|-------------------|------------------------------|--------------------|------------------|------------------------------|-----------------------------|--------------------------------|
| abstimer<br>ecv   | euc_jis_2<br>004_to_<br>utf8 | int2recv           | line_recv        | oidvecto<br>rrecv_ext<br>end | tidrecv                     | utf8_to_<br>koi8u              |
| anyarray<br>_recv | euc_jp_t<br>o_mic            | int2vect<br>orrecv | lseg_rec<br>v    | path_rec<br>v                | time_rec<br>v               | utf8_to_<br>shift_jis_<br>2004 |
| array_re<br>cv    | euc_jp_t<br>o_sjis           | int4recv           | macaddr<br>_recv | pg_node<br>_tree_rec<br>v    | time_tra<br>nsform          | utf8_to_<br>sjis               |
| ascii_to_<br>mic  | euc_jp_t<br>o_utf8           | int8recv           | mic_to_a<br>scii | point_re<br>cv               | timesta<br>mp_recv          | utf8_to_<br>uhc                |
| ascii_to_<br>utf8 | euc_kr_t<br>o_mic            | internal_<br>out   | mic_to_b<br>ig5  | poly_rec<br>v                | timesta<br>mp_tran<br>sform | utf8_to_<br>win                |

|                                    |                      |                    |                |                    |                      |                   |
|------------------------------------|----------------------|--------------------|----------------|--------------------|----------------------|-------------------|
| big5_to_euc_tw                     | euc_kr_to_utf8       | interval_recv      | mic_to_euc_cn  | pound_nexttoken    | timestamp_tz_recv    | uuid_recv         |
| big5_to_mic                        | euc_tw_to_big5       | interval_transform | mic_to_euc_jp  | prsd_nexttoken     | timetz_recv          | varbit_recv       |
| big5_to_utf8                       | euc_tw_to_mic        | iso_to_koi8r       | mic_to_euc_kr  | range_recv         | tinterval_recv       | varbit_transform  |
| bit_recv                           | euc_tw_to_utf8       | iso_to_mic         | mic_to_euc_tw  | rawrecv            | tsquery_recv         | varchar_transform |
| boolrecv                           | float4_recv          | iso_to_win1251     | mic_to_iso     | record_recv        | tsvector_recv        | varchar_recv      |
| box_recv                           | float8_recv          | iso_to_win866      | mic_to_koi8r   | regclass_recv      | txid_snapshot_recv   | void_recv         |
| bpchar_recv                        | gb18030_to_utf8      | iso8859_1_to_utf8  | mic_to_latin1  | regconfig_recv     | uhc_to_utf8          | win_to_utf8       |
| btoidsortsupport                   | gbk_to_utf8          | iso8859_to_utf8    | mic_to_latin2  | regdictionary_recv | unknown_recv         | win1250_to_latin2 |
| bytea_recv                         | gin_extract_tsvector | johab_to_utf8      | mic_to_latin3  | regoperator_recv   | utf8_to_ascii        | win1250_to_mic    |
| byteawithoutorderwithequalcol_recv | gtsvector_compress   | json_recv          | mic_to_latin4  | regoperator_recv   | utf8_to_big5         | win1251_to_iso    |
| cash_recv                          | gtsvector_consistent | koi8r_to_iso       | mic_to_sjis    | regprocedure_recv  | utf8_to_euc_cn       | win1251_to_koi8r  |
| char_recv                          | gtsvector_decompress | koi8r_to_mic       | mic_to_win1250 | regprocedure_recv  | utf8_to_euc_jis_2004 | win1251_to_mic    |
| cidr_recv                          | gtsvector_penalty    | koi8r_to_utf8      | mic_to_win1251 | regtype_recv       | utf8_to_euc_jp       | win1251_to_win866 |
| cid_recv                           | gtsvector_picksplit  | koi8r_to_win1251   | mic_to_win866  | reltimer_recv      | utf8_to_euc_kr       | win866_to_iso     |

|                                            |                                 |                           |                              |                                             |                              |                               |
|--------------------------------------------|---------------------------------|---------------------------|------------------------------|---------------------------------------------|------------------------------|-------------------------------|
| circle_re<br>cv                            | gtsvecto<br>r_same              | koi8r_to<br>_win866       | namerec<br>v                 | shift_jis_<br>2004_to_<br>_euc_jis_<br>2004 | utf8_to_<br>euc_tw           | win866_t<br>o_koi8r           |
| cstring_r<br>ecv                           | gtsvecto<br>r_union             | koi8u_to<br>_utf8         | ngram_n<br>exttoken          | shift_jis_<br>2004_to_<br>_utf8             | utf8_to_<br>gb18030          | win866_t<br>o_mic             |
| date_rec<br>v                              | hll_recv                        | latin1_to<br>_mic         | numeric<br>_recv             | sjis_to_e<br>uc_jp                          | utf8_to_<br>gbk              | win866_t<br>o_win12<br>51     |
| domain_<br>recv                            | hll_trans<br>_recv              | latin2_to<br>_mic         | numeric<br>_transfor<br>m    | sjis_to_<br>mic                             | utf8_to_i<br>so8859          | xidrecv                       |
| euc_cn_t<br>o_mic                          | large_se<br>q_upgra<br>de_ntree | latin2_to<br>_win125<br>0 | nvarchar<br>2recv            | sjis_to_u<br>tf8                            | utf8_to_i<br>so8859_<br>1    | xidrecv4                      |
| euc_cn_t<br>o_utf8                         | inet_recv                       | latin3_to<br>_mic         | oidrecv                      | smalldat<br>etime_re<br>cv                  | utf8_to_j<br>ohab            | xml_recv                      |
| euc_jis_2<br>004_to_s<br>hift_jis_2<br>004 | int1recv                        | latin4_to<br>_mic         | oidvecto<br>rrecv            | textrecv                                    | utf8_to_<br>koi8r            | cstore_ti<br>d_out            |
| i16toi1                                    | int16                           | int16_bo<br>ol            | int16eq                      | int16div                                    | int16ge                      | int16gt                       |
| int16in                                    | int16le                         | int16lt                   | int16mi                      | int16mul                                    | int16ne                      | int16out                      |
| int16pl                                    | int16rec<br>v                   | int16sen<br>d             | numeric<br>_bool             | int2vect<br>orin_ext<br>end                 | int2vect<br>orout_ex<br>tend | int2vect<br>orrecv_e<br>xtend |
| int2vect<br>orsend_e<br>xtend              | tdigest_i<br>n                  | tdigest_<br>merge         | tdigest_<br>merge_t<br>o_one | tdigest_<br>mergep                          | tdigest_<br>out              | -                             |

- Internal functions for aggregation operations

|                                 |                                            |                                  |                                          |                                          |                                          |                                            |
|---------------------------------|--------------------------------------------|----------------------------------|------------------------------------------|------------------------------------------|------------------------------------------|--------------------------------------------|
| array_ag<br>g_finalfn           | array_ag<br>g_transf<br>n                  | bytea_st<br>ring_agg<br>_finalfn | bytea_st<br>ring_agg<br>_transfn         | date_list<br>_agg_no<br>arg2_tra<br>nsfn | date_list<br>_agg_tra<br>nsfn            | float4_li<br>st_agg_n<br>oarg2_tr<br>ansfn |
| float4_li<br>st_agg_t<br>ransfn | float8_li<br>st_agg_n<br>oarg2_tr<br>ansfn | float8_li<br>st_agg_t<br>ransfn  | int2_list<br>_agg_noa<br>rg2_tran<br>sfn | int2_list<br>_agg_tran<br>sfn            | int4_list<br>_agg_noa<br>rg2_tran<br>sfn | int4_list<br>_agg_tran<br>sfn              |

|                               |                                |                                  |                           |                                   |                            |                                      |
|-------------------------------|--------------------------------|----------------------------------|---------------------------|-----------------------------------|----------------------------|--------------------------------------|
| int8_list_agg_noarg2_transfn  | int8_list_agg_transfn          | interval_list_agg_noarg2_transfn | interval_list_agg_transfn | list_agg_finalfn                  | list_agg_noarg2_transfn    | list_agg_transfn                     |
| median_float8_finalfn         | median_interval_finalfn        | median_transfn                   | mode_final                | numeric_list_agg_noarg2_transfn   | numeric_list_agg_transfn   | ordered_set_transition               |
| percentile_cont_float8_final  | percentile_cont_interval_final | string_agg_finalfn               | string_agg_transfn        | timestamp_list_agg_noarg2_transfn | timestamp_list_agg_transfn | timestamp_tz_list_agg_noarg2_transfn |
| timestamp_tz_list_agg_transfn | checksum_text_agg_transfn      | -                                | -                         | -                                 | -                          | -                                    |

- Hash internal functions

|                |            |                 |                |                   |              |                 |
|----------------|------------|-----------------|----------------|-------------------|--------------|-----------------|
| hashbegin_scan | hashbuild  | hashbuild_empty | hashbulkdelete | hashcost_estimate | hashend_scan | hashget_bitmap  |
| hashgettuple   | hashinsert | hashmarkpos     | hashmerge      | hashrescan        | hashretpos   | hashvacuumclean |
| hashvarlena    | -          | -               | -              | -                 | -            | -               |

- Internal functions of the B-tree index

|               |                  |                    |                   |                   |                     |                     |
|---------------|------------------|--------------------|-------------------|-------------------|---------------------|---------------------|
| cbtreebuild   | cbtreecanreturn  | cbtreecostestimate | cbtreegetbitmap   | cbtreegettup      | btbegin_scan        | btbuild             |
| btbuild_empty | btbulkdelete     | btcanreturn        | btcostestimate    | btendscan         | btfloat4sortsupport | btfloat8sortsupport |
| btgetbitmap   | btgettup         | btinsert           | btint2sortsupport | btint4sortsupport | btint8sortsupport   | btmarkpos           |
| btmerge       | btnameortsupport | btrescan           | btrestropos       | bttextsortsupport | btvacuumclean       | cbtreeoptions       |

- Internal functions of the Psort index

|            |                |                   |                |                    |
|------------|----------------|-------------------|----------------|--------------------|
| psortbuild | psortcanreturn | psortcostestimate | psortgetbitmap | psortgettupletuple |
|------------|----------------|-------------------|----------------|--------------------|

- Internal functions of the UB-tree index

|                 |            |               |                  |              |
|-----------------|------------|---------------|------------------|--------------|
| ubtbeginscan    | ubtbuild   | ubtbuildempty | ubtbulkdelete    | ubtcanreturn |
| ubtcostestimate | ubtendscan | ubtgetbitmap  | ubtgettupletuple | ubtinsert    |
| ubtmarkpos      | ubtmerge   | ubtoptions    | ubtrescan        | ubtrestrpos  |
| ubtvacuumclean  | -          | -             | -                | -            |

- PL/pgSQL internal function  
plpgsql\_inline\_handler
- Set-related internal functions

|                      |                      |                         |                      |                     |                    |
|----------------------|----------------------|-------------------------|----------------------|---------------------|--------------------|
| array_indexby_delete | array_indexby_length | array_integer_deleteidx | array_integer_exists | array_integer_first | array_integer_last |
| array_integer_next   | array_integer_prior  | array_varchar_deleteidx | array_varchar_exists | array_varchar_first | array_varchar_last |
| array_varchar_next   | array_varchar_prior  | -                       | -                    | -                   | -                  |

- Foreign table-related internal functions

|                  |               |                       |                    |                  |                    |                 |
|------------------|---------------|-----------------------|--------------------|------------------|--------------------|-----------------|
| dist_fdw_handler | roach_handler | streaming_fdw_handler | dist_fdw_validator | file_fdw_handler | file_fdw_validator | log_fdw_handler |
|------------------|---------------|-----------------------|--------------------|------------------|--------------------|-----------------|

- Auxiliary function for the primary DN to remotely read the data page from the standby DN.

**gs\_read\_block\_from\_remote** is used to read the pages of a non-segment-page table file. By default, only the initial user can view the data. Other users can use the data only after being granted with permissions.

**gs\_read\_segment\_block\_from\_remote** is used to read the pages of a segment-page table file. By default, only the initial user can view the data. Other users can use the data only after being granted with permissions.

- Auxiliary function for the primary DN to remotely read the data file from the standby DN.

The **gs\_read\_file\_from\_remote** command is used to read a specified file. After obtaining the file size using the **gs\_read\_file\_size\_from\_remote** function, **gs\_repair\_file** reads the remote file segment by segment using this function. By default, only the initial user can view the data. Other users can use the data only after being granted with permissions.



The **gs\_read\_file\_size\_from\_remote** command is used to read the size of a specified file. This command is used to read the size of a specified file. Before using the **gs\_repair\_file** function to repair a file, you need to obtain the size of the file from the remote end to verify the missing file information and repair the missing files one by one. By default, only the initial user can view the data. Other users can use the data only after being granted with permissions.

- AI feature functions

|                         |                          |                           |                  |                          |                  |                  |
|-------------------------|--------------------------|---------------------------|------------------|--------------------------|------------------|------------------|
| create_snapshot         | create_snapshot_internal | prepare_snapshot_internal | prepare_snapshot | manage_snapshot_internal | archive_snapshot | publish_snapshot |
| purge_snapshot_internal | purge_snapshot           | sample_snapshot           | -                | -                        | -                | -                |

– Functions of **PKG\_SERVICE**

|                  |                 |   |   |   |   |   |
|------------------|-----------------|---|---|---|---|---|
| isubmit_on_nodes | submit_on_nodes | - | - | - | - | - |
|------------------|-----------------|---|---|---|---|---|

- Other functions

|                            |                               |                   |                             |                    |                   |                            |
|----------------------------|-------------------------------|-------------------|-----------------------------|--------------------|-------------------|----------------------------|
| to_tsvector_for_batch      | value_of_percentile           | disable_conn      | bind_variable               | job_update         | job_cancel        | job_finish                 |
| similar_escape             | table_skewness(unavailable)   | timetz_text       | time_text                   | reltime_text       | abstime_text      | _pg_keys_equal             |
| analyze_query(unavailable) | analyze_workload(unavailable) | ssign_table_type  | gs_comm_proxy_thread_status | gs_txid_oldestxmin | pg_cancel_session | pg_stat_segment_space_info |
| remote_segment_space_info  | set_cost_params               | set_weight_params | start_collect_workload      | tdigest_in         | tdigest_merge     | tdigest_merge_to_one       |
| tdigest_mergep             | tdigest_out                   | pg_get_delta_info | -                           | -                  | -                 | -                          |

- View-related reference functions  
adm\_hist\_sqlstat\_func

adm\_hist\_sqlstat\_idlog\_func

## 7.5.34 Global SysCache Feature Functions

- gs\_gsc\_table\_detail(database\_id default NULL, rel\_id default NULL)

Description: Queries global system cache table metadata in a database. The user who calls this function must have the SYSADMIN permission.

Parameter: Specifies the database and table whose global system cache is to be queried. The default value of **database\_id** is **NULL** or **-1**, indicating all databases. The value **0** indicates a shared table. Other values indicate the specified database and shared table. **rel\_id** indicates the OID of the specified table. The default value **NULL** or **-1** indicates all tables. Other values indicate the specified table. If the table does not exist, an error is reported. If the OID does not exist, a null result is returned.

Return type: Tuple

Example:

```
select * from gs_gsc_table_detail(-1) limit 1;
```

| database_oid | database_name | reloid | relname                 | relnamespace | reltype | reloftype | relowner | relam | relfilenode | reltablespace | relhasindex | relisshared | relkind | relnatts | relhasoids | relhaspkey | parttype | tdhasuids | attnames | extinfo |
|--------------|---------------|--------|-------------------------|--------------|---------|-----------|----------|-------|-------------|---------------|-------------|-------------|---------|----------|------------|------------|----------|-----------|----------|---------|
| 0            |               | 2676   | pg_authid_rolname_index | 11           | 0       | 0         | 10       | 403   |             | 0             |             |             |         |          |            |            |          |           |          |         |
|              | 1664          | f      | t                       | i            | 1       | f         | f        | n     | f           | 'rolname'     |             |             |         |          |            |            |          |           |          |         |

(1 row)

- gs\_gsc\_catalog\_detail(database\_id default NULL, rel\_id default NULL)

Description: Queries the system table row information cached in the global system in a database. The user who calls this function must have the SYSADMIN permission.

Parameter: Specifies the database and table whose global system cache is to be queried. The default value of **database\_id** is **NULL** or **-1**, indicating all databases. The value **0** indicates a shared table. Other values indicate the specified database and shared table. **rel\_id** indicates the ID of the specified table, including all system catalogs that have system caches. The default value **NULL** or **-1** indicates all tables. Other values indicate the specified table. If the database does not exist, an error is reported. If the table does not exist, a null result is returned.

Return type: Tuple

Example:

```
openGauss=#
select * from gs_gsc_catalog_detail(16574, 1260);
```

| database_id | database_name | rel_id | rel_name  | cache_id             | self  | ctid | infomask   | infomask2 | hash_value | refcount |
|-------------|---------------|--------|-----------|----------------------|-------|------|------------|-----------|------------|----------|
| 0           |               | 1260   | pg_authid | 10   (0, 9)   (0, 9) | 10507 | 26   | 531311568  |           |            |          |
| 10          |               | 1260   | pg_authid | 11   (0, 4)   (0, 4) | 2313  | 26   | 365368336  |           |            | 1        |
| 0           |               | 1260   | pg_authid | 11   (0, 9)   (0, 9) | 10507 | 26   | 3911517328 |           |            |          |
| 10          |               | 1260   | pg_authid | 11   (0, 7)   (0, 7) | 2313  | 26   | 1317799983 |           |            |          |
| 1           |               | 1260   | pg_authid | 11   (0, 5)   (0, 5) | 2313  | 26   | 3664347448 |           |            |          |
| 1           |               | 1260   | pg_authid | 11   (0, 1)   (0, 1) | 2313  | 26   | 276477273  |           |            | 1        |

|           |   |                  |                      |      |                 |   |
|-----------|---|------------------|----------------------|------|-----------------|---|
| 1         | 0 | 1260   pg_authid | 11   (0, 3)   (0, 3) | 2313 | 26   2465837659 |   |
| 1         | 0 | 1260   pg_authid | 11   (0, 8)   (0, 8) | 2313 | 26   3205288035 |   |
|           | 0 | 1260   pg_authid | 11   (0, 6)   (0, 6) | 2313 | 26   131811687  | 1 |
|           | 0 | 1260   pg_authid | 11   (0, 2)   (0, 2) | 2313 | 26   1226484587 |   |
| 1         |   |                  |                      |      |                 |   |
| (10 rows) |   |                  |                      |      |                 |   |

- gs\_gsc\_clean(database\_id default NULL)

Description: Clears the global syscache cache. Note that data in use will not be cleared. The user who calls this function must have the SYSADMIN permission.

Parameter: Specifies the database whose global system cache needs to be cleared. The default value **NULL** or **-1** indicates that the global system cache of all databases is cleared. The value **0** indicates that the global system cache of only the shared table is cleared. Other values indicate that the global system cache of the specified database and shared table is cleared. If **database\_id** does not exist, an error is reported.

Return type: Boolean

Example:

```
openGauss=# select * from gs_gsc_clean();
gs_gsc_clean

t
(1 row)
```

- gs\_gsc\_dbstat\_info(database\_id default NULL)

Description: Obtains GSC memory statistics on the local node, including cache query, hit, loading, expiration, and occupied space information of tuples, relationships, and partitions, database-level elimination information, thread reference information, and memory usage information. This parameter can be used to locate performance problems. For example, if the value of the hits/searches array is far less than 1, the value of **global\_syscache\_threshold** may be too small. As a result, the query hit ratio decreases. The user who calls this function must have the SYSADMIN permission.

Parameter: Specifies the global system cache statistics of the database to be queried. **NULL** or **-1** indicates that all databases are queried. **0** indicates that only the information about the shared table is queried. Other values indicate that the information about the specified database and shared table is queried. If an invalid value is entered, an error is reported, indicating that **database\_id** does not exist.

Return type: Tuple

Example:

```
openGauss=# select * from gs_gsc_dbstat_info();
database_id | database_name | tup_searches | tup_hits | tup_miss | tup_count | tup_dead | tup_memory |
rel_searches | rel_hits | rel_miss | rel_count | rel_dead | rel_memory | part_searches | part_hits | part_miss | part_count | part_dead |
part_memory | total_memory | swa_pout_count | refcount
-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | 300 | 235 | 31 | 22 | 2 | 9752 | 598 | 108 |
| 1 | 18 | 0 | 77720 | 0 | 0 | 0 | 0 | 0 | 0 | 752912 |
```

```
 0 | 0
 16574 | testdb | 3368 | 2289 | 329 | 273 | 0 | 92593 | 1113 |
524 | 4
8 | 48 | 0 | 340456 | 0 | 0 | 0 | 0 | 0 | 4124792 |
 0 | 10
(2 rows)
```

## 7.5.35 Data Damage Detection and Repair Functions

- `gs_verify_data_file(verify_segment bool)`

Description: Checks whether files in the current database of the current instance are lost. The verification only checks whether intermediate segments are lost in the main file of the data table. The default value is **false**, indicating that the segment-page table data file is not verified. If this parameter is set to **true**, only segment-page table files are checked. By default, only the initial user, users with the SYSADMIN permission, and users with the OPRADMIN permission in the O&M mode can view the information. Other users can view the information only after being granted permissions.

The returned result is as follows:

- Non-segment-page table: **rel\_oid** and **rel\_name** indicate the table OID and table name of the corresponding file, and **miss\_file\_path** indicates the relative path of the lost file.
- Segment-page table: All tables are stored in the same file. Therefore, **rel\_oid** and **rel\_name** cannot display information about a specific table. For a segment-page table, if the first file is damaged, the subsequent files such as .1 and .2 are not checked. For example, if 3, 3.1, and 3.2 are damaged, only 3 damage can be detected. When the number of segment-page files is less than 5, the files that are not generated are also detected during function detection. For example, if there are only files 1 and 2, files 3, 4, and 5 are detected during segment-page file detection. In the following examples, the first is an example of checking a non-segment-page table, and the second is an example of checking a segment-page table.

Parameter description:

- `verify_segment`

Specifies the range of files to be checked. **false** indicates that non-segment-page tables are verified. **true** indicates that segment-page tables are verified.

The value can be **true** or **false** (default value).

Return type: record.

Example (The abnormal line is displayed only when an exception is detected. Otherwise, no line is displayed.):

Check a non-segment-page table.

```
openGauss=# select * from gs_verify_data_file();
node_name | rel_oid | rel_name | miss_file_path
-----+-----+-----+-----
dn_6001_6002_6003 | 16554 | test | base/16552/24745
```

Verify a segment-page table.

```
openGauss=# select * from gs_verify_data_file(true);
node_name | rel_oid | rel_name | miss_file_path
-----+-----+-----+-----
dn_6001_6002_6003 | 0 | none | base/16573/2
```

- **gs\_repair\_file**(tableoid Oid, path text, timeout int)  
Description: Repairs the file based on the input parameters. Only the primary DN with normal primary/standby connection is supported. The parameter is set based on the OID and path returned by the **gs\_verify\_data\_file** function. The value of table OID for a segment-page table ranges from 0 to 4294967295. (The internal verification determines whether a file is a segment-page table file based on the file path. The table OID is not used for a segment-page table file.) If the repair is successful, **true** is returned. If the repair fails, the failure cause is displayed. By default, only the initial user, users with the SYSADMIN permission, and users with the OPRADMIN permission in the O&M mode on the primary DN can view the information. Other users can view the information only after being granted permissions.

---

**⚠ CAUTION**

1. If a file on a DN is damaged, a verification error at the PANIC level is reported when the DN is promoted to primary. The DN cannot be promoted to primary, which is normal.
2. If a file exists but its size is 0, the file will not be repaired. To repair the file, you need to delete the file whose size is 0 and then repair it.
3. You can delete a file only after the file descriptor is automatically closed. You can manually restart the process or perform a primary/standby switchover.

---

Parameter description:

- **tableoid**  
OID of the table corresponding to the file to be repaired. Set this parameter based on the **rel\_oid** column in the list returned by the **gs\_verify\_data\_file** function.  
Value range: OID ranging from 0 to 4294967295. Note: A negative value will be forcibly converted to a non-negative integer.
- **path**  
Path of the file to be repaired. Set this parameter based on the **miss\_file\_path** column in the list returned by the **gs\_verify\_data\_file** function.  
Value range: a string
- **timeout**  
Specifies the duration for waiting for the standby DN to replay. The repair file needs to wait for the standby DN to be put back to the corresponding location on the current primary DN. Set this parameter based on the replay duration of the standby DN.  
Value range: 60s to 3600s.

Return type: Boolean.

Example (Set **tableoid** and **path** based on the output of **gs\_verify\_data\_file**):

```
openGauss=# select * from gs_repair_file(16554,'base/16552/24745',360);
gs_repair_file

t
```

- `local_bad_block_info()`

Description: Displays the page damage of the instance. You can read the page from the disk and record the page CRC failure. By default, only the initial user, users with the SYSADMIN attribute, users with the MONADMIN attribute, users with the OPRADMIN attribute in the O&M mode, and monitoring users can view the information. Other users can view the information only after being granted permissions.

In the displayed information, **file\_path** indicates the relative path of the damaged file. If the table is a segment-page table, the logical information instead of the actual physical file information is displayed. **block\_num** indicates the number of the page where the file is damaged. The page number starts from 0. **check\_time** indicates the time when the page damage is detected. **repair\_time** indicates the time when the page is repaired.

Return type: record.

Example (Related entries are displayed only when there are damaged records. Otherwise, no log is displayed.):

```
openGauss=# select * from local_bad_block_info();
node_name | spc_node | db_node | rel_node | bucket_node | fork_num | block_num | file_path |
check_time | repair_time
-----+-----+-----+-----+-----+-----+-----+-----+
dn_6001_6002_6003 | 1663 | 16552 | 24745 | -1 | 0 | 0 | base/16552/24745 |
2022-01-13 20:19:08.385004+08 | 2022-01-13 20:19:08.407314+08
```

- `local_clear_bad_block_info()`

Description: Deletes data of repaired pages from **local\_bad\_block\_info**, that is, information whose **repair\_time** is not empty. By default, only the initial user, users with the SYSADMIN attribute, users with the OPRADMIN attribute in the O&M mode, and monitoring users can view the information. Other users can view the information only after being granted permissions.

Return type: Boolean.

Example:

```
openGauss=# select * from local_clear_bad_block_info();
result

t
```

- `gs_verify_and_tryrepair_page` (path text, blocknum oid, verify\_mem bool, is\_segment bool)

Description: Verifies the page specified by the instance. By default, only the initial user, users with the SYSADMIN attribute, and users with the OPRADMIN attribute in O&M mode on the primary DN can view the table. Other users can view the table only after being granted permissions.

In the command output, **disk\_page\_res** indicates the verification result of the page on the disk, **mem\_page\_res** indicates the verification result of the page in the memory, and **is\_repair** indicates whether the repair function is triggered during the verification. **t** indicates that the page is repaired, and **f** indicates that the page is not repaired.

Note:

1. If a page on a DN is damaged, a verification error at the PANIC level is reported when the DN is promoted to primary. The DN cannot be promoted to primary, which is normal. Damaged pages of hash bucket tables cannot be repaired.

2. The repair triggered by this function can only repair pages in the memory. The repair takes effect only after the memory pages are flushed to disks.

Parameters:

- path  
Path of the damaged file. Set this parameter based on the **file\_path** column in the **local\_bad\_block\_info** file. To verify the undo pages of the Ustore table, enter the path of the undo pages to be verified.  
Value range: a string
- blocknum  
Page number of the damaged file. Set this parameter based on the **block\_num** column in the **local\_bad\_block\_info** file. If you want to verify the undo pages of the Ustore table, enter the block number of the undo pages to be verified.  
Value range: OID ranging from 0 to 4294967295. Note: A negative value will be forcibly converted to a non-negative integer.
- verify\_mem  
Specifies whether to verify a specified page in the memory. If this parameter is set to **false**, only pages on the disk are verified. If this parameter is set to **true**, pages in the memory and on the disk are verified. If a page on the disk is damaged, the system verifies the basic information of the page in the memory and flushes the page to the disk to restore the page. If a page is not found in the memory during memory page verification, the page on the disk is read through the memory API. During this process, if the disk page is faulty, the remote read automatic repair function is triggered.  
Value range: The value is of a Boolean type and can be **true** or **false**.
- is\_segment  
Determines whether the table is a segment-page table. Set this parameter based on the value of **bucket\_node** in the **local\_bad\_block\_info** file. If the value of **bucket\_node** is **-1**, the table is not a segment-page table. In this case, set **is\_segment** to **false**. If the value of **bucket\_node** is not **-1**, set **is\_segment** to **true**.  
Value range: The value is of a Boolean type and can be **true** or **false**.

Return type: record.

Examples (Set parameters based on the output of **local\_bad\_block\_info**. Otherwise, an error is reported.):

```
openGauss=# select * from gs_verify_and_tryrepair_page('base/16552/24745',0,false,false);
node_name | path | blocknum | disk_page_res | mem_page_res | is_repair
-----+-----+-----+-----+-----+-----
dn_6001_6002_6003 | base/16552/24745 | 0 | page verification succeeded. | | f
```

- **gs\_repair\_page**(path text, blocknum oid, is\_segment bool, timeout int)  
Description: Restores the specified page of the instance. This function can be used only by the primary DN that is properly connected to the primary and standby DNs. If the page is successfully restored, **true** is returned. If an error occurs during the restoration, an error message is displayed. By default, only the initial user, users with the SYSADMIN attribute, and users with the OPRADMIN attribute in O&M mode on the primary DN can view the table. Other users can view the table only after being granted permissions.

Note: If a page on a DN is damaged, a verification error at the PANIC level is reported when the DN is promoted to primary. The DN cannot be promoted to primary, which is normal. Damaged pages of hash bucket tables cannot be repaired.

Parameter description:

- path  
Path of the damaged page. Set this parameter based on the **file\_path** column in **local\_bad\_block\_info** or the **path** column in the **gs\_verify\_and\_tryrepair\_page** function.  
Value range: a string
- blocknum  
Number of the damaged page. Set this parameter based on the **block\_num** column in **local\_bad\_block\_info** or the **blocknum** column in the **gs\_verify\_and\_tryrepair\_page** function.  
Value range: OID ranging from 0 to 4294967295. Note: A negative value will be forcibly converted to a non-negative integer.
- is\_segment  
Determines whether the table is a segment-page table. The value of this parameter is determined by the value of **bucket\_node** in **local\_bad\_block\_info**. If the value of **bucket\_node** is **-1**, the table is not a segment-page table and **is\_segment** is set to **false**. If the value of **bucket\_node** is not **-1**, **is\_segment** is set to **true**.  
Value range: The value is of a Boolean type and can be **true** or **false**.
- timeout  
Duration of waiting for standby DN replay. The page to be repaired needs to wait for the standby DN to be replayed to the location of the current primary DN. Set this parameter based on the replay duration of the standby DN.  
Value range: 60s to 3600s.

Return type: Boolean.

Examples (Set parameters based on the output of **local\_bad\_block\_info**. Otherwise, an error is reported.):

```
openGauss=# select * from gs_repair_page('base/16552/24745',0,false,60);
result

t
```

- **gs\_verify\_urq**(index\_oid oid, queue\_type text, blocknum bigint, verify\_type text)

Description: Verifies the correctness of the index recycling queue or the performance of obtaining index pages from the recycling queue.

Parameter description: See [Table 7-85](#).

Return type: record.



**Table 7-85** gs\_verify\_urq parameters

| Category         | Parameter   | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------|-------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Input parameter  | index_oid   | oid    | UB-tree index OID.                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Input parameter  | queue_type  | text   | Specifies the queue type: <ul style="list-style-type: none"><li>• <b>empty queue</b>: potential queue</li><li>• <b>free queue</b>: available queue</li><li>• <b>single page</b>: single page of the queue</li></ul>                                                                                                                                                                                                                                    |
| Input parameter  | blocknum    | bigint | Specifies the page number: <ul style="list-style-type: none"><li>• If the queue type is single page, all tuples of <b>blocknum</b> on a single page are verified. The value range is [0,Queue file size/8192).</li><li>• If the queue type is empty queue or free queue and <b>blocknum</b> is not set to <b>0</b>, all tuples on all pages of this queue are verified. If <b>blocknum</b> is set to <b>0</b>, page tuples are not verified.</li></ul> |
| Input parameter  | verify_type | text   | Specifies the verification type: <ul style="list-style-type: none"><li>• <b>physics</b>: verifies the physical structure of a queue.</li><li>• <b>performance</b>: verifies the performance of obtaining pages from the recycling queue.</li></ul>                                                                                                                                                                                                     |
| Output parameter | verify_code | text   | Error code                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Output parameter | detail      | text   | Error description                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Example:**

```
openGauss=# select * from gs_verify_urq(16387,'free queue',1,'physics');
verify_code | detail
-----+-----
```

 **NOTE**

Currently, this API only supports Ustore index tables and does not support partition local indexes.

- gs\_repair\_urq(index\_oid oid)

Description: Rebuilds an index to recycle queues (potential and available queues). If the rebuild is successful, **reinitial the recycle queue of index relation successfully** is displayed.

Parameter description: index\_oid (UB-tree index OID)

Return type: text

Note: The current function can be called only on the primary node.

Example:

```
openGauss=# select * from gs_repair_urq(16387);
 result

reinitial the recycle queue of index relation successfully.
```

#### NOTE

Currently, this API only supports Ustore index tables and does not support partition local indexes.

## 7.5.36 Other System Functions

Other system functions are classified into PG-compatible functions and internal functions. These functions are not recommended. If you need to use them, contact Huawei technical support.

### PG-compatible Functions and Operators

The following table lists the built-in functions and operators of GaussDB that are compatible with PG.

|                             |                       |                        |                 |                    |                   |                       |
|-----------------------------|-----------------------|------------------------|-----------------|--------------------|-------------------|-----------------------|
| _pg_char_max_length         | _pg_char_octet_length | _pg_datetime_precision | _pg_expandarray | _pg_index_position | _pg_interval_type | _pg_numeric_precision |
| _pg_numeric_precision_radix | _pg_numeric_scale     | _pg_truetypid          | _pg_truetypmod  | abbrev             | abs               | abstime               |
| abstimeeq                   | abstimege             | abstimegt              | abstimein       | abstimele          | abstimeelt        | abstimene             |
| abstimeout                  | abstimerectv          | abstimesend            | aclcontains     | acldefault         | aclexplode        | aclinsert             |
| aclitemeq                   | aclitemin             | aclitemout             | aclremove       | acos               | age               | akeys                 |
| any_in                      | any_out               | anyarray_in            | anyarray_out    | anyarray_recv      | anyarray_send     | anyelement_in         |
| anyelement_out              | anyenum_in            | anyenum_out            | anynonarray_in  | anynonarray_out    | anyrange_in       | anyrange_out          |

|                   |                    |                   |                  |                    |                   |                   |
|-------------------|--------------------|-------------------|------------------|--------------------|-------------------|-------------------|
| anytextcat        | area               | areajoinse        | areasel          | array_agg          | array_agg_finalfn | array_agg_transfn |
| array_append      | array_cat          | array_dims        | array_eq         | array_fill         | array_ge          | array_gt          |
| array_in          | array_larger       | array_length      | array_length     | array_lower        | array_lt          | array_ndims       |
| array_ne          | array_out          | array_prepend     | array_recv       | array_send         | array_smaller     | array_to_json     |
| array_to_string   | array_type_analyze | array_upper       | arraycontained   | arraycontains      | arraycontainsjoin | arraycontainsjoin |
| arrayoverlap      | ascii              | asin              | atan             | atan2              | avals             | avg               |
| big5_to_euc_tw    | big5_to_mic        | big5_to_utf8      | bit              | bit_and            | bit_in            | bit_length        |
| bit_or            | bit_out            | bit_recv          | bit_send         | bitand             | bitcat            | bitcmp            |
| biteq             | bitge              | bitgt             | bitle            | bitlt              | bitne             | bitnot            |
| bitor             | bitshiftleft       | bitshiftright     | bittypmodin      | bittypmodout       | bitxor            | bool              |
| bool_and          | bool_or            | booland_statefunc | booleq           | boolge             | boolgt            | boolin            |
| boolle            | boollt             | boolne            | boolor_statefunc | boolout            | boolrecv          | boolsend          |
| box               | box_above          | box_above_eq      | box_adj          | box_below          | box_below_eq      | box_center        |
| box_contain       | box_contain_pt     | box_contained     | box_distance     | box_div            | box_eq            | box_ge            |
| box_gt            | box_in             | box_intersect     | box_le           | box_left           | box_lt            | box_mul           |
| box_out           | box_overabove      | box_overbelow     | box_overlap      | box_overleft       | box_overright     | box_recv          |
| box_right         | box_same           | box_send          | box_sub          | bpchar             | bpchar_larger     | bpchar_pattern_ge |
| bpchar_pattern_gt | bpchar_pattern_le  | bpchar_pattern_lt | bpchar_smaller   | bpchar_sortsupport | bpcharcmp         | bpchareq          |

|                 |                   |                     |                          |                   |                 |                     |
|-----------------|-------------------|---------------------|--------------------------|-------------------|-----------------|---------------------|
| bpcharge        | bpcharget         | bpchariclike        | bpcharicnlike            | bpcharicregexe    | bpcharicregexe  | bpcharin            |
| bpcharle        | bpcharlike        | bpcharlt            | bpcharne                 | bpcharnlike       | bpcharout       | bpcharrecv          |
| bpcharregexe    | bpcharregexe      | bpcharend           | bpchartypmodin           | bpchartypmodout   | broadcast       | btabsstimecmp       |
| btarraycmp      | btbeginscan       | btboolcmp           | btbpchar_pattern_cmp     | btbuild           | btbuildempty    | btbulkdelete        |
| btcanreturn     | btcharcmp         | btcostestimate      | btendscan                | btfloat48cmp      | btfloat4cmp     | btfloat4sortsupport |
| btfloat84cmp    | btfloat8cmp       | btfloat8sortsupport | btgetbitmap              | btgettuple        | btinsert        | btint24cmp          |
| btint28cmp      | btint2cmp         | btint2sortsupport   | btint42cmp               | btint48cmp        | btint4cmp       | btint4sortsupport   |
| btint82cmp      | btint84cmp        | btint8cmp           | btint8sortsupport        | btmarkpos         | btnamecmp       | btnamesortsupport   |
| btoidcmp        | btoidsortsupport  | btoidvectorcmp      | btoptions                | btrecordcmp       | btreltimecmp    | btrescan            |
| btrestnpos      | btrim             | bttext_pattern_cmp  | bttextcmp                | bttextsortsupport | bttidcmp        | btintervalcmp       |
| btvacuumcleanup | bytea_sortsupport | bytea_string_agg_fn | bytea_string_agg_transfn | byteacat          | byteacmp        | byteaeq             |
| byteage         | byteagt           | byteain             | byteale                  | bytealike         | bytealt         | byteane             |
| byteanlike      | byteaout          | bytearecv           | byteasend                | cash_cmp          | cash_div_cash   | cash_div_float4     |
| cash_div_float8 | cash_div_int2     | cash_div_int4       | cash_div_int8            | cash_eq           | cash_ge         | cash_gt             |
| cash_in         | cash_le           | cash_lt             | cash_mi                  | cash_mul_float4   | cash_mul_float8 | cash_mul_int2       |

|                      |                   |                               |                       |                  |                   |                     |
|----------------------|-------------------|-------------------------------|-----------------------|------------------|-------------------|---------------------|
| cash_mul_int4        | cash_mul_int8     | cash_ne                       | cash_out              | cash_pl          | cash_recv         | cash_send           |
| cashlarger           | cashsmaller       | cbrt                          | ceil                  | ceiling          | center            | char                |
| char_length          | character_length  | chareq                        | charge                | charget          | charin            | charle              |
| charlt               | charne            | charout                       | charrecv              | charsend         | chr               | cideq               |
| cidin                | cidout            | cidr                          | cidr_in               | cidr_out         | cidr_recv         | cidr_send           |
| cidrecv              | cidsend           | circle                        | circle_above          | circle_add_pt    | circle_below      | circle_center       |
| circle_contain       | circle_contain_pt | circle_contained              | circle_distance       | circle_dv_pt     | circle_eq         | circle_ge           |
| circle_gt            | circle_in         | circle_le                     | circle_left           | circle_lt        | circle_mul_pt     | circle_ne           |
| circle_out           | circle_overabove  | circle_overbelow              | circle_overlap        | circle_overleft  | circle_overright  | circle_recv         |
| circle_right         | circle_same       | circle_sen                    | circle_sub_pt         | clock_timestamp  | close_lb          | close_ls            |
| close_lseg           | close_pb          | close_pl                      | close_ps              | close_sb         | close_sl          | col_description     |
| concat               | concat_ws         | contjoinsel                   | contsel               | convert          | convert_from      | convert_to          |
| corr                 | cos               | cot                           | count                 | covar_pop        | covar_samp        | cstring_in          |
| cstring_out          | cstring_recv      | cstring_send                  | cume_dist             | current_database | current_query     | current_schema      |
| xpath_exists         | current_setting   | current_user                  | currtid               | currtid2         | currval           | cursor_to_xml       |
| cursor_to_xmlschema  | database_to_xml   | database_to_xml_and_xmlschema | database_to_xmlschema | date             | date_cmp          | date_cmp_timestamp  |
| date_cmp_timestamptz | date_eq           | date_eq_timestamp             | date_eq_timestampz    | date_ge          | date_ge_timestamp | date_ge_timestamptz |

|                         |                       |                         |                             |                                            |                                   |                       |
|-------------------------|-----------------------|-------------------------|-----------------------------|--------------------------------------------|-----------------------------------|-----------------------|
| date_gt                 | date_gt_timestam<br>p | date_gt_timestam<br>ptz | date_in                     | date_lar<br>ger                            | date_l<br>e                       | date_le_tim<br>estamp |
| date_le_tim<br>estamptz | date_lt               | date_lt_timestam<br>p   | date_lt_timest<br>amptz     | date_mi                                    | date_mi_int<br>erval              | date_mii              |
| date_ne                 | date_ne_timesta<br>mp | date_ne_timestam<br>ptz | date_o<br>ut                | date_pl<br>interval                        | date_p<br>li                      | date_recv             |
| date_send               | date_sm<br>aller      | date_sort<br>support    | datera<br>nge_ca<br>nonical | dateran<br>ge_subd<br>iff                  | dateti<br>me_pl                   | datetimetz_<br>pl     |
| dcbt                    | decode                | defined                 | degree<br>s                 | delete                                     | dense_<br>rank                    | dexp                  |
| diagonal                | diameter              | dispell_ini<br>t        | dispell_<br>lexize          | dist_cpo<br>ly                             | dist_lb                           | dist_pb               |
| dist_pc                 | dist_pl               | dist_ppat<br>h          | dist_ps                     | dist_sb                                    | dist_sl                           | div                   |
| dlog1                   | dlog10                | domain_i<br>n           | domai<br>n_recv             | dpow                                       | droun<br>d                        | dsimple_init          |
| dsimple_lexi<br>ze      | dsnowba<br>ll_init    | dsnowbal<br>l_lexize    | dsqrt                       | dsynony<br>m_init                          | dsyno<br>nym_l<br>exize           | dtrunc                |
| each                    | enum_ne               | enum_ou<br>t            | enum_<br>range              | enum_r<br>ecv                              | enum_<br>send                     | enum_small<br>er      |
| eqjoinsel               | eqsel                 | euc_cn_to<br>_mic       | euc_cn<br>_to_utf<br>8      | euc_jis_<br>2004_to<br>_shift_jis<br>_2004 | euc_jis_<br>_2004_<br>to_utf<br>8 | euc_jp_to_m<br>ic     |
| euc_jp_to_sji<br>s      | euc_jp_to<br>_utf8    | euc_kr_to<br>_mic       | euc_kr<br>to_utf8           | euc_tw_<br>to_big5                         | euc_tw_<br>to_mi<br>c             | euc_tw_to_u<br>tf8    |
| every                   | exist                 | exists_all              | exists_<br>any              | exp                                        | factori<br>al                     | family                |
| fdw_handler<br>_in      | fdw_han<br>dler_out   | fetchval                | first_va<br>lue             | float4                                     | float4_<br>accum                  | float48div            |
| float48eq               | float48g<br>e         | float48gt               | float48<br>le               | float48l<br>t                              | float4<br>8mi                     | float48mul            |
| float48ne               | float48pl             | float4abs               | float4d<br>iv               | float4eq                                   | float4<br>ge                      | float4gt              |

|                     |                       |                        |                           |                        |                         |                       |
|---------------------|-----------------------|------------------------|---------------------------|------------------------|-------------------------|-----------------------|
| float4in            | float4larger          | float4le               | float4lt                  | float4mi               | float4mul               | float4ne              |
| float4out           | float4pl              | float4recv             | float4send                | float4smaller          | float4um                | float4up              |
| float8              | float8_accum          | float8_avg             | float8_collect            | float8_corr            | float8_covar_pop        | float8_covar_samp     |
| float8_regr_accum   | float8_regr_avgx      | float8_regr_avgy       | float8_regr_collect       | float8_regr_intercept  | float8_regr_r2          | float8_regr_slope     |
| float8_regr_sxx     | float8_regr_sxy       | float8_regr_syy        | float8_stddev_pop         | float8_stddev_samp     | float8_var_pop          | float8_var_samp       |
| float84div          | float84eq             | float84ge              | float84gt                 | float84le              | float84lt               | float84mi             |
| float84mul          | float84ne             | float84pl              | float8abs                 | float8div              | float8eq                | float8ge              |
| float8gt            | float8in              | float8larger           | float8le                  | float8lt               | float8mi                | float8mul             |
| float8ne            | float8out             | float8pl               | float8recv                | float8send             | float8smaller           | float8um              |
| float8up            | floor                 | flt4_mul_cash          | flt8_mul_cash             | fmgr_validator         | fmgr_internal_validator | fmgr_sql_validator    |
| format              | format_type           | gb18030_to_utf8        | gbk_to_utf8               | generate_series        | generate_subscripts     | get_bit               |
| get_byte            | get_current_ts_config | -                      | -                         | gin_clean_pending_list | gin_cmp_prefix          | gin_cmp_tsl_exeme     |
| gin_extract_tsquery | gin_extract_tsvector  | gin_tsquery_consistent | gin_tsquery_triconsistent | ginarray_consistent    | ginarrayextract         | ginarraytriconsistent |
| ginbeginscan        | ginbuild              | ginbuildempty          | ginbulkdelete             | gincostestimate        | ginendscan              | gingetbitmap          |
| gininsert           | ginmarkpos            | ginoptions             | ginqueryarrayextract      | ginrescan              | ginrestpos              | ginvacuumcleanup      |

|                      |                        |                          |                       |                     |                    |                      |
|----------------------|------------------------|--------------------------|-----------------------|---------------------|--------------------|----------------------|
| gist_box_compress    | gist_box_consistent    | gist_box_decompress      | gist_box_penalty      | gist_box_picksplit  | gist_box_same      | gist_box_union       |
| gist_circle_compress | gist_circle_consistent | gist_point_compress      | gist_point_consistent | gist_point_distance | gist_poly_compress | gist_poly_consistent |
| gistbeginscan        | gistbuild              | gistbuildempty           | gistbulkdelete        | gistcostestimate    | gistendscan        | gistgetbitmap        |
| gistgettupple        | gistinsert             | gistmarkpos              | gistoptions           | gistrescan          | gistresstrpos      | gistvacuumcleanup    |
| gtsquery_compress    | gtsquery_consistent    | gtsquery_decompress      | gtsquery_penalty      | gtsquery_picksplit  | gtsquery_same      | gtsquery_union       |
| gtsvector_compress   | gtsvector_consistent   | gtsvector_decompress     | gtsvector_penalty     | gtsvector_picksplit | gtsvector_same     | gtsvector_union      |
| gtsvectorin          | gtsvectorout           | has_tablespace_privilege | has_type_privilege    | hash_aclitem        | hashbeginscan      | hashbuild            |
| hashbuildempty       | hashbulkdelete         | hashcostestimate         | hashendscan           | hashgetbitmap       | hashgettupple      | hashinsert           |
| hashint2vector       | hashint4               | hashint8                 | hashmarkaddr          | hashmarkpos         | hashname           | hashoid              |
| hashoidvector        | hashoptions            | hashrescan               | hashrestrpos          | hashtext            | hashvacuumcleanup  | hashvarlena          |
| host                 | hostmask               | iclikejoinsel            | iclikejoin            | icnlikejoinsel      | icnlikejoin        | icregexejoin         |
| icregexejoin         | icregexejoin           | icregexejoin             | inet_client_addr      | inet_client_port    | inet_in            | inet_out             |
| inet_recv            | inet_send              | inet_server_addr         | inet_server_port      | inetand             | inetmi             | inetmi_int8          |
| inetnot              | inetor                 | inetpl                   | initcap               | int2_accum          | int2_avg_accum     | int2_mul_cash        |
| int2_sum             | int24div               | int24eq                  | int24ge               | int24gt             | int24le            | int24lt              |
| int24mi              | int24mul               | int24ne                  | int24pl               | int28div            | int28eq            | int28ge              |



|                  |                |             |                       |                  |                     |                   |
|------------------|----------------|-------------|-----------------------|------------------|---------------------|-------------------|
| int28gt          | int28le        | int28lt     | int28mi               | int28mul         | int28ne             | int28pl           |
| int2abs          | int2and        | int2div     | int2eq                | int2ge           | int2gt              | int2in            |
| int2larger       | int2le         | int2lt      | int2mi                | int2mod          | int2mul             | int2ne            |
| int2not          | int2or         | int2out     | int2pl                | int2recv         | int2send            | int2shl           |
| int2shr          | int2smaller    | int2um      | int2up                | int2vectorreq    | int2vectorin        | int2vectorout     |
| int2vectorrecv   | int2vectorsend | int2xor     | int4_accum            | int4_avg_accum   | int4_mul_cash       | int4_sum          |
| int42div         | int42eq        | int42ge     | int42gt               | int42le          | int42lt             | int42mi           |
| int42mul         | int42ne        | int42pl     | int48div              | int48eq          | int48ge             | int48gt           |
| int48le          | int48lt        | int48mi     | int48mul              | int48ne          | int48pl             | int4abs           |
| int4and          | int4div        | int4eq      | int4ge                | int4gt           | int4in              | int4inc           |
| int4larger       | int4le         | int4lt      | int4mi                | int4mod          | int4mul             | int4ne            |
| int4not          | int4or         | int4out     | int4pl                | int4range        | int4range_canonical | int4range_subdiff |
| int4recv         | int4send       | int4shl     | int4shr               | int4smaller      | int4um              | int4up            |
| int4xor          | int8           | int8_avg    | int8_avg_accum        | int8_avg_collect | int8_mul_cash       | int8_sum          |
| int8_sum_to_int8 | int8_accum     | int82div    | int82eq               | int82ge          | int82gt             | int82le           |
| int82lt          | int82mi        | int82mul    | int82ne               | int82pl          | int84div            | int84eq           |
| int84ge          | int84gt        | int84le     | int84lt               | int84mi          | int84mul            | int84ne           |
| int84pl          | int8abs        | int8and     | int8div               | int8eq           | int8ge              | int8gt            |
| int8in           | int8inc        | int8inc_any | int8inc_float8_float8 | int8larger       | int8le              | int8lt            |

|                          |                       |                      |                             |                            |                                   |                             |
|--------------------------|-----------------------|----------------------|-----------------------------|----------------------------|-----------------------------------|-----------------------------|
| int8mi                   | int8mod               | int8mul              | int8ne                      | int8not                    | int8or                            | int8out                     |
| int8pl                   | int8pl_in<br>et       | int8range            | int8ran<br>ge_can<br>onical | int8rang<br>e_subdif<br>f  | int8rec<br>v                      | int8send                    |
| int8shl                  | int8shr               | int8small<br>er      | int8um                      | int8up                     | int8xor                           | integer_pl_d<br>ate         |
| inter_lb                 | inter_sb              | inter_sl             | interna<br>l_in             | internal<br>_out           | interva<br>l                      | interval_acc<br>um          |
| interval_avg             | interval_<br>cmp      | interval_c<br>ollect | interval<br>_div            | interval<br>_eq            | interva<br>l_ge                   | interval_gt                 |
| interval_has<br>h        | interval_i<br>n       | interval_l<br>arger  | interval<br>_le             | interval<br>_lt            | interva<br>l_mi                   | interval_mul                |
| interval_ne              | interval_<br>out      | interval_p<br>l      | interval<br>_pl_dat<br>e    | interval<br>_pl_time       | interva<br>l_pl_ti<br>mesta<br>mp | interval_pl_t<br>imestamptz |
| interval_pl_t<br>imetz   | interval_r<br>ecv     | interval_s<br>end    | interval<br>_smalle<br>r    | interval<br>_transfo<br>rm | interva<br>l_um                   | intervaltyp<br>modin        |
| intervaltyp<br>modout    | intinterv<br>al       | isexists             | ishoriz<br>ontal            | iso_to_k<br>oi8r           | iso_to_<br>mic                    | iso_to_win1<br>251          |
| iso_to_win8<br>66        | iso8859_<br>1_to_utf8 | iso8859_t<br>o_utf8  | isparall<br>el              | isperp                     | isvertic<br>al                    | johab_to_ut<br>f8           |
| jsonb_in                 | jsonb_ou<br>t         | jsonb_rec<br>v       | jsonb_s<br>end              | -                          | -                                 | -                           |
| json_in                  | json_out              | json_rcv             | json_se<br>nd               | justify_d<br>ays           | justify<br>_hours                 | justify_inter<br>val        |
| koi8r_to_iso             | koi8r_to_<br>mic      | koi8r_to_<br>utf8    | koi8r_t<br>o_win1<br>251    | koi8r_to<br>_win866        | koi8u_<br>to_utf<br>8             | language_h<br>andler_in     |
| language_h<br>andler_out | latin1_to<br>_mic     | latin2_to_<br>mic    | latin2_t<br>o_win1<br>250   | latin3_t<br>o_mic          | latin4_<br>to_mic                 | like_escape                 |
| likejoinsel              | likesel               | line                 | line_dis<br>tance           | line_eq                    | line_h<br>orizont<br>al           | line_in                     |
| line_interpt             | line_inter<br>sect    | line_out             | line_pa<br>rallel           | line_per<br>p              | line_re<br>cv                     | line_send                   |
| line_vertical            | ln                    | lo_close             | lo_crea<br>t                | lo_creat<br>e              | lo_exp<br>ort                     | lo_import                   |

|                |                                                                                                                                                          |               |               |                 |               |                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|---------------|-----------------|---------------|----------------|
| lo_lseek       | lo_open                                                                                                                                                  | lo_tell       | lo_truncate   | lo_unlink       | log           | loread         |
| lower          | lower_inc                                                                                                                                                | lower_inf     | lowrite       | lpad            | lseg          | lseg_center    |
| lseg_distance  | lseg_eq                                                                                                                                                  | lseg_ge       | lseg_gt       | lseg_horizontal | lseg_in       | lseg_interpt   |
| lseg_intersect | lseg_le                                                                                                                                                  | lseg_length   | lseg_lt       | lseg_ne         | lseg_out      | lseg_parallel  |
| lseg_perp      | lseg_rcv                                                                                                                                                 | lseg_send     | lseg_vertical | ltrim           | macaddr_and   | macaddr_cmp    |
| macaddr_eq     | macaddr_ge                                                                                                                                               | macaddr_gt    | macaddr_in    | macaddr_le      | macaddr_lt    | macaddr_ne     |
| macaddr_not    | macaddr_or                                                                                                                                               | macaddr_out   | macaddr_rcv   | macaddr_send    | makeaclitem   | masklen        |
| max            | md5<br>(The MD5 encryption algorithm has lower security and poses security risks. Therefore, you are advised to use a more secure encryption algorithm.) | mic_to_big5   | mic_to_euc_cn | mic_to_euc_jp   | mic_to_euc_kr | mic_to_euc_tw  |
| mic_to_iso     | mic_to_koi8r                                                                                                                                             | mic_to_latin1 | mic_to_latin2 | mic_to_latin3   | mic_to_latin4 | mic_to_sjis    |
| mic_to_win1250 | mic_to_win1251                                                                                                                                           | mic_to_win866 | min           | mktinterval     | money         | mul_d_interval |
| name           | nameeq                                                                                                                                                   | namege        | namegt        | nameiclike      | nameicnlike   | nameicregexeq  |

|                       |                     |                             |                         |                            |                                    |                            |
|-----------------------|---------------------|-----------------------------|-------------------------|----------------------------|------------------------------------|----------------------------|
| nameicrege<br>xne     | namein              | namele                      | nameli<br>ke            | namelt                     | namen<br>e                         | namenlike                  |
| nameout               | namerec<br>v        | namereg<br>exeq             | namer<br>egexne         | name<br>se<br>nd           | neqjoi<br>nsel                     | neqsel                     |
| network_cm<br>p       | network_<br>eq      | network_<br>ge              | networ<br>k_gt          | network<br>_le             | netwo<br>rk_lt                     | network_ne                 |
| network_su<br>b       | network_<br>subeq   | network_<br>sup             | networ<br>k_supe<br>q   | nlikejoin<br>sel           | nlike<br>se<br>l                   | numeric                    |
| numeric_ab<br>s       | numeric_<br>accum   | numeric_<br>add             | numeri<br>c_avg         | numeric<br>_avg_ac<br>cum  | numeri<br>c_avg_<br>collec<br>t    | numeric_cm<br>p            |
| numeric_col<br>lect   | numeric_<br>div     | numeric_<br>div_trunc       | numeri<br>c_eq          | numeric<br>_exp            | numeri<br>c_fac                    | numeric_ge                 |
| numeric_gt            | numeric_<br>in      | numeric_<br>inc             | numeri<br>c_large<br>r  | numeric<br>_le             | numeri<br>c_ln                     | numeric_log                |
| numeric_lt            | numeric_<br>mod     | numeric_<br>mul             | numeri<br>c_ne          | numeric<br>_out            | numeri<br>c_pow<br>er              | numeric_rec<br>v           |
| numeric_se<br>nd      | numeric_<br>smaller | numeric_<br>sortsuppo<br>rt | numeri<br>c_sqrt        | numeric<br>_stddev<br>_pop | numeri<br>c_stddev_<br>sa<br>mp    | numeric_su<br>b            |
| numeric_tra<br>nsform | numeric_<br>uminus  | numeric_<br>uplus           | numeri<br>c_var_p<br>op | numeric<br>_var_sa<br>mp   | numeri<br>c_ttypm<br>odin          | numeri<br>c_ttyp<br>modout |
| numrange_s<br>ubdiff  | oid                 | oideq                       | oidge                   | oidgt                      | oidin                              | oidlarger                  |
| oidle                 | oidlt               | oidne                       | oidout                  | oidrecv                    | oidsen<br>d                        | oidsmaller                 |
| oidvectoreq           | oidvector<br>ge     | oidvector<br>gt             | oidvect<br>orin         | oidvecto<br>rle            | oidvec<br>torlt                    | oidvectorne                |
| oidvectorou<br>t      | oidvector<br>recv   | oidvector<br>send           | oidvect<br>ortypes      | on_pb                      | on_pl                              | on_ppath                   |
| on_ps                 | on_sb               | on_sl                       | opaque<br>_in           | opaque<br>_out             | ordere<br>d_set_t<br>ransiti<br>on | overlaps                   |

|                                  |                                           |                                   |                                |                                |                             |                                      |
|----------------------------------|-------------------------------------------|-----------------------------------|--------------------------------|--------------------------------|-----------------------------|--------------------------------------|
| overlay                          | path                                      | path_add                          | path_add_pt                    | path_center                    | path_contains_pt            | path_distance                        |
| path_div_pt                      | path_in                                   | path_inter                        | path_length                    | path_mul_pt                    | path_neq                    | path_nge                             |
| path_n_gt                        | path_n_le                                 | path_n_lt                         | path_n_points                  | path_out                       | path_recv                   | path_send                            |
| path_sub_pt                      | percentile_cont                           | percentile_cont_float8_final      | percentile_cont_interval_final | pg_char_to_encoding            | pg_cursor                   | pg_encoding_max_length               |
| pg_encoding_to_char              | -                                         | -                                 | -                              | pg_node_tree_in                | pg_node_tree_out            | pg_node_tree_recv                    |
| pg_node_tree_send                | pg_prepared_statement                     | pg_prepared_xact                  | -                              | -                              | pg_show_all_settings        | pg_stat_get_bgwriter_stat_reset_time |
| pg_stat_get_buf_fsync_backend    | pg_stat_get_checkpoint_sync_time          | pg_stat_get_checkpoint_write_time | pg_stat_get_dbb_blk_read_time  | pg_stat_get_dbb_blk_write_time | pg_stat_get_db_conflict_all | pg_stat_get_db_conflict_bufferpin    |
| pg_stat_get_db_conflict_snapshot | pg_stat_get_db_conflict_start_up_deadlock | pg_switch_xlog                    | xpath                          | pg_timezone_abbrs              | pg_timezone_names           | pg_stat_get_wal_receiver             |
| plpgsql_call_handler             | plpgsql_inline_handler                    | plpgsql_validator                 | point_above                    | point_add                      | point_below                 | point_distance                       |
| point_div                        | point_eq                                  | point_horiz                       | point_in                       | point_left                     | point_mul                   | point_ne                             |
| point_out                        | point_recv                                | point_right                       | point_send                     | point_sub                      | point_vert                  | poly_above                           |
| poly_below                       | poly_center                               | poly_contain                      | poly_contains_pt               | poly_contained                 | poly_distance               | poly_in                              |
| poly_left                        | poly_npoints                              | poly_out                          | poly_overabove                 | poly_overbelow                 | poly_overlap                | poly_overleft                        |
| poly_overright                   | poly_recv                                 | poly_right                        | poly_same                      | poly_send                      | polygon                     | position                             |

|                        |                      |                                  |                                   |                               |                                   |                                    |
|------------------------|----------------------|----------------------------------|-----------------------------------|-------------------------------|-----------------------------------|------------------------------------|
| positionjoin<br>sel    | positions<br>el      | postgresq<br>l_fdw_vali<br>dator | pow                               | power                         | prsd_e<br>nd                      | prsd_headli<br>ne                  |
| prsd_lextype           | prsd_nex<br>ttoken   | prsd_star<br>t                   | pt_cont<br>ained_c<br>ircle       | pt_cont<br>ained_p<br>oly     | query_<br>to_xml                  | query_to_x<br>ml_and_xml<br>schema |
| query_to_x<br>mlschema | quote_id<br>ent      | quote_lit<br>eral                | quote_<br>nullabl<br>e            | radians                       | radius                            | random                             |
| range_adjac<br>ent     | range_af<br>ter      | range_be<br>fore                 | range_<br>cmp                     | range_c<br>ontaine<br>d_by    | range_<br>contai<br>ns            | range_conta<br>ins_elem            |
| range_eq               | range_ge             | range_gis<br>t_compre<br>ss      | range_<br>gist_co<br>nsisten<br>t | range_g<br>ist_deco<br>mpress | range_<br>gist_pe<br>nalty        | range_gist_p<br>icksplit           |
| range_gist_s<br>ame    | range_gi<br>st_union | range_gt                         | range_i<br>n                      | range_i<br>ntersect           | range_<br>le                      | range_lt                           |
| range_minu<br>s        | range_ne             | range_ou<br>t                    | range_<br>overlap<br>s            | range_o<br>verleft            | range_<br>overrig<br>ht           | range_rcv                          |
| range_send             | range_ty<br>panalyze | range_uni<br>on                  | rank                              | record_e<br>q                 | record_<br>ge                     | record_gt                          |
| record_in              | record_le            | record_lt                        | record_<br>ne                     | record_<br>out                | record_<br>rcv                    | record_send                        |
| regclass               | regclassi<br>n       | regclasso<br>ut                  | regclas<br>srecv                  | regclass<br>send              | regconf<br>igin                   | regconfigou<br>t                   |
| regconfigrec<br>v      | regconfig<br>send    | regdictio<br>naryin              | regdicti<br>onaryou<br>t          | regdicti<br>onaryre<br>cv     | regdict<br>ionary<br>send         | regexeqjoins<br>el                 |
| regexeqsel             | regexnej<br>oinsel   | regexnes<br>el                   | regexp_<br>match<br>es            | regexp_<br>replace            | regexp_<br>split_t<br>o_arra<br>y | regexp_split<br>_to_table          |
| regoperatori<br>n      | regopera<br>torout   | regoperat<br>orrecv              | regope<br>ratorse<br>nd           | regoperi<br>n                 | regope<br>rout                    | regoperrecv                        |
| regopersend            | regproce<br>durein   | regproce<br>dureout              | regproc<br>edurer<br>ecv          | regproc<br>edurese<br>nd      | regpro<br>cin                     | regprocout                         |
| regprocrecv            | regprocs<br>end      | regr_avgx                        | regr_av<br>gy                     | regr_co<br>unt                | regr_in<br>tercept                | regr_r2                            |

|                                   |                                  |                                   |                                                |                                     |                                       |                          |
|-----------------------------------|----------------------------------|-----------------------------------|------------------------------------------------|-------------------------------------|---------------------------------------|--------------------------|
| regr_slope                        | regr_sxx                         | regr_sxy                          | regr_sy<br>y                                   | regtypei<br>n                       | regtyp<br>eout                        | regtyperecv              |
| regtypesend                       | reltime                          | reltimeeq                         | reltime<br>ge                                  | reltimeg<br>t                       | reltim<br>ein                         | reltimele                |
| reltimelt                         | reltimen<br>e                    | reltimeou<br>t                    | reltime<br>recv                                | reltimes<br>end                     | repeat                                | replace                  |
| reverse                           | RI_FKey_<br>cascade_<br>del      | RI_FKey_<br>ascade_u<br>pd        | RI_FKe<br>y_chec<br>k_ins                      | RI_FKey_<br>_check_<br>upd          | RI_FKe<br>y_noac<br>tion_d<br>el      | RI_FKey_noa<br>ction_upd |
| RI_FKey_restr<br>ict_del          | RI_FKey_<br>restrict_u<br>pd     | RI_FKey_s<br>etdefault_<br>del    | RI_FKe<br>y_setde<br>fault_u<br>pd             | RI_FKey_<br>_setnull_<br>del        | RI_FKe<br>y_setn<br>ull_up<br>d       | right                    |
| round                             | row_num<br>ber                   | row_to_js<br>on                   | rpadd                                          | rtrim                               | scalarg<br>tjoinse<br>l               | scalargtsel              |
| scalarltjoins<br>el               | scalarltse<br>l                  | schema_t<br>o_xml                 | schem<br>a_to_x<br>ml_and_<br>_xmlsc<br>hema   | schema<br>_to_xml<br>schema         | sessio<br>n_user                      | set_bit                  |
| set_byte                          | set_confi<br>g                   | set_maskl<br>en                   | shift_jis<br>_2004_<br>to_euc<br>_jis_20<br>04 | shift_jis<br>_2004_<br>to_<br>_utf8 | sjis_to<br>_euc_j<br>p                | sjis_to_mic              |
| sjis_to_utf8                      | smgrin                           | smgrout                           | spg_kd_<br>_choos<br>e                         | spg_kd_<br>config                   | spg_kd_<br>_inner_<br>_consis<br>tent | spg_kd_pick<br>split     |
| spg_quad_c<br>hoose               | spg_qua<br>d_config              | spg_quad<br>_inner_co<br>nsistent | spg_qu<br>ad_leaf_<br>_consis<br>tent          | spg_qua<br>d_picksp<br>lit          | spg_te<br>xt_cho<br>ose               | spg_text_co<br>nfig      |
| spg_text_inn<br>er_consisten<br>t | spg_text_<br>leaf_cons<br>istent | spg_text_<br>picksplit            | spgbeg<br>inscan                               | spgbuild                            | spgbui<br>ldemp<br>ty                 | spgbulkdele<br>te        |
| spgcanretur<br>n                  | spgcoste<br>stimate              | spgendsc<br>an                    | spgget<br>bitmap                               | spggett<br>uple                     | spgins<br>ert                         | spgmarkpos               |
| spgoptions                        | spgresca<br>n                    | spgrestrp<br>os                   | spgvac<br>uumcle<br>anup                       | stddev                              | stddev<br>_pop                        | stddev_sam<br>p          |

|                            |                          |                          |                  |                   |                                    |                           |
|----------------------------|--------------------------|--------------------------|------------------|-------------------|------------------------------------|---------------------------|
| string_agg                 | string_agg_finalfn       | string_agg_transfn       | strip            | sum               | suppress_redundant_updates_trigger | table_to_xml              |
| table_to_xml_and_xmlschema | table_to_xmlschema       | tan                      | text             | text_ge           | text_gt                            | text_larger               |
| text_le                    | text_lt                  | text_pattern_ge          | text_pattern_gt  | text_pattern_le   | text_pattern_lt                    | text_smaller              |
| textanycat                 | textcat                  | texteq                   | texticlike       | texticnlike       | texticregexe                       | texticregexne             |
| textin                     | textlike                 | textne                   | textnlike        | textout           | textrecv                           | textregexe                |
| textregexne                | textsend                 | thesaurus_init           | thesaurus_lexize | tideq             | tidge                              | tidgt                     |
| tidin                      | tidlarger                | tidle                    | tidlt            | tidne             | tidout                             | tidrecv                   |
| tidsend                    | tidsmaller               | time                     | time_cmp         | time_eq           | time_ge                            | time_gt                   |
| time_hash                  | time_in                  | time_larger              | time_le          | time_lt           | time_mi_interval                   | time_mi_time              |
| time_ne                    | time_out                 | time_pl_interval         | time_recv        | time_send         | time_smaller                       | time_transform            |
| timedate_pl                | timemi                   | timepl                   | timestamp        | timestamp_cmp     | timestamp_cmp_date                 | timestamp_cmp_timestamptz |
| timestamp_eq               | timestamp_eq_date        | timestamp_eq_timestamptz | timestamp_ge     | timestamp_ge_date | timestamp_ge_timestamptz           | timestamp_gt              |
| timestamp_gt_date          | timestamp_gt_timestamptz | timestamp_hash           | timestamp_in     | timestamp_larger  | timestamp_le                       | timestamp_le_date         |



|                                   |                                      |                                  |                                      |                            |                               |                                 |
|-----------------------------------|--------------------------------------|----------------------------------|--------------------------------------|----------------------------|-------------------------------|---------------------------------|
| timestamp_l<br>e_timestamp<br>ptz | timesta<br>mp_lt                     | timestamp<br>p_lt_date           | timesta<br>mp_lt_t<br>imesta<br>mptz | timesta<br>mp_mi           | timestamp_<br>mi_int<br>erval | timestamp_<br>ne                |
| timestamp_<br>ne_date             | timesta<br>mp_ne_t<br>imestamp<br>tz | timestamp<br>p_out               | timesta<br>mp_pl_<br>interval        | timesta<br>mp_recv         | timestamp_<br>send            | timestamp_<br>smaller           |
| timestamp_<br>sortsupport         | timesta<br>mp_trans<br>form          | timestamp<br>typmodi<br>n        | timesta<br>mp_typ<br>modou<br>t      | timesta<br>mptz            | timestamp_<br>cmp             | timestamp<br>z_cmp_date         |
| timestamp<br>z_cmp_time<br>stamp  | timesta<br>mptz_eq                   | timestamp<br>tz_eq_d<br>ate      | timesta<br>mptz_e<br>q_time<br>stamp | timesta<br>mptz_g<br>e     | timestamp_<br>ge_da<br>te     | timestamp<br>z_ge_time<br>stamp |
| timestamp<br>z_gt                 | timesta<br>mptz_gt_<br>date          | timestamp<br>tz_gt_t<br>imestamp | timesta<br>mptz_i<br>n               | timesta<br>mptz_l<br>arger | timestamp_<br>le              | timestamp<br>z_le_date          |
| timestamp<br>z_le_time<br>stamp   | timesta<br>mptz_lt                   | timestamp<br>tz_lt_d<br>ate      | timesta<br>mptz_lt_<br>time<br>stamp | timesta<br>mptz_m<br>i     | timestamp_<br>mi_in<br>terval | timestamp<br>z_ne               |
| timestamp<br>z_ne_date            | timesta<br>mptz_ne_<br>time<br>stamp | timestamp<br>tz_out              | timesta<br>mptz_p<br>l_int<br>erval  | timesta<br>mptz_r<br>ecv   | timestamp_<br>send            | timestamp<br>z_smaller          |
| timestamp<br>ztypmodin            | timesta<br>mptztyp<br>modout         | timetypm<br>odin                 | timetyp<br>pmodou<br>t               | timetz                     | timetz_<br>cmp                | timetz_eq                       |
| timetz_ge                         | timetz_gt                            | timetz_h<br>ash                  | timetz_<br>in                        | timetz_l<br>arger          | timetz_<br>le                 | timetz_lt                       |
| timetz_mi_<br>interval            | timetz_n<br>e                        | timetz_o<br>ut                   | timetz_<br>pl_int<br>erval           | timetz_r<br>ecv            | timetz_<br>send               | timetz_sma<br>ller              |
| timetzdate_<br>pl                 | timetzty<br>pmodin                   | timetztyp<br>modout              | timezo<br>ne<br>(2069)               | timezon<br>e<br>(1159)     | timezo<br>ne<br>(2037)        | timezone<br>(2070)              |
| timezone<br>(1026)                | timezone<br>(2038)                   | tintervalc<br>t                  | tinterv<br>aleq                      | tinterval<br>ge            | tinterv<br>algt               | tintervalin                     |
| tintervalle                       | tintervall<br>eneq                   | tintervall<br>enge               | tinterv<br>allengt                   | tinterval<br>lenle         | tinterv<br>allenlt            | tintervallen<br>ne              |

|                        |                              |                       |                            |                                 |                                                |                       |
|------------------------|------------------------------|-----------------------|----------------------------|---------------------------------|------------------------------------------------|-----------------------|
| tintervallt            | tinterval<br>ne              | tintervalo<br>ut      | tinterv<br>alov            | tinterval<br>recv               | tinterv<br>alsam<br>e                          | tintervalsen<br>d     |
| tintervalstar<br>t     | to_ascii<br>(1845)           | to_ascii<br>(1847)    | to_ascii<br>(1846)         | trigger_i<br>n                  | trigger<br>_out                                | ts_match_qv           |
| ts_match_tq            | ts_match<br>_tt              | ts_match<br>_vq       | ts_rank                    | ts_rank_<br>cd                  | ts_rew<br>rite                                 | ts_stat               |
| ts_token_ty<br>pe      | ts_typan<br>alyze            | tsmatchj<br>oinsel    | tsmatc<br>hsel             | tsq_mco<br>ntained              | tsq_mc<br>ontain<br>s                          | tsquery_and           |
| tsquery_cm<br>p        | tsquery_<br>eq               | tsquery_g<br>e        | tsquery<br>_gt             | tsquery_<br>le                  | tsquer<br>y_lt                                 | tsquery_ne            |
| tsquery_not            | tsquery_<br>or               | tsqueryin             | tsquery<br>out             | tsqueryr<br>ecv                 | tsquer<br>ysend                                | tsrange               |
| tsrange_sub<br>diff    | tstzrange                    | tstzrange<br>_subdiff | tsvecto<br>r_cmp           | tsvector<br>_concat             | tsvecto<br>r_eq                                | tsvector_ge           |
| tsvector_gt            | tsvector_<br>le              | tsvector_l<br>t       | tsvecto<br>r_ne            | tsvector<br>_update<br>_trigger | tsvecto<br>r_upda<br>te_trig<br>ger_co<br>lumn | tsvectorin            |
| tsvectorout            | tsvectorr<br>ecv             | tsvectors<br>end      | txid_cu<br>rrent           | txid_cur<br>rent_sn<br>apshot   | txid_sn<br>apshot<br>_in                       | txid_snapsh<br>ot_out |
| txid_snapsh<br>ot_recv | txid_snap<br>shot_sen<br>d   | txid_snap<br>shot_xip | txid_sn<br>apshot<br>_xmax | txid_sna<br>pshot_x<br>min      | txid_vi<br>sible_i<br>n_snap<br>shot           | uhc_to_utf8           |
| unique_key_<br>recheck | unknown<br>in                | unknown<br>out        | unkno<br>wnrecv            | unknow<br>nsend                 | unnest                                         | utf8_to_big5          |
| utf8_to_euc<br>_cn     | utf8_to_e<br>uc_jis_20<br>04 | utf8_to_e<br>uc_jp    | utf8_to<br>_euc_kr         | utf8_to_<br>euc_tw              | utf8_t<br>o_gb1<br>8030                        | utf8_to_gbk           |
| utf8_to_iso8<br>859    | utf8_to_i<br>so8859_1        | utf8_to_j<br>ohab     | utf8_to<br>_koi8r          | utf8_to_<br>koi8u               | utf8_t<br>o_shift<br>_jis_20<br>04             | utf8_to_sjis          |
| utf8_to_uhc            | utf8_to_<br>win              | uuid_cmp              | uuid_e<br>q                | uuid_ge                         | uuid_g<br>t                                    | uuid_hash             |
| uuid_in                | uuid_le                      | uuid_lt               | uuid_n<br>e                | uuid_ou<br>t                    | uuid_r<br>ecv                                  | uuid_send             |

|                    |                            |                             |                   |                 |                   |                  |
|--------------------|----------------------------|-----------------------------|-------------------|-----------------|-------------------|------------------|
| var_pop            | var_samp                   | varbit                      | varbit_in         | varbit_out      | varbit_recv       | varbit_send      |
| varbit_transform   | varbitcmp                  | varbiteq                    | varbitge          | varbitgt        | varbitle          | varbitlt         |
| varbitne           | varbittypmodin             | varbittypmodout             | vchar             | vchar_transform | vcharin           | vcharout         |
| vcharrecv          | vcharsend                  | vchartypmodin               | vchartypmodout    | variance        | void_in           | void_out         |
| void_recv          | void_send                  | win_to_utf8                 | win1250_to_latin2 | win1250_to_mic  | win1251_to_iso    | win1251_to_koi8r |
| win1251_to_mic     | win1251_to_win866          | win866_to_iso               | win866_to_koi8r   | win866_to_mic   | win866_to_win1251 | xideq            |
| xideqint4          | xidin                      | xidout                      | xidrecv           | xidsend         | xml               | xml_in           |
| xml_is_well_formed | xml_is_well_formed_content | xml_is_well_formed_document | xml_out           | xml_recv        | xml_send          | xmlagg           |
| xmlcomment         | xmlconcat2                 | xmlexists                   | xmlvalidate       | pg_notify       | -                 | -                |

## Internal Functions

The following lists the functions used by GaussDB to implement internal system functions. You are advised not to use these functions. If you need to use them, contact Huawei technical support.

- locktag\_decode(locktag text)  
Description: Parses lock details from **locktag**.  
Return type: text
- smgreq(a smgr, b smgr)  
Description: Compares two smgrs to check whether they are the same.  
Parameter: smgr, smgr  
Return type: Boolean
- smgrne(a smgr, b smgr)  
Description: Checks whether the two smgrs are different.  
Parameter: smgr, smgr  
Return type: Boolean

- `xidin4`  
Description: Inputs a 4-byte xid.  
Parameter: `cstring`  
Return type: `xid32`
- `set_hashbucket_info`  
Description: Sets hash bucket information.  
Parameter: `text`  
Return type: `Boolean`
- `int1send`  
Description: Packs unsigned 1-byte integers into the internal data buffer stream.  
Parameter: `tinyint`  
Return type: `bytea`
- `listagg`  
Description: Specifies aggregate functions of the list type.  
Parameter: `smallint`, `text`  
Return type: `text`
- `log_fdw_validator`  
Description: Specifies validation functions.  
Parameter: `text[]`, `oid`  
Return type: `void`
- `nvarchar2typmodin`  
Description: Obtains the typmod information of the varchar type.  
Parameter: `cstring[]`  
Return type: `integer`
- `nvarchar2typmodout`  
Description: Obtains the typmod information of the varchar type, constructs a character string, and returns the character string.  
Parameter: `integer`  
Return type: `cstring`
- `read_disable_conn_file`  
Description: Reads forbidden connection files.  
Parameter: `nan`  
Return type: `disconn_mode text`, `disconn_host text`, `disconn_port text`, `local_host text`, `local_port text`, `redo_finished text`
- `update_pgjob`  
Description: Updates a job.  
Parameter: `bigint`, `"char"`, `bigint`, `timestamp without time zone`, `timestamp without time zone`, `timestamp without time zone`, `timestamp without time zone`, `timestamp without time zone`, `smallint`, `text`

Return type: void

- `enum_cmp`

Description: Specifies the enumeration comparison function, which is used to determine whether two enumeration classes are equal and determine their relative sizes.

Parameter: anyenum, anyenum

Return type: integer

- `enum_eq`

Description: Specifies the enumeration comparison function, which is used to implement the equal sign (=).

Parameter: anyenum, anyenum

Return type: Boolean

- `enum_first`

Description: Returns the first element in the enumeration class.

Parameter: anyenum

Return type: anyenum

- `enum_ge`

Description: Specifies the enumeration comparison function, which is used to implement the greater-than sign (>) and equal sign (=).

Parameter: anyenum, anyenum

Return type: Boolean

- `enum_gt`

Description: Specifies the enumeration comparison function, which is used to implement the greater-than sign (>).

Parameter: anyenum, anyenum

Return type: Boolean

- `enum_in`

Description: Specifies the enumeration comparison function, which is used to determine whether an element is in an enumeration class.

Parameter:cstring, oid

Return type: anyenum

- `enum_larger`

Description: Specifies the enumeration comparison function, which is used to implement the greater-than sign (>).

Parameter: anyenum, anyenum

Return type: anyenum

- `enum_last`

Description: Returns the last element in the enumeration class.

Parameter: anyenum

Return type: anyenum

- `enum_le`  
Description: Specifies the enumeration comparison function, which is used to implement the less-than sign (<) and equal sign (=).  
Parameter: anyenum, anyenum  
Return type: Boolean
- `enum_lt`  
Description: Specifies the enumeration comparison function, which is used to implement the less-than sign (<).  
Parameter: anyenum, anyenum  
Return type: Boolean
- `enum_smaller`  
Description: Specifies the enumeration comparison function, which is used to implement the less-than sign (<).  
Parameter: anyenum, anyenum  
Return type: Boolean
- `node_oid_name`  
Description: Not supported.  
Parameter: oid  
Return type: cstring
- `pg_buffercache_pages`  
Description: Reads data from the shared buffer.  
Parameter: nan  
Return type: bufferid integer, relfilenode oid, bucketid smallint, reltablespace oid, reldatabase oid, relforknumber smallint, relblocknumber bigint, isdirty Boolean, usage\_count smallint
- `pg_check_xidlimit`  
Description: Checks whether nextxid is greater than or equal to xidwarnlimit.  
Parameter: nan  
Return type: Boolean
- `pg_comm_delay`  
Description: Displays the delay status of the communications library of a single DN.  
Parameter: nan  
Return type: text, text, integer, integer, integer, integer
- `pg_comm_rcv_stream`  
Description: Displays the receiving stream status of all communication libraries on a single DN.  
Parameter: nan  
Return type: text, bigint, text, bigint, integer, integer, integer, text, bigint, integer, integer, integer, bigint, bigint, bigint, bigint, bigint

- `pg_comm_send_stream`  
Description: Displays the sending stream status of all communication libraries on a single DN.  
Parameter: nan  
Return type: text, bigint, text, bigint, integer, integer, integer, text, bigint, integer, integer, integer, bigint, bigint, bigint, bigint
- `pg_comm_status`  
Description: Displays the communication status of a single DN.  
Parameter: nan  
Return type: text, integer, integer, bigint, bigint, bigint, bigint, bigint, integer, integer, integer, integer
- `pg_log_comm_status`  
Description: Prints some logs on the DN.  
Parameter: nan  
Return type: Boolean
- `pg_parse_clog`  
Description: Parses clog to obtain the status of xid.  
Parameter: nan  
Return type: xid xid, status text
- `pg_pool_ping`  
Description: Sets PoolerPing.  
Parameter: Boolean  
Return type: SETOF Boolean
- `pg_resume_bkp_flag`  
Description: Obtains the delay Xlog flag for backup and restoration.  
Parameter: slot\_name name  
Return type: start\_backup\_flag Boolean, to\_delay Boolean, ddl\_delay\_recycle\_ptr text, rewind\_time text
- `psortoptions`  
Description: Returns the psort attribute.  
Parameter: text[], Boolean  
Return type: bytea
- `xideq4`  
Description: Compares two values of the xid type to check whether they are the same.  
Parameter: xid32, xid32  
Return type: Boolean
- `xideqint8`  
Description: Compares values of the xid type and int8 type to check whether they are the same.  
Parameter: xid, bigint

- Return type: Boolean
- xidlt**  
 Description: Returns whether xid1 < xid2 is true.  
 Parameter: xid, xid  
 Return type: Boolean
  - xidlt4**  
 Description: Returns whether xid1 < xid2 is true.  
 Parameter: xid32, xid32  
 Return type: Boolean

## 7.5.37 Obsolete Functions

The following functions in GaussDB have been deprecated in the latest version:

|                               |                                    |                                       |                                        |                                          |                                         |                                    |
|-------------------------------|------------------------------------|---------------------------------------|----------------------------------------|------------------------------------------|-----------------------------------------|------------------------------------|
| gs_wlm_get_session_info       | gs_wlm_get_user_session_info       | pgxc_get_csn                          | pgxc_get_stat_dirty_tables             | pgxc_get_thread_wait_statuses            | pgxc_get_m_snapshots_hot_statuses       | pgxc_is_committed                  |
| pgxc_lock_for_backup          | pgxc_lock_for_sp_database          | pgxc_lock_for_transfer                | pgxc_log_comm_status                   | pgxc_max_datanode_size                   | pgxc_node_str                           | pgxc_pool_check                    |
| pgxc_pool_connection_statuses | pgxc_pool_reload                   | pgxc_prepared_xact                    | pgxc_snapshot_status                   | pgxc_stat_dirty_tables                   | pgxc_unlock_for_snapshot_database       | pgxc_unlock_for_transfer           |
| pgxc_version                  | array_extend                       | prepare_statement_status              | remote_rt_status                       | dbe_perf.global_slow_query_info          | dbe_perf.global_slow_query_info_by_time | dbe_perf.global_slow_query_history |
| pg_stat_get_pooler_status     | pg_stat_get_wlm_node_resource_info | pg_stat_get_wlm_session_info_internal | DBE_PERF.get_wlm_controlgroup_config() | DBE_PERF.get_wlm_user_resource_runtime() | global_space_shrink                     | pg_pool_validate                   |



|                    |                              |                                             |                                               |                                        |   |   |
|--------------------|------------------------------|---------------------------------------------|-----------------------------------------------|----------------------------------------|---|---|
| gs_stat_u<br>store | table_ske<br>wness(te<br>xt) | table_ske<br>wness(te<br>xt, text,<br>text) | gs_w<br>lm_g<br>et_se<br>ssion<br>_inf<br>o() | gs_wlm_get_<br>user_session_i<br>nfo() | - | - |
|--------------------|------------------------------|---------------------------------------------|-----------------------------------------------|----------------------------------------|---|---|

## 7.6 Expressions

### 7.6.1 Simple Expressions

#### Logical Expressions

[Logical Operators](#) lists the operators and calculation rules of logical expressions.

#### Comparative Expressions

[Comparison Operators](#) lists the common comparative operators.

In addition to comparative operators, you can also use the following sentence structure:

- BETWEEN operator  
**a BETWEEN x AND y** is equivalent to **a >= x AND a <= y**.  
**a NOT BETWEEN x AND y** is equivalent to **a < x OR a > y**.
- To check whether a value is null, use:  
expression IS NULL  
expression IS NOT NULL  
or an equivalent (non-standard) sentence structure:  
expression ISNULL  
expression NOTNULL

---

#### NOTICE

Do not write **expression=NULL** or **expression<>(!=)NULL**, because **NULL** represents an unknown value, and these expressions cannot determine whether two unknown values are equal.

---

- is distinct from/is not distinct from
  - is distinct from  
If the data types and values of A and B are different, the value is **true**.  
If the data types and values of A and B are the same, the value is **false**.  
Empty values are considered the same.
  - is not distinct from

If the data types and values of A and B are different, the value is **false**.  
If the data types and values of A and B are the same, the value is **true**.  
Empty values are considered the same.

## Pseudocolumn

### ROWNUM

**ROWNUM** is a pseudocolumn that returns a number indicating the row number of the result obtained from the query. The value of **ROWNUM** in the first row is **1**, the value of **ROWNUM** in the second row is **2**, and so on. The return type of **ROWNUM** is BIGINT. **ROWNUM** can be used to limit the total number of rows returned by a query. For example, the following statement limits the maximum number of records returned from the table **Students** to 10.

```
openGauss=# CREATE TABLE Students (name varchar(20), id int) with (STORAGE_TYPE = USTORE);
openGauss=# INSERT INTO Students VALUES ('Jack', 35);
openGauss=# INSERT INTO Students VALUES ('Leon', 15);
openGauss=# INSERT INTO Students VALUES ('James', 24);
openGauss=# INSERT INTO Students VALUES ('Taker', 81);
openGauss=# INSERT INTO Students VALUES ('Mary', 25);
openGauss=# INSERT INTO Students VALUES ('Rose', 64);
openGauss=# INSERT INTO Students VALUES ('Perl', 18);
openGauss=# INSERT INTO Students VALUES ('Under', 57);
openGauss=# INSERT INTO Students VALUES ('Angel', 101);
openGauss=# INSERT INTO Students VALUES ('Frank', 20);
openGauss=# INSERT INTO Students VALUES ('Charlie', 40);

-- Output the first 10 rows of data records in the Students table.
openGauss=# SELECT * FROM Students WHERE rownum <= 10;
 name | id
-----+-----
 Jack | 35
 Leon | 15
 James | 24
 Taker | 81
 Mary | 25
 Rose | 64
 Perl | 18
 Under | 57
 Angel | 101
 Frank | 20
(10 rows)

openGauss=# DROP TABLE Students;
DROP TABLE
```

## Examples

```
openGauss=# SELECT 2 BETWEEN 1 AND 3 AS RESULT;
 result

 t
(1 row)

openGauss=# SELECT 2 >= 1 AND 2 <= 3 AS RESULT;
 result

 t
(1 row)

openGauss=# SELECT 2 NOT BETWEEN 1 AND 3 AS RESULT;
 result

 f
```

```
(1 row)

openGauss=# SELECT 2 < 1 OR 2 > 3 AS RESULT;
result

f
(1 row)

openGauss=# SELECT 2+2 IS NULL AS RESULT;
result

f
(1 row)

openGauss=# SELECT 2+2 IS NOT NULL AS RESULT;
result

t
(1 row)

openGauss=# SELECT 2+2 ISNULL AS RESULT;
result

f
(1 row)

openGauss=# SELECT 2+2 NOTNULL AS RESULT;
result

t
(1 row)

openGauss=# SELECT 2+2 IS DISTINCT FROM NULL AS RESULT;
result

t
(1 row)

openGauss=# SELECT 2+2 IS NOT DISTINCT FROM NULL AS RESULT;
result

f
(1 row)
```

## 7.6.2 Condition Expressions

Data that meets the requirements specified by conditional expressions are filtered during SQL statement execution.


Conditional expressions include the following types:

- CASE

**CASE** expressions are similar to the **CASE** statements in other coding languages.

**Figure 7-1** shows the syntax of a **CASE** expression.

**Figure 7-1** case::=



A **CASE** clause can be used in a valid expression. **condition** is an expression that returns a value of Boolean type.

- If the result is true, the result of the **CASE** expression is the required result.
- If the result is false, the following **WHEN** or **ELSE** clauses are processed in the same way.
- If every **WHEN condition** is false, the result of the expression is the result of the **ELSE** clause. If the **ELSE** clause is omitted and has no match condition, the result is NULL.

Example:

```
openGauss=# CREATE TABLE case_when_t1(CW_COL1 INT);

openGauss=# INSERT INTO case_when_t1 VALUES (1), (2), (3);

openGauss=# SELECT * FROM case_when_t1;
cw_col1

1
2
3
(3 rows)


openGauss=# SELECT CW_COL1, CASE WHEN CW_COL1=1 THEN 'one' WHEN CW_COL1=2 THEN
'two' ELSE 'other' END FROM case_when_t1 ORDER BY 1;
cw_col1 | case
-----+-----
1 | one
2 | two
3 | other
(3 rows)

openGauss=# DROP TABLE case_when_t1;
```

- **DECODE**

**Figure 7-2** shows the syntax of a **DECODE** expression.

**Figure 7-2** decode::=



Compare each following **compare(n)** with **base\_expr**. **value(n)** is returned if a **compare(n)** matches the **base\_expr** expression. If no match occurs, default is returned.


**Conditional Expression Functions** describes the examples.

```
openGauss=# SELECT DECODE('A','A',1,'B',2,0);
case

1
(1 row)
```

- **COALESCE**

**Figure 7-3** shows the syntax of a **COALESCE** expression.

**Figure 7-3** coalesce::=

**COALESCE** returns its first not-**NULL** value. If all the parameters are **NULL**, **NULL** is returned. This value is replaced by the default value when data is displayed. Like a **CASE** expression, **COALESCE** only evaluates the parameters that are needed to determine the result. That is, parameters to the right of the first non-null parameter are not evaluated.

Example:

```
openGauss=# CREATE TABLE c_tabl(description varchar(10), short_description varchar(10), last_value
varchar(10)) ;
```

```
openGauss=# INSERT INTO c_tabl VALUES('abc', 'efg', '123');
openGauss=# INSERT INTO c_tabl VALUES(NULL, 'efg', '123');
```

```
openGauss=# INSERT INTO c_tabl VALUES(NULL, NULL, '123');
```

```
openGauss=# SELECT description, short_description, last_value, COALESCE(description,
short_description, last_value) FROM c_tabl ORDER BY 1, 2, 3, 4;
```

| description | short_description | last_value | coalesce |
|-------------|-------------------|------------|----------|
| abc         | efg               | 123        | abc      |
|             | efg               | 123        | efg      |
|             |                   | 123        | 123      |

(3 rows)

```
openGauss=# DROP TABLE c_tabl;
```


If **description** is not **NULL**, the value of **description** is returned. Otherwise, parameter **short\_description** is calculated. If **short\_description** is not **NULL**, the value of **short\_description** is returned. Otherwise, parameter **last\_value** is calculated. If **last\_value** is not **NULL**, the value of **last\_value** is returned. Otherwise, **none** is returned.

```
openGauss=# SELECT COALESCE(NULL,'Hello World');
coalesce

Hello World
(1 row)
```

- **NULLIF**

**Figure 7-4** shows the syntax of a **NULLIF** expression.

**Figure 7-4** nullif::=

Only if **value1** is equal to **value2** can **NULLIF** return the **NULL** value. Otherwise, **value1** is returned.

Example:

```
openGauss=# CREATE TABLE null_if_t1 (
NI_VALUE1 VARCHAR(10),
NI_VALUE2 VARCHAR(10))
```

```
);

openGauss=# INSERT INTO null_if_t1 VALUES('abc', 'abc');
openGauss=# INSERT INTO null_if_t1 VALUES('abc', 'efg');

openGauss=# SELECT NI_VALUE1, NI_VALUE2, NULLIF(NI_VALUE1, NI_VALUE2) FROM null_if_t1
ORDER BY 1, 2, 3;

ni_value1 | ni_value2 | nullif
-----+-----+-----
abc | abc |
abc | efg | abc
(2 rows)
openGauss=# DROP TABLE null_if_t1;
```

If the value of **value1** is equal to that of **value2**, **NULL** is returned. Otherwise, the value of **value1** is returned.


```
openGauss=# SELECT NULLIF('Hello','Hello World');
nullif

Hello
(1 row)
```

- GREATEST (maximum value) and LEAST (minimum value)

**Figure 7-5** shows the syntax of a **GREATEST** expression.

**Figure 7-5** greatest::=




You can select the maximum value from any numerical expression list.

```
openGauss=# SELECT greatest(9000,15555,2.01);
greatest

15555
(1 row)
```

**Figure 7-6** shows the syntax of a **LEAST** expression.

**Figure 7-6** least::=



You can select the minimum value from any numerical expression list.

Each of the preceding numeric expressions can be converted into a common data type, which will be the data type of the result.

The NULL values in the list will be ignored. The result is **NULL** only if the results of all expressions are **NULL**.

```
openGauss=# SELECT least(9000,2);
least


2
(1 row)
```

**Conditional Expression Functions** describes the examples.

- NVL

**Figure 7-7** shows the syntax of an **NVL** expression.

**Figure 7-7** nvl::=



If the value of **value1** is **NULL**, the value of **value2** is returned. Otherwise, the value of **value1** is returned.

Example:

```
openGauss=# SELECT nvl(null,1);
nvl

1
(1 row)

openGauss=# SELECT nvl ('Hello World',1);
nvl

Hello World
(1 row)
```


## 7.6.3 Subquery Expressions

Subquery expressions include the following types:

- EXISTS/NOT EXISTS

**Figure 7-8** shows the syntax of an **EXISTS/NOT EXISTS** expression.

**Figure 7-8** EXISTS/NOT EXISTS::=



The parameter of an **EXISTS** expression is an arbitrary **SELECT** statement, or subquery. The subquery is evaluated to determine whether it returns any rows. If it returns at least one row, the result of **EXISTS** is true. If the subquery returns no rows, the result of **EXISTS** is false.

The subquery will generally only be executed long enough to determine whether at least one row is returned, not all the way to completion.

Example:

```
openGauss=# CREATE TABLE exists_t1(a int, b int);
openGauss=# INSERT INTO exists_t1 VALUES(1, 2),(2, 3),(3, 4),(4, 5);

openGauss=# CREATE TABLE exists_t2(a int, c int);
openGauss=# INSERT INTO exists_t2 VALUES(3, 4),(4, 5),(5, 6),(6, 7);

openGauss=# SELECT * FROM exists_t1 t1 WHERE EXISTS (SELECT * FROM exists_t2 t2 WHERE t2.a =
t1.a);
```


```
a | b
---+---
3 | 4
4 | 5
(2 rows)

openGauss=# DROP TABLE exists_t1, exists_t2;
```

- IN/NOT IN

**Figure 7-9** shows the syntax of an **IN/NOT IN** expression.

**Figure 7-9** IN/NOT IN::=



The right-hand side is a parenthesized subquery, which must return exactly one column. The left-hand expression is evaluated and compared to each row of the subquery result. The result of **IN** is true if any equal subquery row is found. The result is false if no equal row is found (including the case where the subquery returns no rows).

This is in accordance with SQL normal rules for Boolean combinations of null values. If the columns corresponding to two rows equal and are not empty, the two rows are equal to each other. If any columns corresponding to the two rows do not equal and are not empty, the two rows are not equal to each other. Otherwise, the result is **NULL**. If there are no equal right-hand values and at least one right-hand row yields null, the result of **IN** will be null, not false.

Example:

```
openGauss=# CREATE TABLE in_t1(a int, b int);
openGauss=# INSERT INTO in_t1 VALUES(1, 2),(2, 3),(3, 4),(4, 5);

openGauss=# CREATE TABLE in_t2(a int, c int);
openGauss=# INSERT INTO in_t2 VALUES(3, 4),(4, 5),(5, 6),(6, 7);


openGauss=# SELECT * FROM in_t1 t1 WHERE t1.a IN (SELECT t2.a FROM in_t2 t2);
a | b
---+---
3 | 4
4 | 5
(2 rows)

openGauss=# DROP TABLE in_t1, in_t2;
```

- ANY/SOME

**Figure 7-10** shows the syntax of an **ANY/SOME** expression.

**Figure 7-10** any/some::=





The right-hand side is a parenthesized subquery, which must return exactly one column. The left-hand expression is evaluated and compared to each row of the subquery result using the given operator, which must yield a Boolean result. The result of **ANY** is "true" if any true result is obtained. The result is "false" if no true result is found (including the case where the subquery returns no rows). **SOME** is a synonym of **ANY**. **IN** can be equivalently replaced with **ANY**.

Example:

```
openGauss=# CREATE TABLE any_t1(a int, b int);
openGauss=# INSERT INTO any_t1 VALUES(1, 2),(2, 3),(3, 4),(4, 5);

openGauss=# CREATE TABLE any_t2(a int, c int);
openGauss=# INSERT INTO any_t2 VALUES(3, 4),(4, 5),(5, 6),(6, 7);


openGauss=# SELECT * FROM any_t1 t1 WHERE t1.a < ANY(SELECT t2.a FROM any_t2 t2 where t2.a =
3 or t2.a = 4);
 a | b
----+----
 1 | 2
 2 | 3
 3 | 4
(3 rows)

openGauss=# DROP TABLE any_t1, any_t2;
```

- **ALL**

**Figure 7-11** shows the syntax of an **ALL** expression.

**Figure 7-11** all::=



The right-hand side is a parenthesized subquery, which must return exactly one column. The left-hand expression is evaluated and compared to each row of the subquery result using the given operator, which must yield a Boolean result. The result of **ALL** is true if all values are true (including the case where the subquery returns no rows). The result is false if any false result is found.

Example:

```
openGauss=# CREATE TABLE all_t1(a int, b int);
openGauss=# INSERT INTO all_t1 VALUES(1, 2),(2, 3),(3, 4),(4, 5);

openGauss=# CREATE TABLE all_t2(a int, c int);
openGauss=# INSERT INTO all_t2 VALUES(3, 4),(4, 5),(5, 6),(6, 7);

openGauss=# SELECT * FROM all_t1 t1 WHERE t1.a < ALL(SELECT t2.a FROM all_t2 t2 where t2.a = 3
or t2.a = 4);
 a | b
----+----
 1 | 2
 2 | 3
(2 rows)

openGauss=# DROP TABLE all_t1, all_t2;
```

## 7.6.4 Array Expressions

### IN

*expression* **IN** (*value* [, ...])

The parentheses on the right contain an expression list. The expression result on the left is compared with the content in the expression list. If the content in the list meets the expression result on the left, the result of **IN** is **true**. If no result meets the requirements, the result of **IN** is **false**.

Example:

```
openGauss=# SELECT 8000+500 IN (10000, 9000) AS RESULT;
result

f
(1 row)
```

#### NOTE

If the expression result is null or the expression list does not meet the expression conditions and at least one empty value is returned for the expression list on the right, the result of **IN** is **null** rather than **false**. This method is consistent with the Boolean rules used when SQL statements return empty values.

### NOT IN

*expression* **NOT IN** (*value* [, ...])

The parentheses on the right contain an expression list. The expression result on the left is compared with the content in the expression list. If the content in the list does not meet the expression result on the left, the result of **NOT IN** is **true**. If any content meets the expression result, the result of **NOT IN** is **false**.

Example:

```
openGauss=# SELECT 8000+500 NOT IN (10000, 9000) AS RESULT;
result

t
(1 row)
```

#### NOTE

- If the query statement result is null or the expression list does not meet the expression conditions and at least one empty value is returned for the expression list on the right, the result of **NOT IN** is null rather than false. This method is consistent with the Boolean rules used when SQL statements return empty values.
- In all situations, X **NOT IN** Y equals to **NOT**(X **IN** Y).

### ANY/SOME (array)

*expression operator* **ANY** (*array expression*)

*expression operator* **SOME** (*array expression*)

The right side is a parenthesized expression, which must yield an array value. The result of the expression on the left uses operators to compute and compare the results in each row of the array expression. The comparison result must be a Boolean value.

```
openGauss=# SELECT 8000+500 < SOME (array[10000,9000]) AS RESULT;
result

t
(1 row)

openGauss=# SELECT 8000+500 < ANY (array[10000,9000]) AS RESULT;
result

t
(1 row)
```

#### NOTE

- If at least one comparison result is true, the result of ANY is true.
- If no comparison result is true, the result of ANY is false.
- If no comparison result is true and the array expression generates at least one null value, the value of **ANY** is **NULL** rather than **false**. This method is consistent with the Boolean rules used when SQL statements return empty values.
- SOME is a synonym of ANY.

## ALL (array)

*expression operator ALL (array expression)*

The right side is a parenthesized expression, which must yield an array value. The result of the expression on the left uses operators to compute and compare the results in each row of the array expression. The comparison result must be a Boolean value.

- The result of **ALL** is **true** if all comparisons yield **true** (including the case where the array has zero elements).
- The result of **ALL** is **false** if one or multiple comparisons yield **false**.
- If the array expression yields a null array, the result of **ALL** will be null. If the left-hand expression yields null, the result of **ALL** is ordinarily null (though a non-strict comparison operator could possibly yield a different result). Also, if the array on the right contains any null elements and no false comparison result is found, the result of **ALL** is **NULL**, not true (again, assuming a strict comparison operator). This method is consistent with the Boolean rules used when SQL statements return empty values.

```
openGauss=# SELECT 8000+500 < ALL (array[10000,9000]) AS RESULT;
result

t
(1 row)
```

## 7.6.5 Row Expressions

The syntax is as follows:

*row\_constructor operator row\_constructor*

Both sides of the row expression are row constructors. The values of both rows must have the same number of fields and they are compared with each other. The row comparison allows operators including =, <>, <, <=, and >= or a similar operator.

For operators <, <=, >, and >=, the columns in rows are compared from left to right until a pair of columns that are not equal or are empty are detected. If the pair of columns contains at least one null value, the comparison result is null. Otherwise, the comparison result of this pair of columns is the final result. If no unequal or empty column is found, the values in the two rows are equal. The final result is determined based on the operator meaning.

Operations on XML data are not supported.

Example:

```
openGauss=# SELECT ROW(1,2,NULL) < ROW(1,3,0) AS RESULT;
result

t
(1 row)

openGauss=# select (4,5,6) > (3,2,1) as result;
result

t
(1 row)

openGauss=# select (4,1,1) > (3,2,1) as result;
result

t
(1 row)

openGauss=# select ('test','data') > ('data','data') as result;
result

t
(1 row)

openGauss=# select (4,1,1) > (3,2,null) as result;
result

t
(1 row)

openGauss=# select (null,1,1) > (3,2,1) as result;
result

(1 row)

openGauss=# select (null,5,6) > (null,5,6) as result;
result

(1 row)

openGauss=# select (4,5,6) > (4,5,6) as result;
result

f
(1 row)

openGauss=# select (2,2,5) >= (2,2,3) as result;
result

t
(1 row)

openGauss=# select (2,2,1) <= (2,2,3) as result;
result

```

```
t
(1 row)
```

The use of operators = and <> is slightly different from other operators. If all columns in the two rows are not empty and meet the operator condition, the two rows meet the operator condition. If any column in the two rows is not empty and does not meet the operator condition, the two rows do not meet the operator condition. If any column in the two rows is empty, the comparison result is null.

Example:

```
openGauss=# select (1,2,3) = (1,2,3) as result;
result

t
(1 row)

openGauss=# select (1,2,3) <> (2,2,3) as result;
result

t
(1 row)

openGauss=# select (2,2,3) <> (2,2,null) as result;
result

(1 row)

openGauss=# select (null,5,6) <> (null,5,6) as result;
result

(1 row)
```

## 7.7 Type Conversion

### 7.7.1 Overview

#### Context

SQL is a typed language. That is, every data item has an associated data type which determines its behavior and allowed usage. GaussDB has an extensible type system that is more general and flexible than other SQL implementations. Hence, most type conversion behaviors in GaussDB are governed by general rules. This allows the use of mixed-type expressions.

The GaussDB scanner/parser divides lexical elements into five fundamental categories: integers, floating-point numbers, strings, identifiers, and keywords. Constants of most non-numeric types are first classified as strings. The SQL language definition allows specifying type names with constant strings. For example, the query:

```
openGauss=# SELECT text 'Origin' AS "label", point '(0,0)' AS "value";
label | value
-----+-----
Origin | (0,0)
(1 row)
```

has two literal constants, of type **text** and **point**. If a type is not specified for a string literal, then the placeholder type **unknown** is assigned initially.

There are four fundamental SQL constructs requiring distinct type conversion rules in GaussDB parser:

- **Function calls**  
Much of the SQL type system is built around a rich set of functions. Functions can have one or more arguments. Since SQL permits function overloading, the function name alone does not uniquely identify the function to be called. The parser must select the right function based on the data types of the supplied arguments.
- **Operators**  
SQL allows expressions with prefix and postfix unary (one-argument) operators, as well as binary (two-argument) operators. Like functions, operators can be overloaded, so the same problem of selecting the right operator exists.
- **Value storage**  
The INSERT and UPDATE statements place the results of expressions into a table. The expressions in the statement must be matched up with, and perhaps converted to, the types of the target columns.
- **UNION, CASE, and Related Constructs**  
Since all query results from a unionized **SELECT** statement must appear in a single set of columns, the types of the results of each **SELECT** clause must be matched up and converted to a uniform set. Similarly, the result expressions of a **CASE** construct must be converted to a common type so that the **CASE** expression as a whole has a known output type. The same holds for **ARRAY** constructs, and for the **GREATEST** and **LEAST** functions.

The system catalog `pg_cast` stores information about which conversions, or casts, exist between which data types, and how to perform those conversions. For details, see [PG\\_CAST](#).

The return type and conversion behavior of an expression are determined during semantic analysis. Data types are divided into several basic type categories, including Boolean, **numeric**, **string**, **bitstring**, **datetime**, **timespan**, **geometric**, and **network**. Within each category there can be one or more preferred types, which are preferred when there is a choice of possible types. With careful selection of preferred types and available implicit casts, it is possible to ensure that ambiguous expressions (those with multiple candidate parsing solutions) can be resolved in a useful way.

All type conversion rules are designed based on the following principles:

- Implicit conversions should never have surprising or unpredictable outcomes.
- There should be no extra overhead in the parser or executor if a query does not need implicit type conversion. That is, if a query is well-formed and the types already match, then the query should execute without spending extra time in the parser and without introducing unnecessary implicit conversion calls in the query.
- Additionally, if a query usually requires an implicit conversion for a function, and if then the user defines a new function with the correct argument types, the parser should use this new function.

## 7.7.2 Operators

### Operator Type Resolution

1. Select the operators to be considered from the **pg\_operator** system catalog. Considered operators are those with the matching name and argument count. If the search path finds multiple available operators, only the most suitable one is considered.
2. Look for the best match.
  - a. Discard candidate operators for which the input types do not match and cannot be converted (using an implicit conversion) to match. **unknown** literals are assumed to be convertible to anything for this purpose. If only one candidate remains, use it; else continue to the next step.
  - b. Run through all candidates and keep those with the most exact matches on input types. In this case, the domain types are considered the same as their basic types. Keep all candidates if none survives these tests. If only one candidate remains, use it; else continue to the next step.
  - c. Run through all candidates and keep those that accept preferred types (of the input data type's type category) at the most positions where type conversion will be required. Keep all candidates if none accepts preferred types. If only one candidate remains, use it; else continue to the next step.
  - d. If any input arguments are **unknown**, check the type categories accepted at those argument positions by the remaining candidates. At each position, select the string category if any candidate accepts that category. (This bias towards string is appropriate since an unknown-type literal looks like a string.) Otherwise, if all the remaining candidates accept the same type category, select that category; otherwise fail because the correct choice cannot be deduced without more clues. Now discard candidates that do not accept the selected type category. Furthermore, if any candidate accepts a preferred type in that category, discard candidates that accept non-preferred types for that argument. Keep all candidates if none survives these tests. If only one candidate remains, use it; else continue to the next step.
  - e. If there are both **unknown** and known-type arguments, and all the known-type arguments have the same type, assume that the **unknown** arguments are also of that type, and check which candidates can accept that type at the unknown-argument positions. If exactly one candidate passes this test, use it. Otherwise, an error occurs.

---

#### CAUTION

After an operator is found, if the type of the input parameter is different from that of the operator, implicit type conversion may occur. After the conversion, unpredictable behavior may occur. If the behavior is incorrect after implicit conversion, you can use explicit type conversion to avoid this problem. For example, after the fixed-length `bpchar` type is converted to the variable-length text type, spaces at the end of the character string are deleted. If the character string is compared with other character strings, errors may occur.

---

## Examples

Example 1: Use factorial operator type resolution. There is only one factorial operator (postfix !) defined in the system catalog, and it takes an argument of type **bigint**. The scanner assigns an initial type of **bigint** to the argument in this query expression:

```
openGauss=# SELECT 40 ! AS "40 factorial";
 40 factorial

8159152832478977343456112695961158942720000000000
(1 row)
```

So the parser does a type conversion on the operand and the query is equivalent to:

```
openGauss=# SELECT CAST(40 AS bigint) ! AS "40 factorial";
```

Example 2: String concatenation operator type resolution. A string-like syntax is used for working with string types and for working with complex extension types. Strings with unspecified type are matched with likely operator candidates. An example with one unspecified argument:

```
openGauss=# SELECT text 'abc' || 'def' AS "text and unknown";
text and unknown

abcdef
(1 row)
```

In this example, the parser looks for an operator whose parameters are of the text type. Such an operator is found.

Here is a concatenation of two values of unspecified types:

```
openGauss=# SELECT 'abc' || 'def' AS "unspecified";
unspecified

abcdef
(1 row)
```

### NOTE

In this case there is no initial hint for which type to use, since no types are specified in the query. So, the parser looks for all candidate operators and finds that there are candidates accepting both string-category and bit-string-category inputs. Since string category is preferred when available, that category is selected, and then the preferred type for strings, **text**, is used as the specific type to resolve the unknown-type literals as.

Example 3: Absolute-value and negation operator type resolution. The GaussDB operator catalog has several entries for the prefix operator @. All the entries implement absolute-value operations for various numeric data types. One of these entries is for type **float8**, which is the preferred type in the numeric category. Therefore, GaussDB will use that entry when faced with an unknown input:

```
openGauss=# SELECT @ '-4.5' AS "abs";
abs

4.5
(1 row)
```

Here the system has implicitly resolved the unknown-type literal as type **float8** before applying the chosen operator.



Example 4: Use the array inclusion operator type resolution as an example. Here is another example of resolving an operator with one known and one unknown input:

```
openGauss=# SELECT array[1,2] <@ '{1,2,3}' as "is subset";
is subset

t
(1 row)
```

#### NOTE

The GaussDB operator catalog has several entries for the infix operator <@, but the only two that could possibly accept an integer array on the left side are array inclusion (anyarray <@ anyarray) and range inclusion (anyelement <@ anyrange). Since none of these polymorphic pseudo-types (see [Pseudo-Types](#)) is considered preferred, the parser cannot resolve the ambiguity on that basis. However, the last resolution rule tells it to assume that the unknown-type literal is of the same type as the other input, that is, integer array. Now only one of the two operators can match, so array inclusion is selected. (Had range inclusion been selected, we would have gotten an error, because the string does not have the right format to be a range literal.)

## 7.7.3 Functions

### Function Type Resolution

1. Select the functions to be considered from the **pg\_proc** system catalog. If a non-schema-qualified function name was used, the functions in the current search path are considered. If a qualified function name was given, only functions in the specified schema are considered.  
If the search path finds multiple functions of different argument types, a proper function in the path is considered.
2. Check for a function accepting exactly the input argument types. If the function exists, use it. Cases involving **unknown** will never find a match at this step.
3. If no exact match is found, see if the function call appears to be a special type conversion request.
4. Look for the best match.
  - a. Discard candidate functions for which the input types do not match and cannot be converted (using an implicit conversion) to match. **unknown** literals are assumed to be convertible to anything for this purpose. If only one candidate remains, use it; else continue to the next step.
  - b. Run through all candidates and keep those with the most exact matches on input types. Domains are considered the same as their base type for this purpose. Keep all candidates if none has exact matches. If only one candidate remains, use it; else continue to the next step.
  - c. Run through all candidates and keep those that accept preferred types at the most positions where type conversion will be required. Keep all candidates if none accepts preferred types. If only one candidate remains, use it; else continue to the next step.
  - d. If any input arguments are **unknown**, check the type categories accepted at those argument positions by the remaining candidates. At each position, select the string category if any candidate accepts that category. (This bias towards string is appropriate since an unknown-type literal

looks like a string.) Otherwise, if all the remaining candidates accept the same type category, select that category; otherwise fail because the correct choice cannot be deduced without more clues. Now discard candidates that do not accept the selected type category. Furthermore, if any candidate accepts a preferred type in that category, discard candidates that accept non-preferred types for that argument. Keep all candidates if none survives these tests. If only one candidate remains, use it; else continue to the next step.

- e. If there are both **unknown** and known-type arguments, and all the known-type arguments have the same type, assume that the **unknown** arguments are also of that type, and check which candidates can accept that type at the **unknown**-argument positions. If exactly one candidate passes this test, use it. Otherwise, an error occurs.

## Examples

Example 1: Use the rounding function argument type resolution as the first example. There is only one **round** function that takes two arguments; it takes a first argument of type **numeric** and a second argument of type **integer**. So the following query automatically converts the first argument of type **integer** to **numeric**:

```
openGauss=# SELECT round(4, 4);
round

4.0000
(1 row)
```

That query is actually transformed by the parser to:

```
openGauss=# SELECT round(CAST (4 AS numeric), 4);
```

Since numeric constants with decimal points are initially assigned the type **numeric**, the following query will require no type conversion and therefore might be slightly more efficient:

```
openGauss=# SELECT round(4.0, 4);
```

Example 2: Use the substring function type resolution as the second example. There are several **substr** functions, one of which takes types **text** and **integer**. If called with a string constant of unspecified type, the system chooses the candidate function that accepts an argument of the preferred category **string** (namely of type **text**).

```
openGauss=# SELECT substr('1234', 3);
substr

34
(1 row)
```

If the string is declared to be of type **varchar**, as might be the case if it comes from a table, then the parser will try to convert it to become **text**:

```
openGauss=# SELECT substr(varchar '1234', 3);
substr

34
(1 row)
```

This is transformed by the parser to effectively become:

```
openGauss=# SELECT substr(CAST (varchar '1234' AS text), 3);
```

#### NOTE

The parser learns from the **pg\_cast** catalog that **text** and **varchar** are binary-compatible, meaning that one can be passed to a function that accepts the other without doing any physical conversion. Therefore, no type conversion is really inserted in this case.

And, if the function is called with an argument of type **integer**, the parser will try to convert that to **text**:

```
openGauss=# SELECT substr(1234, 3);
substr

34
(1 row)
```

This is transformed by the parser to effectively become:

```
openGauss=# SELECT substr(CAST (1234 AS text), 3);
substr

34
(1 row)
```

## 7.7.4 Value Storage

### Value Storage Type Resolution

1. Search for an exact match with the target column.
2. Try to convert the expression to the target type. This will succeed if there is a registered cast between the two types. If the expression is an unknown-type literal, the content of the literal string will be fed to the input conversion routine for the target type.
3. Check to see if there is a sizing cast for the target type. A sizing cast is a cast from that type to itself. If one is found in the **pg\_cast** catalog, apply it to the expression before storing into the destination column. The implementation function for such a cast always takes an extra parameter of type **integer**. The parameter receives the destination column's **atttypmod** value (typically its declared length, although the interpretation of **atttypmod** varies for different data types), and may take a third Boolean parameter that says whether the cast is explicit or implicit. The cast function is responsible for applying any length-dependent semantics such as size checking or truncation.

### Examples

Use the **character** storage type conversion as an example. For a target column declared as **character(20)** the following statement shows that the stored value is sized correctly:

```
openGauss=# CREATE SCHEMA tpcds;
openGauss=# CREATE TABLE tpcds.value_storage_t1 (
VS_COL1 CHARACTER(20)
);
openGauss=# INSERT INTO tpcds.value_storage_t1 VALUES('abcdef');
openGauss=# SELECT VS_COL1, octet_length(VS_COL1) FROM tpcds.value_storage_t1;
vs_col1 | octet_length
-----+-----
abcdef | 20
(1 row)
```

```
)
openGauss=# DROP TABLE tpceds.value_storage_t1;
openGauss=# DROP SCHEMA tpceds;
```

#### NOTE

What has happened here is that the two unknown literals are resolved to **text** by default, allowing the || operator to be resolved as **text** concatenation. Then the **text** result of the operator is converted to **bpchar** ("blank-padded char", the internal name of the **character** data type) to match the target column type. Since the conversion from **text** to **bpchar** is binary-coercible, this conversion does not insert any real function call. Finally, the sizing function **bpchar(bpchar, integer, Boolean)** is found in the system catalog and used for the operator's result and the stored column length. This type-specific function performs the required length check and addition of padding spaces.

## 7.7.5 UNION, CASE, and Related Constructs

SQL **UNION** constructs must match up possibly dissimilar types to become a single result set. The resolution algorithm is applied separately to each output column of a union query. The **INTERSECT** and **EXCEPT** construct resolve dissimilar types in the same way as **UNION**. The **CASE**, **ARRAY**, **VALUES**, **GREATEST** and **LEAST** constructs use the identical algorithm to match up their component expressions and select a result data type.

### Type Resolution for UNION, CASE, and Related Constructs

- If all inputs are of the same type, and it is not **unknown**, resolve as that type.
- If all inputs are of type **unknown**, resolve as type **text** (the preferred type of the string category). Otherwise, **unknown** inputs are ignored.
- If the inputs are not all of the same type category, a failure will be resulted. (Type **unknown** is not included in this case.)
- If the inputs are all of the same type category, choose the top preferred type in that category. (Exception: The **UNION** operation regards the type of the first branch as the selected type.)

#### NOTE

**typcategory** in the **pg\_type** system catalog indicates the data type category.  
**typispreferred** indicates whether a type is preferred in **typcategory**.

- Convert all inputs to the selected type. (Retain the original lengths of strings). Fail if there is not an implicit conversion from a given input to the selected type.
- If the input contains the json, txid\_snapshot, sys\_refcursor, or geometry type, **UNION** cannot be performed.

### Type Resolution for CASE and COALESCE in TD Compatibility Type

- If all inputs are of the same type, and it is not **unknown**, resolve as that type.
- If all inputs are of type **unknown**, resolve as type **text**.
- If inputs are of the string type (including **unknown** which is resolved as type **text**) and digit type, resolve as the string type. If the inputs are not of the two types, an error will be reported.
- If the inputs are all of the same type category, choose the top preferred type in that category.

- Convert all inputs to the selected type. Fail if there is not an implicit conversion from a given input to the selected type.

## Type Resolution for CASE in ORA Compatibility Type

**decode(expr, search1, result1, search2, result2, ..., defresult):** When the **sql\_beta\_feature** is set to **a\_style\_coerce**, the final return value type of the expression is set to the data type of result1 or a higher-precision data type in the same type as result1, as that in ORA-compatible mode. (For example, numeric and int are both numeric data types, but numeric has higher precision and priority than int.) For CASE WHEN, the behavior is the same as the default behavior in ORA-compatible mode.

- If all inputs are of the same type, and it is not **unknown**, resolve as that type. Otherwise, proceed to the next step.
- Set the data type of result1 to the final return value type preferType, which belongs to preferCategory.
- Consider the data types of result2, result3, and defresult in sequence. If the type category is also preferCategory, which is the same as that of result1, check whether the precision (priority) is higher than that of preferType. If it is, update preferType to a data type with a higher precision. If the type category is not preferCategory, check whether the category can be implicitly converted to preferType. If it cannot, an error is reported.
- Uses the data type recorded by preferType as the return value type of the expression. The expression result is implicitly converted to this data type.

Note:

There is a special case where the character type of a super-large number is converted to the numeric type, for example, **select decode(1, 2, 2, '53465465676465454657567678676')**, in which the large number exceeds the range of the bigint and double types. If result1 is of the numeric type and does not meet the condition that all inputs are of the same type, the type of the return value is set to numeric to be compatible with this special case.

Note 2:

Priority of the numeric types: numeric > float8 > float4 > int8 > int4 > int2 > int1

Priority of the character types: text > varchar = nvarchar2 > bpchar > char

Priority of date types: timestamptz > timestamp > smalldatetime > date > abstime > timetz > time

Priority of date span types: interval > tinterval > reltime.

Note 3:

The following figure shows the supported implicit type conversion when **set sql\_beta\_feature** is set to '**a\_style\_coerce**' in ORA-compatible mode. \ indicates that conversion is not required, **yes** indicates that conversion is supported, and the blank value indicates that conversion is not supported.

|               | bool | int1 | int2 | int4 | int8 | float4 | float8 | numeric | money | char | bpchar | varchar2 | nvarchar2 | text/clob | raw | blob | date | time | timetz | timestamp | timestampz | smalldatetime | interval | retime | abstime |
|---------------|------|------|------|------|------|--------|--------|---------|-------|------|--------|----------|-----------|-----------|-----|------|------|------|--------|-----------|------------|---------------|----------|--------|---------|
| bool          | \    |      |      |      |      |        |        |         |       |      |        |          |           |           |     |      |      |      |        |           |            |               |          |        |         |
| int1          |      | \    | yes  | yes  | yes  | yes    | yes    | yes     |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| int2          |      | yes  | \    | yes  | yes  | yes    | yes    | yes     |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| int4          |      |      | yes  | yes  | \    | yes    | yes    | yes     |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| int8          |      |      | yes  | yes  | yes  | \      | yes    | yes     |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| float4        |      |      | yes  | yes  | yes  | yes    | \      | yes     |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| float8        |      |      | yes  | yes  | yes  | yes    | yes    | \       | yes   |      | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| numeric       |      |      | yes  | yes  | yes  | yes    | yes    | \       |       | yes  | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| money         |      |      |      |      |      |        |        |         | \     |      |        |          |           |           |     |      |      |      |        |           |            |               |          |        |         |
| char          |      | yes  | yes  | yes  | yes  | yes    | yes    | yes     |       | \    | yes    | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| bpchar        |      | yes  | yes  | yes  | yes  | yes    | yes    | yes     |       | yes  | \      | yes      | yes       | yes       |     |      |      |      |        |           |            |               |          |        |         |
| varchar2      |      | yes  | yes  | yes  | yes  | yes    | yes    | yes     |       | yes  | yes    | \        | yes       | yes       | yes |      |      |      |        |           |            |               |          |        |         |
| nvarchar2     |      | yes  | yes  | yes  | yes  | yes    | yes    | yes     |       | yes  | yes    | yes      | \         | yes       |     |      |      |      |        |           |            |               |          |        |         |
| text/clob     |      | yes  | yes  | yes  | yes  | yes    | yes    | yes     |       | yes  | yes    | yes      | yes       | \         |     |      |      |      |        |           |            |               |          |        |         |
| raw           |      |      |      |      |      |        |        |         |       |      |        | yes      |           | yes       | \   | yes  |      |      |        |           |            |               |          |        |         |
| blob          |      |      |      |      |      |        |        |         |       |      |        |          |           | yes       | \   |      |      |      |        |           |            |               |          |        |         |
| date          |      |      |      |      |      |        |        |         |       |      | yes    | yes      | yes       | yes       |     |      | \    |      | yes    | yes       | yes        | yes           |          |        | yes     |
| time          |      |      |      |      |      |        |        |         |       |      |        |          |           | yes       |     |      |      | \    | yes    |           |            |               |          |        |         |
| timetz        |      |      |      |      |      |        |        |         |       |      |        |          |           | yes       |     |      |      | yes  | \      |           |            |               |          |        |         |
| timestamp     |      |      |      |      |      |        |        |         |       |      | yes    | yes      | yes       | yes       |     | yes  |      |      | \      |           | yes        | yes           |          |        | yes     |
| timestampz    |      |      |      |      |      |        |        |         |       |      |        |          | yes       | yes       |     | yes  |      |      | yes    | \         | yes        |               |          |        | yes     |
| smalldatetime |      |      |      |      |      |        |        |         |       |      |        | yes      |           | yes       |     | yes  |      |      | yes    | yes       | \          |               |          |        | yes     |
| interval      |      |      |      |      |      |        |        |         |       |      |        | yes      | yes       | yes       |     |      |      |      |        |           |            | \             | yes      |        |         |
| retime        |      |      |      |      |      |        |        |         |       |      |        |          |           | yes       |     |      |      |      |        |           |            | yes           | \        |        |         |
| abstime       |      |      |      |      |      |        |        |         |       |      |        |          |           | yes       |     | yes  |      |      |        | yes       | yes        | yes           |          |        | \       |

## Examples

Example 1: Use type resolution with underspecified types in a union as the first example. Here, the unknown-type literal **'b'** will be resolved to type **text**.

```
openGauss=# SELECT text 'a' AS "text" UNION SELECT 'b';
text

a
b
(2 rows)
```

Example 2: Use type resolution in a simple union as the second example. The literal **1.2** is of type **numeric**, and the **integer** value **1** can be cast implicitly to **numeric**, so that type is used.

```
openGauss=# SELECT 1.2 AS "numeric" UNION SELECT 1;
numeric

1
1.2
(2 rows)
```

Example 3: Use type resolution in a transposed union as the third example. Since type **real** cannot be implicitly cast to **integer**, but **integer** can be implicitly cast to **real**, the union result type is resolved as **real**.

```
openGauss=# SELECT 1 AS "real" UNION SELECT CAST('2.2' AS REAL);
real

1
2.2
(2 rows)
```

Example 4: In the **TD** type, if input parameters for **COALESCE** are of **int** and **varchar** types, resolve as type **varchar**. In the **A** type, an error is reported.

```
-- In the A type, create the a_1 database compatible with A.
openGauss=# CREATE DATABASE a_1 dbcompatibility = 'A';
```

```
-- Switch to the a_1 database.
openGauss=# \c a_1

-- Create the t1 table.
a_1=# CREATE TABLE t1(a int, b varchar(10));

-- Show the execution plan of a statement for querying the types int and varchar of input parameters for COALESCE.
a_1=# EXPLAIN SELECT coalesce(a, b) FROM t1;
ERROR: COALESCE types integer and character varying cannot be matched
LINE 1: EXPLAIN SELECT coalesce(a, b) FROM t1;
 ^
CONTEXT: referenced column: coalesce

-- Delete the table.
a_1=# DROP TABLE t1;

-- Switch to the testdb database.
a_1=# \c testdb

-- In the TD type, create the td_1 database compatible with Teradata.
openGauss=# CREATE DATABASE td_1 dbcompatibility = 'C';

-- Switch to the td_1 database.
openGauss=# \c td_1

-- Create the t2 table.
td_1=# CREATE TABLE t2(a int, b varchar(10));

-- Show the execution plan of a statement for querying the types int and varchar of input parameters for COALESCE.
td_1=# EXPLAIN VERBOSE select coalesce(a, b) from t2;
 QUERY PLAN

Data Node Scan (cost=0.00..0.00 rows=0 width=0)
 Output: (COALESCE((t2.a)::character varying, t2.b))
 Node/s: All dbnodes
 Remote query: SELECT COALESCE(a::character varying, b) AS "coalesce" FROM public.t2
(4 rows)

-- Delete the table.
td_1=# DROP TABLE t2;

-- Switch to the testdb database.
td_1=# \c testdb

-- Delete databases in A and TD types.
openGauss=# DROP DATABASE a_1;
openGauss=# DROP DATABASE td_1;
```

Example 5: In ORA-compatible mode, set the final return value type of the expression to the data type of result1 or a higher-precision data type whose category is the same as that of the data type of result1.

```
-- In the ORA type, create the ora_1 database compatible with ORA.
openGauss=# CREATE DATABASE ora_1 dbcompatibility = 'A';

-- Switch to the ora_1 database.
openGauss=# \c ora_1

-- Enable the decode compatibility parameters.
set sql_beta_feature='a_style_coerce';

-- Create the t1 table.
ora_1=# CREATE TABLE t1(c_int int, c_float8 float8, c_char char(10), c_text text, c_date date);

-- Insert data.
ora_1=# INSERT INTO t1 VALUES(1, 2, '3', '4', date '12-10-2010');
```

```
-- The data type of result1 is char and that of defresult is text. The precision of text is higher, and the type
of the return value is changed to text from char.
ora_1=# SELECT decode(1, 2, c_char, c_text) AS result, pg_typeof(result) FROM t1;
result | pg_typeof
-----+-----
4 | text
(1 row)

-- The data type of result1 is int, which is a numeric type. The type of the return value is set to numeric.
ora_1=# SELECT decode(1, 2, c_int, c_float8) AS result, pg_typeof(result) FROM t1;
result | pg_typeof
-----+-----
2 | numeric
(1 row)

-- The implicit conversion from the data type of defresult to that of result1 does not exist. An error is
reported.
ora_1=# SELECT decode(1, 2, c_int, c_date) FROM t1;
ERROR: CASE types integer and timestamp without time zone cannot be matched
LINE 1: SELECT decode(1, 2, c_int, c_date) FROM t1;
 ^
CONTEXT: referenced column: c_date

-- Disable the decode compatibility parameters.
set sql_beta_feature='none';

-- Delete the table.
ora_1=# DROP TABLE t1;
DROP TABLE

-- Switch to the testdb database.
ora_1=# \c testdb

-- Delete the ORA-compatible database.
openGauss=# DROP DATABASE ora_1;
DROP DATABASE
```

## 7.8 System Operation

GaussDB runs SQL statements to perform different system operations, such as setting variables, displaying the execution plan, and collecting garbage data.

### Setting Variables

For details about how to set various parameters for a session or transaction, see [SET](#).

### Displaying the Execution Plan

For details about how to display the execution plan that GaussDB makes for SQL statements, see [EXPLAIN](#).

### Specifying a Checkpoint in Transaction Logs

By default, WALs periodically specify checkpoints in a transaction log. **CHECKPOINT** forces an immediate checkpoint when the related command is issued, without waiting for a regular checkpoint scheduled by the system. For details, see [CHECKPOINT](#).



## Collecting Unnecessary Data

For details about how to collect garbage data and analyze a database as required, see [VACUUM](#).

## Collecting Statistics

For details about how to collect statistics on tables in databases, see [ANALYZE | ANALYZE](#).

## Setting the Constraint Check Mode for the Current Transaction

For details about how to set the constraint check mode for the current transaction, see [SET CONSTRAINTS](#).

## Shutting Down The Current Database Node

For details about shutting down the current database node, see [SHUTDOWN](#).

# 7.9 Controlling Transactions

A transaction is a user-defined sequence of database operations, which form an integral unit of work.

## Starting a Transaction

GaussDB starts a transaction using **START TRANSACTION** and **BEGIN**. For details, see [START TRANSACTION](#) and [BEGIN](#).

## Setting a Transaction

GaussDB sets a transaction using **SET TRANSACTION** or **SET LOCAL TRANSACTION**. For details, see [SET TRANSACTION](#).

## Committing a Transaction

GaussDB commits all operations of a transaction using **COMMIT** or **END**. For details, see [COMMIT | END](#).

## Rolling Back a Transaction

If a fault occurs during a transaction and the transaction cannot proceed, the system performs rollback to cancel all the completed database operations related to the transaction. See [ROLLBACK](#).

### NOTE

If an execution request (not in a transaction block) received in the database contains multiple statements, the request is packed into a transaction. If one of the statements fails, the entire request will be rolled back.

## 7.10 DDL Syntax Overview

Data definition language (DDL) is used to define or modify an object in a database, such as a table, an index, or a view.

### NOTE

GaussDB does not support DDL when the primary node of the database is incomplete. For example, if the primary node of the database in the database is faulty, creating a database or a table will fail.

### Defining a Database

A database is the warehouse for organizing, storing, and managing data. Defining a database includes creating a database, altering the database attributes, and deleting the database. For details about related SQL statements, see [Table 7-86](#).

**Table 7-86** SQL statements for defining a database

| Function                      | SQL Statement                   |
|-------------------------------|---------------------------------|
| Creating a database           | <a href="#">CREATE DATABASE</a> |
| Altering a database attribute | <a href="#">ALTER DATABASE</a>  |
| Deleting a database           | <a href="#">DROP DATABASE</a>   |

### Defining a schema

A schema is the set of a group of database objects and is used to control the access to the database objects. For details about related SQL statements, see [Table 7-87](#).

**Table 7-87** SQL statements for defining a schema

| Function                    | SQL Statement                 |
|-----------------------------|-------------------------------|
| Creating a schema           | <a href="#">CREATE SCHEMA</a> |
| Altering a schema attribute | <a href="#">ALTER SCHEMA</a>  |
| Deleting a schema           | <a href="#">DROP SCHEMA</a>   |

### Defining a Tablespace

A tablespace is used to manage data objects and corresponds to a catalog on a disk. For details about related SQL statements, see [Table 7-88](#).

**Table 7-88** SQL statements for defining a tablespace

| Function                        | SQL Statement                     |
|---------------------------------|-----------------------------------|
| Creating a tablespace           | <a href="#">CREATE TABLESPACE</a> |
| Altering a tablespace attribute | <a href="#">ALTER TABLESPACE</a>  |
| Deleting a tablespace           | <a href="#">DROP TABLESPACE</a>   |

## Defining a Table

A table is a special data structure in a database and is used to store data objects and relationship between data objects. For details about related SQL statements, see [Table 7-89](#).

**Table 7-89** SQL statements for defining a table

| Function                   | SQL Statement                |
|----------------------------|------------------------------|
| Creating a table           | <a href="#">CREATE TABLE</a> |
| Altering a table attribute | <a href="#">ALTER TABLE</a>  |
| Deleting a table           | <a href="#">DROP TABLE</a>   |

## Defining a Partitioned Table

A partitioned table is a logical table used to improve query performance and does not store data (data is stored in ordinary tables). For details about related SQL statements, see [Table 7-90](#).

**Table 7-90** SQL statements for defining a partitioned table

| Function                               | SQL Statement                          |
|----------------------------------------|----------------------------------------|
| Creating a partitioned table           | <a href="#">CREATE TABLE PARTITION</a> |
| Creating a partition                   | <a href="#">ALTER TABLE PARTITION</a>  |
| Altering a partitioned table attribute | <a href="#">ALTER TABLE PARTITION</a>  |
| Deleting a partition                   | <a href="#">ALTER TABLE PARTITION</a>  |
| Deleting a partitioned table           | <a href="#">DROP TABLE</a>             |

## Defining an Index

An index indicates the sequence of values in one or more columns in a database table. It is a data structure that improves the speed of data access to specific

information in a database table. For details about related SQL statements, see [Table 7-91](#).

**Table 7-91** SQL statements for defining an index

| Function                    | SQL Statement                |
|-----------------------------|------------------------------|
| Creating an index           | <a href="#">CREATE INDEX</a> |
| Altering an index attribute | <a href="#">ALTER INDEX</a>  |
| Deleting an index           | <a href="#">DROP INDEX</a>   |
| Rebuilding an index         | <a href="#">REINDEX</a>      |

## Defining a Stored Procedure

A stored procedure is a set of SQL statements for achieving specific functions and is stored in the database after compiling. Users can specify a name and provide parameters (if necessary) to execute the stored procedure. For details about related SQL statements, see [Table 7-92](#).

**Table 7-92** SQL statements for defining a stored procedure

| Function                    | SQL Statement                    |
|-----------------------------|----------------------------------|
| Creating a stored procedure | <a href="#">CREATE PROCEDURE</a> |
| Deleting a stored procedure | <a href="#">DROP PROCEDURE</a>   |

## Defining a Function

In GaussDB, a function is similar to a stored procedure, which is a set of SQL statements. The function and stored procedure are used the same. For details about related SQL statements, see [Table 7-93](#).

**Table 7-93** SQL statements for defining a function

| Function                      | SQL Statement                   |
|-------------------------------|---------------------------------|
| Creating a function           | <a href="#">CREATE FUNCTION</a> |
| Altering a function attribute | <a href="#">ALTER FUNCTION</a>  |
| Deleting a function           | <a href="#">DROP FUNCTION</a>   |

## Defining a Package

A package consists of the package specification and package body. It is used to manage stored procedures and functions by class, which is similar to classes in languages such as Java and C++.

**Table 7-94** SQL statements for defining a package

| Function                     | SQL Statement         |
|------------------------------|-----------------------|
| Creating a package           | <b>CREATE PACKAGE</b> |
| Deleting a package           | <b>DROP PACKAGE</b>   |
| Altering a package attribute | <b>ALTER PACKAGE</b>  |

## Defining a View

A view is a virtual table exported from one or more basic tables. It is used to control data accesses of users. [Table 7-95](#) lists the related SQL statements.

**Table 7-95** SQL statements for defining a view

| Function        | SQL Statement      |
|-----------------|--------------------|
| Creating a view | <b>CREATE VIEW</b> |
| Deleting a view | <b>DROP VIEW</b>   |

## Defining a Cursor

To process SQL statements, the stored procedure process allocates a memory fragment to store the context. Cursors are handles or pointers to context regions. With a cursor, the stored procedure can control alterations in context areas. For details, see [Table 7-96](#).

**Table 7-96** SQL statements for defining a cursor

| Function                    | SQL Statement |
|-----------------------------|---------------|
| Creating a cursor           | <b>CURSOR</b> |
| Moving a cursor             | <b>MOVE</b>   |
| Fetching data from a cursor | <b>FETCH</b>  |
| Closing a cursor            | <b>CLOSE</b>  |

## Defining an Aggregate Function

**Table 7-97** SQL statements for defining an aggregate function

| Function                        | SQL Statement                    |
|---------------------------------|----------------------------------|
| Creating an aggregate function  | <a href="#">CREATE AGGREGATE</a> |
| Modifying an aggregate function | <a href="#">ALTER AGGREGATE</a>  |
| Deleting an aggregate function  | <a href="#">DROP AGGREGATE</a>   |

## Defining Data Type Conversion

**Table 7-98** SQL statements for defining a data type

| Function                                   | SQL Statement               |
|--------------------------------------------|-----------------------------|
| Creating user-defined data type conversion | <a href="#">CREATE CAST</a> |
| Deleting user-defined data type conversion | <a href="#">DROP CAST</a>   |

## Defining an Operator

**Table 7-99** SQL statements for defining an operator

| Function              | SQL Statement                   |
|-----------------------|---------------------------------|
| Creating an operator  | <a href="#">CREATE OPERATOR</a> |
| Modifying an operator | <a href="#">ALTER OPERATOR</a>  |
| Deleting an operator  | <a href="#">DROP OPERATOR</a>   |

## Defining a Data Type

**Table 7-100** SQL statements for defining a data type

| Function              | SQL Statement               |
|-----------------------|-----------------------------|
| Creating a data type  | <a href="#">CREATE TYPE</a> |
| Modifying a data type | <a href="#">ALTER TYPE</a>  |
| Deleting a data type  | <a href="#">DROP TYPE</a>   |

## 7.11 DML Syntax Overview

Data manipulation language (DML) is used to perform operations on data in database tables, such as inserting, updating, querying, or deleting data.

### Inserting Data

Inserting data refers to adding one or multiple records to a database table. For details, see [INSERT](#).

### Updating Data

Updating data refers to modifying one or multiple records in a database table. For details, see [UPDATE](#).

### Querying Data

The database query statement **SELECT** is used to search required information in a database. For details, see [SELECT](#).

### Deleting Data

GaussDB provides two statements for deleting data from database tables. To delete data meeting specified conditions from a database table, see [DELETE](#). To delete all data from a database table, see [TRUNCATE](#).

**TRUNCATE** can quickly delete all data from a database table, which achieves the effect same as that running **DELETE** to delete data without specifying conditions from each table. Deletion efficiency using **TRUNCATE** is faster because **TRUNCATE** does not scan tables. Therefore, **TRUNCATE** is useful in large tables.

### Copying Data

GaussDB provides a statement for copying data between tables and files. For details, see [COPY](#).

### Locking a Table

GaussDB provides multiple lock modes to control concurrent accesses to table data. For details, see [LOCK](#).

### Calling a Function

GaussDB provides three statements for calling functions. These statements are the same in the syntax structure. For details, see [CALL](#).

### Session Management

A session is a connection established between the user and the database. [Table 7-101](#) lists the related SQL statements.

**Table 7-101** SQL statements related to sessions

| Function           | SQL Statement                             |
|--------------------|-------------------------------------------|
| Altering a session | <a href="#">ALTER SESSION</a>             |
| Killing a session  | <a href="#">ALTER SYSTEM KILL SESSION</a> |

## 7.12 DCL Syntax Overview

Data control language (DCL) is used to create users and roles and set or modify database users or role rights.

### Defining a Role

A role is used to manage permissions. For database security, management and operation permissions can be granted to different roles. For details about related SQL statements, see [Table 7-102](#).

**Table 7-102** SQL statements for defining a role

| Description              | SQL Statement               |
|--------------------------|-----------------------------|
| Creating a role          | <a href="#">CREATE ROLE</a> |
| Altering role attributes | <a href="#">ALTER ROLE</a>  |
| Dropping a role          | <a href="#">DROP ROLE</a>   |

### Defining a User

A user is used to log in to a database. Different permissions can be granted to users for managing data accesses and operations of the users. For details about related SQL statements, see [Table 7-103](#).

**Table 7-103** SQL statements for defining a user

| Description              | SQL Statement               |
|--------------------------|-----------------------------|
| Creating a User          | <a href="#">CREATE USER</a> |
| Altering user attributes | <a href="#">ALTER USER</a>  |
| Dropping a user          | <a href="#">DROP USER</a>   |

### Granting Rights

GaussDB provides a statement for granting rights to data objects and roles. For details, see [GRANT](#).



## Revoking Rights

GaussDB provides a statement for revoking rights. For details, see [REVOKE](#).

## Setting Default Rights

GaussDB allows users to set rights for objects that will be created. For details, see [ALTER DEFAULT PRIVILEGES](#)

## Shutting Down The Current Node

GaussDB allows users to run the **shutdown** command to shut down the current database node. For details, see [SHUTDOWN](#).

# 7.13 SQL Syntax

## 7.13.1 ABORT

### Description

Rolls back the current transaction and cancels the changes in the transaction.

This command is equivalent to [ROLLBACK](#), and is present only for historical reasons. Now ROLLBACK is recommended.

### Precautions

ABORT has no impact outside a transaction, but will throw a NOTICE message.

### Syntax

```
ABORT [WORK | TRANSACTION] ;
```

### Parameters

- **WORK | TRANSACTION**  
Specifies an optional keyword, which has no effect except increasing readability.

### Examples

```
-- Create the customer_demographics_t1 table.
openGauss=# CREATE TABLE customer_demographics_t1
(
 CD_DEMO_SK INTEGER NOT NULL,
 CD_GENDER CHAR(1) ,
 CD_MARITAL_STATUS CHAR(1) ,
 CD_EDUCATION_STATUS CHAR(20) ,
 CD_PURCHASE_ESTIMATE INTEGER ,
 CD_CREDIT_RATING CHAR(10) ,
 CD_DEP_COUNT INTEGER ,
 CD_DEP_EMPLOYED_COUNT INTEGER ,
 CD_DEP_COLLEGE_COUNT INTEGER ,
)
;
```

```
-- Insert data.
openGauss=# INSERT INTO customer_demographics_t1 VALUES(1920801,'M', 'U', 'DOCTOR DEGREE', 200,
'GOOD', 1, 0,0);

-- Start a transaction.
openGauss=# START TRANSACTION;

-- Update the column.
openGauss=# UPDATE customer_demographics_t1 SET cd_education_status= 'Unknown';

-- Abort the transaction. All updates are rolled back.
openGauss=# ABORT;

-- Query data.
openGauss=# SELECT * FROM customer_demographics_t1 WHERE cd_demo_sk = 1920801;
cd_demo_sk | cd_gender | cd_marital_status | cd_education_status | cd_purchase_estimate | cd_credit_rating
| cd_dep_count | cd_dep_employed_count | cd_dep_college_count
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1920801 | M | U | DOCTOR DEGREE | 200 | GOOD | 1
| 0 | 0
(1 row)
```

```
-- Drop the table.
openGauss=# DROP TABLE customer_demographics_t1;
```

## Helpful Links

[SET TRANSACTION](#), [COMMIT | END](#), and [ROLLBACK](#)

## 7.13.2 ALTER AGGREGATE

### Description

Alters the definition of an aggregate function, including the name, owner, and schema.

### Precautions

To use ALTER AGGREGATE, you must be the owner of the aggregate function. To change the schema of an aggregate function, you must have the CREATE permission on the new schema. To change the owner, you must be a direct or indirect member of the new role, and the role must have the CREATE permission on the aggregate function's schema. (This restricts the owner from doing anything except for deleting and rebuilding aggregate functions. However, a user with the SYSADMIN permission can change the ownership of an aggregate function in any way.)

### Syntax

```
ALTER AGGREGATE name (argtype [, ...]) RENAME TO new_name
ALTER AGGREGATE name (argtype [, ...]) OWNER TO new_owner
ALTER AGGREGATE name (argtype [, ...]) SET SCHEMA new_schema
```

### Parameters

- **name**  
Name (optionally schema-qualified) of an existing aggregate function.

- **argtype**  
Input data type of the aggregate function. To reference a zero-parameter aggregate function, you can write an asterisk (\*) instead of a list of input data types.
- **new\_name**  
New name of the aggregate function.
- **new\_owner**  
New owner of the aggregate function.
- **new\_schema**  
New schema of the aggregate function.

## Examples

Rename the aggregate function myavg that accepts integer-type parameters to **my\_average**.

```
ALTER AGGREGATE myavg(integer) RENAME TO my_average;
```

Change the owner of the aggregate function myavg that accepts integer-type parameters to **joe**.

```
ALTER AGGREGATE myavg(integer) OWNER TO joe;
```

Move the aggregate function myavg that accepts integer-type parameters to **myschema**.

```
ALTER AGGREGATE myavg(integer) SET SCHEMA myschema;
```

## Compatibility

The SQL standard does not contain the ALTER AGGREGATE statement.

## Helpful Links

[CREATE AGGREGATE](#) and [DROP AGGREGATE](#)

## 7.13.3 ALTER AUDIT POLICY

### Description

Alters the unified audit policy.

### Precautions

- Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.
- The unified audit policy takes effect only after **enable\_security\_policy** is enabled.

### Syntax

```
ALTER AUDIT POLICY [IF EXISTS] policy_name { ADD | REMOVE } { [privilege_audit_clause]
[access_audit_clause] };
ALTER AUDIT POLICY [IF EXISTS] policy_name MODIFY (filter_group_clause);
```

```
ALTER AUDIT POLICY [IF EXISTS] policy_name DROP FILTER;
ALTER AUDIT POLICY [IF EXISTS] policy_name COMMENTS policy_comments;
ALTER AUDIT POLICY [IF EXISTS] policy_name { ENABLE | DISABLE };
```

- **privilege\_audit\_clause**  
PRIVILEGES ({ DDL | ALL } [ ON LABEL ( resource\_label\_name [, ... ] ) ] )
- **access\_audit\_clause**  
ACCESS ({ DML | ALL } [ ON LABEL ( resource\_label\_name [, ... ] ) ] )
- **filter\_group\_clause**  
FILTER ON { filter\_type ( filter\_value [, ... ] ) } [, ... ]

## Parameters

- **policy\_name**  
Specifies the audit policy name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **DDL**  
Specifies the operations that are audited in the database: **CREATE**, **ALTER**, **DROP**, **ANALYZE**, **COMMENT**, **GRANT**, **REVOKE**, **SET**, **SHOW**, **LOGIN\_ANY**, **LOGIN\_FAILURE**, **LOGIN\_SUCCESS**, and **LOGOUT**.
- **DML**  
Specifies the operations that are audited in the database: **SELECT**, **COPY**, **DEALLOCATE**, **DELETE**, **EXECUTE**, **INSERT**, **PREPARE**, **REINDEX**, **TRUNCATE**, and **UPDATE**.
- **ALL**  
Specifies all operations supported by the specified DDL or DML statements in the database. When the form is { DDL | ALL }, **ALL** indicates all DDL operations. When the form is { DML | ALL }, **ALL** indicates all DML operations.
- **filter\_type**  
Specifies the types of information to be filtered by the policy: **IP**, **ROLES**, and **APP**.
- **filter\_value**  
Specifies the detailed information to be filtered.
- **policy\_comments**  
Records description information of the audit policy.
- **ENABLE|DISABLE**  
Enables or disables the unified audit policy. If **ENABLE|DISABLE** is not specified, **ENABLE** is used by default.

## Examples

See [Examples](#) in "CREATE AUDIT POLICY."

## Helpful Links

[CREATE AUDIT POLICY](#) and [DROP AUDIT POLICY](#)

## 7.13.4 ALTER DATABASE

### Description

Alters a database, including its name, owner, object isolation, and connection limitation.

### Precautions

- Only the database owner or a user granted with the ALTER permission can run the **ALTER DATABASE** command. System administrators have this permission by default. The following are permission restrictions depending on attributes to be modified:
  - To change the database name, you must have the CREATEDB permission.
  - To change a database owner, you must be a database owner or system administrator and a member of the new owner role, with the CREATEDB permission.
  - To change the default tablespace of the database, the user must have the CREATE permission for creating tablespaces. This statement physically migrates tables and indexes in a default tablespace to a new tablespace. Note that tables and indexes outside the default tablespace are not affected.
- You are not allowed to rename a database in use. To rename it, connect to another database.

### Syntax

- Modify the maximum number of connections to the database.

```
ALTER DATABASE database_name
[WITH] CONNECTION LIMIT connlimit;
```

- Rename the database.

```
ALTER DATABASE database_name
RENAME TO new_name;
```

- Change the database owner.

```
ALTER DATABASE database_name
OWNER TO new_owner;
```

- Change the default tablespace of the database.

```
ALTER DATABASE database_name
SET TABLESPACE new_tablespace;
```

#### NOTE

If some tables or objects in the database have been created in **new\_tablespace**, the default tablespace of the database cannot be changed to **new\_tablespace**. An error will be reported during the execution.

- Modify the session parameter value of the database.

```
ALTER DATABASE database_name
SET configuration_parameter { { TO | = } { value | DEFAULT } | FROM CURRENT };
```

- Reset the database configuration parameter.

```
ALTER DATABASE database_name RESET
{ configuration_parameter | ALL };
```

- Modify the object isolation attribute of the database.

```
ALTER DATABASE database_name [WITH] { ENABLE | DISABLE } PRIVATE OBJECT;
```

 **NOTE**

- To modify the object isolation attribute of a database, the database must be connected. Otherwise, the modification will fail.
- For a new database, the object isolation attribute is disabled by default. After this attribute is enabled, common users can view only the objects (such as tables, functions, views, and columns) that they have the permission to access. This attribute does not take effect for administrators. After this attribute is enabled, administrators can still view all database objects.

## Parameters

- **database\_name**  
Specifies the name of the database whose attributes are to be modified.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **connlimit**  
Specifies the maximum number of concurrent connections that can be made to this database (excluding administrators' connections).  
Value range: an integer ranging from  $-1$  to  $2^{31} - 1$ . You are advised to set this parameter to an integer ranging from 1 to 50. The default value  $-1$  indicates that there is no restriction on the number of concurrent connections.
- **new\_name**  
Specifies the new name of a database.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **new\_owner**  
Specifies the new owner of a database.  
Value range: a string. It must be a valid username.
- **new\_tablespace**  
Specifies the new default tablespace of a database. The tablespace exists in the database. The default tablespace is pg\_default.  
Value range: a string. It must be a valid tablespace name.
- **configuration\_parameter**
  - **value**  
Sets a specified database session parameter to a specified value. If the value is **DEFAULT** or **RESET**, the default setting is used in the new session. **OFF** disables the setting.  
Value range: a string.
    - **DEFAULT**
    - **OFF**
    - **RESET**
  - **FROM CURRENT**  
Uses the value of **configuration\_parameter** of the current session.
- **FROM CURRENT**  
Sets the value of the database based on the current connected session.

- **RESET configuration\_parameter**  
Resets the specified database session parameter.
- **RESET ALL**  
Resets all database session parameters.

 **NOTE**

- Modify the default tablespace of a database by moving the table or index in the old tablespace into the new tablespace. This operation does not affect the tables or indexes in other non-default tablespaces.
- The modified database session parameter values will take effect in the next session.

## Examples

See [Examples](#) in "CREATE DATABASE."

## Helpful Links

[CREATE DATABASE](#) and [DROP DATABASE](#)

## 7.13.5 ALTER DEFAULT PRIVILEGES

### Description

Allows you to set the permissions that will be applied to objects created in the future. (It does not affect permissions granted to existing objects.)

### Precautions

Only the permissions for tables (including views), sequences, functions, and types can be changed.

### Syntax

```
ALTER DEFAULT PRIVILEGES
[FOR { ROLE | USER } target_role [, ...]]
[IN SCHEMA schema_name [, ...]]
abbreviated_grant_or_revoke;
```

- The abbreviated\_grant\_or\_revoke clause grants or revokes permissions on some objects.

```
grant_on_tables_clause
| grant_on_sequences_clause
| grant_on_functions_clause
| grant_on_types_clause
| revoke_on_tables_clause
| revoke_on_sequences_clause
| revoke_on_functions_clause
| revoke_on_types_clause
```

- grant\_on\_tables\_clause grants permissions on tables.  
GRANT { { SELECT | INSERT | UPDATE | DELETE | TRUNCATE | REFERENCES | ALTER | DROP |  
COMMENT | INDEX | VACUUM }  
[, ...] | ALL [ PRIVILEGES ] }  
ON TABLES  
TO { [ GROUP ] role\_name | PUBLIC } [, ...]  
[ WITH GRANT OPTION ]
- grant\_on\_sequences\_clause grants permissions on sequences.

```
GRANT { { SELECT | UPDATE | USAGE | ALTER | DROP | COMMENT }
[, ...] | ALL [PRIVILEGES] }
ON SEQUENCES
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION]
```

- **grant\_on\_functions\_clause** grants permissions on functions.

```
GRANT { { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON FUNCTIONS
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION]
```

- **grant\_on\_types\_clause** grants permissions on types.

```
GRANT { { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON TYPES
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION]
```

- **revoke\_on\_tables\_clause** revokes permissions on tables.

```
REVOKE [GRANT OPTION FOR]
{ { SELECT | INSERT | UPDATE | DELETE | TRUNCATE | REFERENCES | ALTER | DROP | COMMENT |
INDEX | VACUUM }
[, ...] | ALL [PRIVILEGES] }
ON TABLES
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT | CASCADE CONSTRAINTS]
```

- **revoke\_on\_sequences\_clause** revokes permissions on sequences.

```
REVOKE [GRANT OPTION FOR]
{ { SELECT | UPDATE | USAGE | ALTER | DROP | COMMENT }
[, ...] | ALL [PRIVILEGES] }
ON SEQUENCES
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT | CASCADE CONSTRAINTS]
```

- **revoke\_on\_functions\_clause** revokes permissions on functions.

```
REVOKE [GRANT OPTION FOR]
{ { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON FUNCTIONS
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT | CASCADE CONSTRAINTS]
```

- **revoke\_on\_types\_clause** revokes permissions on types.

```
REVOKE [GRANT OPTION FOR]
{ { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON TYPES
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT | CASCADE CONSTRAINTS]
```

## Parameters

- **target\_role**

Specifies the name of an existing role. If **FOR ROLE/USER** is omitted, the current role is assumed.

Value range: an existing role name.

- **schema\_name**

Specifies the name of an existing schema.

**target\_role** must have the CREATE permission for **schema\_name**.

Value range: an existing schema name.

- **role\_name**

Specifies the name of an existing role to grant or revoke permissions for.

Value range: an existing role name.



**NOTICE**

To drop a role for which the default permissions have been granted, reverse the changes in its default permissions or use DROP OWNED BY to get rid of the default permission entry for the role.

## Examples

```
-- Grant the SELECT permission on all the tables (and views) in tpcds to every user.
openGauss=# ALTER DEFAULT PRIVILEGES IN SCHEMA tpcds GRANT SELECT ON TABLES TO PUBLIC;

-- Create a common user jack.
openGauss=# CREATE USER jack PASSWORD '*****';

-- Grant the INSERT permission on all the tables in tpcds to the user jack.
openGauss=# ALTER DEFAULT PRIVILEGES IN SCHEMA tpcds GRANT INSERT ON TABLES TO jack;

-- Revoke the preceding permissions.
openGauss=# ALTER DEFAULT PRIVILEGES IN SCHEMA tpcds REVOKE SELECT ON TABLES FROM PUBLIC;
openGauss=# ALTER DEFAULT PRIVILEGES IN SCHEMA tpcds REVOKE INSERT ON TABLES FROM jack;

-- Drop user jack.
openGauss=# DROP USER jack CASCADE;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds;
```

## Helpful Links

[GRANT](#) and [REVOKE](#)

## 7.13.6 ALTER DIRECTORY

### Description

Alters directory attributes.

### Precautions

- Currently, only the directory owner can be changed.
- When **enable\_access\_server\_directory** is set to **off**, only the initial user is allowed to change the directory owner. When **enable\_access\_server\_directory** is set to **on**, users with the SYSADMIN permission and the directory object owner can change the directory object owner, and the user who changes the owner is required to be a member of the new owner.

### Syntax

```
ALTER DIRECTORY directory_name
 OWNER TO new_owner;
```

### Parameters

- directory\_name**  
Specifies the name of a directory to be modified. The value must be an existing directory name.

- **new\_owner**  
Specifies the new owner of the directory.

## Examples

```
-- Create a directory.
openGauss=# CREATE OR REPLACE DIRECTORY dir as '/tmp/';

-- Create a user.
openGauss=# CREATE USER system PASSWORD '*****!';

-- Change the owner of the directory.
openGauss=# ALTER DIRECTORY dir OWNER TO system;

-- Drop a directory.
openGauss=# DROP DIRECTORY dir;
```

## Helpful Links

[CREATE DIRECTORY](#) and [DROP DIRECTORY](#)

## 7.13.7 ALTER FUNCTION

### Description

Alters the attributes of a custom function.

### Precautions

Only the function owner or a user granted the ALTER permission can run the **ALTER FUNCTION** command. System administrators have this permission by default. The following are permission restrictions depending on attributes to be modified:

- If a function involves operations on temporary tables, ALTER FUNCTION cannot be used.
- To modify the owner or schema of a function, you must be a function owner or system administrator and a member of the new owner role.
- Only system administrators and the initial user can change the schema of a function to public.

### Syntax

- Modify the additional parameters of the user-defined function.  
**ALTER FUNCTION** function\_name ( [ { [ argname ] [ argmode ] argtype } [, ...] ] )  
action [ ... ] [ RESTRICT ];

The syntax of the ACTION clause is as follows:

```
{CALLED ON NULL INPUT | RETURNS NULL ON NULL INPUT | STRICT}
| {IMMUTABLE | STABLE | VOLATILE}
| {SHIPPABLE | NOT SHIPPABLE}
| {NOT FENCED | FENCED}
| [NOT] LEAKPROOF
| { [EXTERNAL] SECURITY INVOKER | [EXTERNAL] SECURITY DEFINER }
| AUTHID { DEFINER | CURRENT_USER }
| COST execution_cost
| ROWS result_rows
| SET configuration_parameter { { TO | = } { value | DEFAULT } } FROM CURRENT}
| RESET {configuration_parameter | ALL}
```

- Rename the user-defined function.  

```
ALTER FUNCTION function_name ([{ [argname] [argmode] argtype} [, ...]])
 RENAME TO new_name;
```
- Change the owner of the user-defined function.  

```
ALTER FUNCTION function_name ([{ [argname] [argmode] argtype} [, ...]])
 OWNER TO new_owner;
```
- Modify the schema of the user-defined function.  

```
ALTER FUNCTION function_name ([{ [argname] [argmode] argtype} [, ...]])
 SET SCHEMA new_schema;
```

## Parameters

- **function\_name**  
Specifies the name of the function to be modified.  
Value range: an existing function name.
- **argmode**  
Specifies whether a parameter is an input or output parameter.  
Value range:
  - **IN**: declares input parameters.
  - **OUT**: declares output parameters.
  - **INOUT**: declares input and output parameters.
  - **VARIADIC**: declares parameters of the array type.
- **argname**  
Specifies the parameter name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **argtype**  
Specifies the parameter type of the function.
- **CALLED ON NULL INPUT**  
Declares that some parameters of the function can be called in normal mode if the parameter values are **NULL**. Omitting this parameter is the same as specifying it.
- **RETURNS NULL ON NULL INPUT**  
**STRICT**  
Specifies that a function always returns **NULL** when the value of any of its parameters is **NULL**. If this parameter is specified, the function is not executed when there is **NULL** parameter; instead a NULL result is assumed automatically.  
**RETURNS NULL ON NULL INPUT** and **STRICT** have the same functions.
- **IMMUTABLE**  
Specifies that the function always returns the same result if the parameter values are the same.
- **STABLE**  
Specifies that the function cannot modify the database, and that within a single table scan it will consistently return the same result for the same parameter value, but its result varies by SQL statements.
- **VOLATILE**

Specifies that the function value can change in a single table scan and no optimization is performed.

- **LEAKPROOF**

Specifies that the function has no side effect and the parameter contains only the return value. **LEAKPROOF** can be set only by system administrators.

- **EXTERNAL**

(Optional) The purpose is to be compatible with SQL. This feature applies to all functions, not only external functions.

- **SECURITY INVOKER**

**AUTHID CURRENT\_USER**

Specifies that the function will be executed with the permissions of the user who calls it. Omitting this parameter is the same as specifying it.

**SECURITY INVOKER** and **AUTHID CURRENT\_USER** have the same functions.

- **SECURITY DEFINER**

**AUTHID DEFINER**

Specifies that the function will be executed with the permissions of the user who created it.

**AUTHID DEFINER** and **SECURITY DEFINER** have the same functions.

- **COST execution\_cost**

Estimates the execution cost of a function.

The unit of **execution\_cost** is **cpu\_operator\_cost**.

Value range: a positive integer.

- **ROWS result\_rows**

Estimates the number of rows returned by the function. This is only allowed when the function is declared to return a set.

Value range: a positive number. The default value is **1000**.

- **configuration\_parameter**

- **value**

Sets a specified database session parameter to a specified value. If the value is **DEFAULT** or **RESET**, the default setting is used in the new session. **OFF** disables the setting.

Value range: a string.

- **DEFAULT**

- **OFF**

- **RESET**

- Specified default value

- **FROM CURRENT**

Uses the value of **configuration\_parameter** of the current session.

- **new\_name**

Specifies the new name of a function. To change the schema of a function, you must have the CREATE permission on the new schema.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **new\_owner**

Specifies the new owner of a function. To change the owner of a function, the new owner must have the CREATE permission on the schema to which the function belongs.

Value range: an existing user role.

- **new\_schema**

Specifies the new schema of a function.

Value range: an existing schema.

## Examples

See [Examples](#) in "CREATE FUNCTION."

## Helpful Links

[CREATE FUNCTION](#) and [DROP FUNCTION](#)

## 7.13.8 ALTER GLOBAL CONFIGURATION

### Description

Adds or changes the **key-value** of the gs\_global\_config system catalog. You can modify an existing parameter or add a new one.

### Precautions

- Only the initial database user can run this command.
- The parameter name cannot be **weak\_password** or **undostoragetype**.

### Syntax

```
ALTER GLOBAL CONFIGURATION with(name=value, name=value...);
```

### Parameters

- **name**  
Parameter name, which is of the text type. The value cannot be **weak\_password** or **undostoragetype**.
- **value**  
Parameter value, which is of the text type.

## Examples

```
-- Insert.
openGauss=# ALTER GLOBAL CONFIGURATION with(redis_is_ok = true);

-- Query.
openGauss=# SELECT * FROM gs_global_config;
 name | value
-----+-----
 buckets_len | 16384
undostoragetype | page
```

```
redis_is_ok | true
(3 rows)

-- Alter.
openGauss=# ALTER GLOBAL CONFIGURATION with(redis_is_ok = false);

-- Query.
openGauss=# SELECT * FROM gs_global_config;
 name | value
-----+-----
 buckets_len | 16384
 undostoragetype | page
 redis_is_ok | false
(3 rows)

-- Drop.
openGauss=# DROP GLOBAL CONFIGURATION redis_is_ok;

-- Query.
openGauss=# SELECT * FROM gs_global_config;
 name | value
-----+-----
 buckets_len | 16384
 undostoragetype | page
(2 rows)
```

## Helpful Links

### [DROP GLOBAL CONFIGURATION](#)

## 7.13.9 ALTER GROUP

### Description

Alters the attributes of a user group.

### Precautions

ALTER GROUP is an alias for ALTER ROLE, and it is not a standard SQL command and not recommended. Users can use ALTER ROLE directly.

### Syntax

- Add users to a group.  

```
ALTER GROUP group_name
 ADD USER user_name [, ...];
```
- Remove users from a group.  

```
ALTER GROUP group_name
 DROP USER user_name [, ...];
```
- Change the name of the group.  

```
ALTER GROUP group_name
 RENAME TO new_name;
```

### Parameters

See [Parameters](#) in "ALTER ROLE."

### Examples

```
-- Create a user group.
openGauss=# CREATE GROUP super_users WITH PASSWORD '*****';
```

```
-- Create a user.
openGauss=# CREATE ROLE lche WITH PASSWORD '*****';

-- Create a user.
openGauss=# CREATE ROLE jim WITH PASSWORD '*****';

-- Add users to a group.
openGauss=# ALTER GROUP super_users ADD USER lche, jim;

-- Remove users from a group.
openGauss=# ALTER GROUP super_users DROP USER jim;

-- Change the name of the group.
openGauss=# ALTER GROUP super_users RENAME TO normal_users;

-- Drop users.
openGauss=# DROP ROLE lche, jim;

-- Drop the user group.
openGauss=# DROP GROUP normal_users;
```

## Helpful Links

[CREATE GROUP](#), [DROP GROUP](#), and [ALTER ROLE](#)

## 7.13.10 ALTER INDEX

### Description

ALTER INDEX modifies the definition of an existing index.

It has the following forms:

- IF EXISTS  
Sends a notice instead of an error if the specified index does not exist.
- RENAME TO  
Changes only the name of the index. The stored data is not affected.
- SET TABLESPACE  
This option changes the index tablespace to the specified tablespace and moves index-related data files to the new tablespace.
- SET ( { STORAGE\_PARAMETER = value } [, ...] )  
Changes one or more index-method-specific storage parameters of an index. Note that the index content will not be modified immediately by this statement. You may need to use REINDEX to rebuild the index based on different parameters to achieve the expected effect.
- RESET ( { storage\_parameter } [, ...] )  
Resets one or more index-method-specific storage parameters of an index to the default value. Similar to the SET statement, REINDEX may be used to completely update the index.
- [ MODIFY PARTITION index\_partition\_name ] UNUSABLE  
Sets the indexes on a table or index partition to be unavailable.
- REBUILD [ PARTITION index\_partition\_name ]  
Rebuilds indexes on a table or an index partition.

- **RENAME PARTITION**  
Renames an index partition.
- **MOVE PARTITION**  
Modifies the tablespace to which an index partition belongs.

## Precautions

Only the index owner or a user who has the INDEX permission on the table where the index resides can run the **ALTER INDEX** command. System administrators have this permission by default.

## Syntax

- Rename a table index.  

```
ALTER INDEX [IF EXISTS] index_name
 RENAME TO new_name;
```
- Alter the tablespace to which a table index belongs.  

```
ALTER INDEX [IF EXISTS] index_name
 SET TABLESPACE tablespace_name;
```
- Modify the storage parameter of a table index.  

```
ALTER INDEX [IF EXISTS] index_name
 SET ({storage_parameter = value} [, ...]);
```
- Reset the storage parameter of a table index.  

```
ALTER INDEX [IF EXISTS] index_name
 RESET (storage_parameter [, ...]);
```
- Set a table index or an index partition to be unavailable.  

```
ALTER INDEX [IF EXISTS] index_name
 [MODIFY PARTITION index_partition_name] UNUSABLE;
```
- Rebuild a table index or index partition.  

```
ALTER INDEX index_name
 REBUILD [PARTITION index_partition_name];
```
- Rename an index partition.  

```
ALTER INDEX [IF EXISTS] index_name
 RENAME PARTITION index_partition_name TO new_index_partition_name;
```
- Modify the tablespace to which an index partition belongs.  

```
ALTER INDEX [IF EXISTS] index_name
 MOVE PARTITION index_partition_name TABLESPACE new_tablespace;
```

## Parameters

- **index\_name**  
Specifies the index name to be modified.
- **new\_name**  
Specifies the new name of the index.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **tablespace\_name**  
Specifies the tablespace name.  
Value range: an existing tablespace name.
- **storage\_parameter**  
Specifies the name of an index-method-specific parameter.
- **value**



Specifies the new value for an index-method-specific storage parameter. This might be a number or a word depending on the parameter.

- **new\_index\_partition\_name**  
Specifies the new name of an index partition.
- **index\_partition\_name**  
Specifies the name of an index partition.
- **new\_tablespace**  
Specifies a new tablespace.

## Examples

See [Examples](#) in "CREATE INDEX."

## Helpful Links

[CREATE INDEX](#), [DROP INDEX](#), and [REINDEX](#)

## 7.13.11 ALTER LANGUAGE

This version does not support this syntax.

## 7.13.12 ALTER MASKING POLICY

### Description

Alters a masking policy.

### Precautions

- Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.
- The masking policy takes effect only after **enable\_security\_policy** is enabled.

### Syntax

- Modify the policy description.  
`ALTER MASKING POLICY policy_name COMMENTS policy_comments;`
- Modify the masking method.  
`ALTER MASKING POLICY policy_name [ADD | REMOVE | MODIFY] masking_actions[, ...]*;`  
The syntax of **masking\_action**:  
`masking_function ON LABEL(label_name[, ...]*)`
- Modify the scenarios where the masking policy takes effect.  
`ALTER MASKING POLICY policy_name MODIFY(FILTER ON FILTER_TYPE(filter_value[, ...]*)[, ...]*);`
- Remove the filters of the masking policy.  
`ALTER MASKING POLICY policy_name DROP FILTER;`
- Enable or disable the masking policy.  
`ALTER MASKING POLICY policy_name [ENABLE | DISABLE];`

### Parameters

- **policy\_name**  
Specifies the name of a masking policy, which must be unique.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **policy\_comments**

Adds or modifies description of the masking policy.

- **masking\_function**

Specifies eight preset masking methods or user-defined functions. Schemas are supported.

The maskall function is not preset. It is hard-coded and cannot be displayed by running `\df`.

The masking methods during presetting are as follows:

```
maskall | randommasking | creditcardmasking | basicemailmasking | fullemailmasking |
shufflemasking | alldigitsmasking | regexprmasking
```

- **label\_name**

Specifies the resource label name.

- **FILTER\_TYPE**

Specifies the types of information to be filtered by the masking policy, including **IP**, **ROLES**, and **APP**.

- **filter\_value**

Specifies the detailed information to be filtered, such as the IP address, app name, and username.

- **ENABLE|DISABLE**

Enables or disables the masking policy. If **ENABLE|DISABLE** is not specified, **ENABLE** is used by default.

## Examples

```
-- Create users dev_mask and bob_mask.
openGauss=# CREATE USER dev_mask PASSWORD '*****';
openGauss=# CREATE USER bob_mask PASSWORD '*****';

-- Create table tb_for_masking.
openGauss=# CREATE TABLE tb_for_masking(col1 text, col2 text, col3 text);

-- Create a resource label for sensitive column col1.
openGauss=# CREATE RESOURCE LABEL mask_lb1 ADD COLUMN(tb_for_masking.col1);

-- Create a resource label for sensitive column col2.
openGauss=# CREATE RESOURCE LABEL mask_lb2 ADD COLUMN(tb_for_masking.col2);

-- Create a masking policy for the operation of accessing sensitive column col1.
openGauss=# CREATE MASKING POLICY maskpol1 maskall ON LABEL(mask_lb1);

-- Add description for masking policy maskpol1.
openGauss=# ALTER MASKING POLICY maskpol1 COMMENTS 'masking policy for tb_for_masking.col1';

-- Modify masking policy maskpol1 to add a masking method.
openGauss=# ALTER MASKING POLICY maskpol1 ADD randommasking ON LABEL(mask_lb2);

-- Modify masking policy maskpol1 to remove a masking method.
openGauss=# ALTER MASKING POLICY maskpol1 REMOVE randommasking ON LABEL(mask_lb2);

-- Modify masking policy maskpol1 to modify a masking method.
openGauss=# ALTER MASKING POLICY maskpol1 MODIFY randommasking ON LABEL(mask_lb1);

-- Modify masking policy maskpol1 so that it takes effect only for scenarios where users are dev_mask and bob_mask, the client tool is gsql, and the IP addresses are 10.20.30.40 and 127.0.0.0/24.
openGauss=# ALTER MASKING POLICY maskpol1 MODIFY (FILTER ON ROLES(dev_mask, bob_mask),
APP(gsql), IP('10.20.30.40', '127.0.0.0/24'));
```

```
-- Modify masking policy maskpol1 so that it takes effect for all user scenarios.
openGauss=# ALTER MASKING POLICY maskpol1 DROP FILTER;

-- Disable masking policy maskpol1.
openGauss=# ALTER MASKING POLICY maskpol1 DISABLE;

-- Drop a masking policy.
openGauss=# DROP MASKING POLICY maskpol1;

-- Drop resource labels.
openGauss=# DROP RESOURCE LABEL mask_lb1, mask_lb2;

-- Drop the tb_for_masking table.
openGauss=# DROP TABLE tb_for_masking;

-- Drop the dev_mask and bob_mask users.
openGauss=# DROP USER dev_mask, bob_mask;
```

## Helpful Links

[CREATE MASKING POLICY](#) and [DROP MASKING POLICY](#)

## 7.13.13 ALTER MATERIALIZED VIEW

### Description

Alters multiple auxiliary attributes of an existing materialized view.

Statements and actions that can be used for ALTER MATERIALIZED VIEW are a subset of ALTER TABLE and have the same meaning when used for materialized views. For details, see [ALTER TABLE](#).

### Precautions

- Only the owner of a materialized view or a system administrator has the ALTER MATERIALIZED VIEW permission.
- The materialized view structure cannot be modified.

### Syntax

- Alter the owner of the materialized view.  

```
ALTER MATERIALIZED VIEW [IF EXISTS] mv_name
OWNER TO new_owner;
```
- Modify the column of a materialized view.  

```
ALTER MATERIALIZED VIEW [IF EXISTS] mv_name
RENAME [COLUMN] column_name TO new_column_name;
```
- Rename a materialized view.  

```
ALTER MATERIALIZED VIEW [IF EXISTS] mv_name
RENAME TO new_name;
```

### Parameters

- **mv\_name**  
Specifies the name of an existing materialized view, which can be schema-qualified.  
Value range: a string that complies with the [Identifier Naming Conventions](#).

- **column\_name**  
Specifies the name of a new or existing column.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **new\_column\_name**  
Specifies the new name of an existing column.
- **new\_owner**  
Specifies the username of the new owner of a materialized view.
- **new\_name**  
Specifies the new name of a materialized view.

## Examples

```
-- Create a table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int) WITH(STORAGE_TYPE=ASTORE);

-- Create a complete-refresh materialized view.
openGauss=# CREATE MATERIALIZED VIEW foo AS SELECT * FROM my_table;

-- Rename the materialized view foo to bar.
openGauss=# ALTER MATERIALIZED VIEW foo RENAME TO bar;

-- Drop a complete-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW bar;

-- Drop the my_table table.
openGauss=# DROP TABLE my_table;
```

## Helpful Links

[CREATE INCREMENTAL MATERIALIZED VIEW](#), [CREATE MATERIALIZED VIEW](#),  
[DROP MATERIALIZED VIEW](#), [REFRESH INCREMENTAL MATERIALIZED VIEW](#),  
and [REFRESH MATERIALIZED VIEW](#)

## 7.13.14 ALTER OPERATOR

### Description

Alters the definition of an operator.

### Precautions

ALTER OPERATOR changes the definition of an operator. Currently, the only function available is to change the owner of the operator.

To use ALTER OPERATOR, you must be the owner of the operator. To modify the owner, you must also be a direct or indirect member of the new owning role, and that member must have CREATE permission on the operator's schema. (These restrict that the owner cannot be changed by doing anything other than deleting or rebuilding the operator. However, a user with the SYSADMIN permission can change the ownership of any operator in any way.)

### Syntax

```
ALTER OPERATOR name ({ left_type | NONE } , { right_type | NONE }) OWNER TO new_owner
ALTER OPERATOR name ({ left_type | NONE } , { right_type | NONE }) SET SCHEMA new_schema
```

## Parameters

- **name**  
Name of an existing operator.
- **left\_type**  
Data type of the left operand for the operator; if there is no left operand, write **NONE**.
- **right\_type**  
Data type of the right operand for the operator; if there is no right operand, write **NONE**.
- **new\_owner**  
New owner of the operator.
- **new\_schema**  
New schema name of the operator.

## Examples

Alter a user-defined operator for text a @@ b:

```
ALTER OPERATOR @@ (text, text) OWNER TO joe;
```

## Compatibility

The SQL standard does not contain the ALTER OPERATOR statement.

## Helpful Links

[CREATE OPERATOR](#), [CREATE OPERATOR CLASS](#), and [DROP OPERATOR](#)

## 7.13.15 ALTER PACKAGE

### Description

Alters the attributes of a package.

### Precautions

Currently, only users with the ALTER PACKAGE OWNER permission can run this command. By default, system administrators have the permission. The restrictions are as follows:

- The current user must be the owner of the package or a system administrator and a member of the new owner role.

### Syntax

- Alter the owner of a package.  

```
ALTER PACKAGE package_name OWNER TO new_owner;
```

### Parameters

- **package\_name**

Specifies the name of the package to be modified.

Value range: an existing package name. Only one package can be modified at a time.

- **new\_owner**

Specifies the new owner of a package. To change the owner of a package, the new owner must have the CREATE permission on the schema to which the package belongs.

Value range: an existing user role.

## Examples

For details, see [CREATE PACKAGE](#).

## Helpful Links

[CREATE PACKAGE](#) and [DROP PACKAGE](#)

## 7.13.16 ALTER PROCEDURE

### Description

Modify the owner of a user-defined stored procedure.

### Precautions

Only the owner of a stored procedure or a user granted with the ALTER permission can run the **ALTER PROCEDURE** command. System administrators have this permission by default. The following are permission restrictions depending on attributes to be modified:

- If a stored procedure involves operations on temporary tables, ALTER PROCEDURE cannot be used.
- To modify the owner or schema of a stored procedure, you must be the owner of the stored procedure or system administrator and a member of the new owner role.
- Only system administrators and the initial user can change the schema of a stored procedure to a public schema.

### Syntax

- Modify the additional parameters of a user-defined stored procedure.

```
ALTER PROCEDURE procedure_name ([{ [argname] [argmode] argtype } [, ...]])
action [...] [RESTRICT];
```

The syntax of the ACTION clause is as follows:

```
{ CALLED ON NULL INPUT | STRICT }
| { IMMUTABLE | STABLE | VOLATILE }
| { SHIPPABLE | NOT SHIPPABLE }
| { NOT FENCED | FENCED }
| [NOT] LEAKPROOF
| { [EXTERNAL] SECURITY INVOKER | [EXTERNAL] SECURITY DEFINER }
| AUTHID { DEFINER | CURRENT_USER }
| COST execution_cost
| ROWS result_rows
| SET configuration_parameter { { TO | = } { value | DEFAULT } } FROM CURRENT}
| RESET { configuration_parameter | ALL }
```

- Modify the name of a user-defined stored procedure.  

```
ALTER PROCEDURE proname ([{ [argname] [argmode] argtype} [, ...]])
 RENAME TO new_name;
```
- Modify the owner of a user-defined stored procedure.  

```
ALTER PROCEDURE proname ([{ [argname] [argmode] argtype} [, ...]])
 OWNER TO new_owner;
```
- Modify the schema of a user-defined stored procedure.  

```
ALTER PROCEDURE proname ([{ [argname] [argmode] argtype} [, ...]])
 SET SCHEMA new_schema;
```

## Parameters

- **procedure\_name**  
Specifies the name of the stored procedure to be modified.  
Value range: an existing stored procedure name.
- **argmode**  
Specifies whether a parameter is an input or output parameter.  
Value range: **IN**, **OUT**, **INOUT**, and **VARIADIC**.
- **argname**  
Specifies the parameter name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **argtype**  
Specifies the type of the stored procedure parameter.
- **CALLED ON NULL INPUT**  
Declares that some parameters of the stored procedure can be called in normal mode if the parameter values are **NULL**. Omitting this parameter is the same as specifying it.
- **IMMUTABLE**  
Specifies that the stored procedure always returns the same result if the parameter values are the same.
- **STABLE**  
Specifies that the stored procedure cannot modify the database, and that within a single table scan it will consistently return the same result for the same parameter value, but its result varies by SQL statements.
- **VOLATILE**  
Specifies that the stored procedure value can change in a single table scan and no optimization is performed.
- **LEAKPROOF**  
Specifies that the stored procedure has no side effect and the parameter contains only the return value. **LEAKPROOF** can be set only by system administrators.
- **EXTERNAL**  
(Optional) The purpose is to be compatible with SQL. This feature applies to all functions, not only external functions.
- **SECURITY INVOKER**  
**AUTHID CURRENT\_USER**

Specifies that the stored procedure will be executed with the permissions of the user who calls it. Omitting this parameter is the same as specifying it.

**SECURITY INVOKER** and **AUTHID CURRENT\_USER** have the same functions.

- **SECURITY DEFINER**

**AUTHID DEFINER**

Specifies that the stored procedure will be executed with the permissions of the user who created it.

**AUTHID DEFINER** and **SECURITY DEFINER** have the same functions.

- **COST execution\_cost**

Estimates the execution cost of the stored procedure.

The unit of **execution\_cost** is **cpu\_operator\_cost**.

Value range: a positive integer.

- **ROWS result\_rows**

Estimates the number of rows returned by the stored procedure. This is only allowed when the stored procedure is declared to return a set.

Value range: a positive number. The default value is **1000**.

- **configuration\_parameter**

- **value**

Sets a specified database session parameter to a specified value. If the value is **DEFAULT** or **RESET**, the default setting is used in the new session. **OFF** disables the setting.

Value range: a string.

- **DEFAULT**
    - **OFF**
    - **RESET**
    - Specified default value

- **FROM CURRENT**

Uses the value of **configuration\_parameter** of the current session.

- **new\_name**

Specifies the new name of the stored procedure. To change the schema of a stored procedure, you must have the CREATE permission on the new schema.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **new\_owner**

Specifies the new owner of the stored procedure. To change the owner of a stored procedure, the new owner must have the CREATE permission on the schema to which the stored procedure belongs.

Value range: an existing user role.

- **new\_schema**

Specifies the new schema of the stored procedure.

Value range: an existing schema.



## Examples

See [Examples](#) in "CREATE FUNCTION."

## Helpful Links

[CREATE PROCEDURE](#) and [DROP PROCEDURE](#)

## 7.13.17 ALTER RESOURCE LABEL

### Description

Alters resource labels.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

### Syntax

```
ALTER RESOURCE LABEL label_name (ADD|REMOVE)
label_item_list[, ...]*;
```

- **label\_item\_list**  
resource\_type(resource\_path[, ...]\*)
- **resource\_type**  
TABLE | COLUMN | SCHEMA | VIEW | FUNCTION

### Parameters

- **label\_name**  
Specifies the resource label name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **resource\_type**  
Specifies the type of database resources to be labeled.
- **resource\_path**  
Specifies the path of database resources.

## Examples

```
-- Create basic table table_for_label.
openGauss=# CREATE TABLE table_for_label(col1 int, col2 text);

-- Create resource label table_label.
openGauss=# CREATE RESOURCE LABEL table_label ADD COLUMN(table_for_label.col1);

-- Attach resource label table_label to col2.
openGauss=# ALTER RESOURCE LABEL table_label ADD COLUMN(table_for_label.col2)

-- Remove table_label from an item.
openGauss=# ALTER RESOURCE LABEL table_label REMOVE COLUMN(table_for_label.col1);

-- Drop the resource label table_label.
openGauss=# DROP RESOURCE LABEL table_label;
```

```
-- Drop the base table table_for_label.
openGauss=# DROP TABLE table_for_label;
```

## Helpful Links

[CREATE RESOURCE LABEL](#) and [DROP RESOURCE LABEL](#)

## 7.13.18 ALTER ROLE

### Description

Alters role attributes.

### Precautions

None

### Syntax

- Modify the permissions of a role.

```
ALTER ROLE role_name [[WITH] option [...]];
```

The option clause for granting permissions is as follows:

```
{CREATEDB | NOCREATEDB}
| {CREATEROLE | NOCREATEROLE}
| {INHERIT | NOINHERIT}
| {AUDITADMIN | NOAUDITADMIN}
| {SYSADMIN | NOSYSADMIN}
| {MONADMIN | NOMONADMIN}
| {OPRADMIN | NOOPRADMIN}
| {POLADMIN | NOPOLADMIN}
| {USEFT | NOUSEFT}
| {LOGIN | NOLOGIN}
| {REPLICATION | NOREPLICATION}
| {VCADMIN | NOVCADMIN}
| {PERSISTENCE | NOPERSISTENCE}
| CONNECTION LIMIT connlimit
| [ENCRYPTED | UNENCRYPTED] PASSWORD 'password' [EXPIRED]
| [ENCRYPTED | UNENCRYPTED] IDENTIFIED BY 'password' [REPLACE 'old_password' | EXPIRED]
| [ENCRYPTED | UNENCRYPTED] PASSWORD { 'password' | DISABLE | EXPIRED }
| [ENCRYPTED | UNENCRYPTED] IDENTIFIED BY { 'password' [REPLACE 'old_password'] |
DISABLE }
| VALID BEGIN 'timestamp'
| VALID UNTIL 'timestamp'
| PERM SPACE 'spacelimit'
| PGUSER
```

- Rename a role.

```
ALTER ROLE role_name
RENAME TO new_name;
```

- Lock or unlock.

```
ALTER ROLE role_name
ACCOUNT { LOCK | UNLOCK };
```

- Set parameters for a role.

```
ALTER ROLE role_name [IN DATABASE database_name]
SET configuration_parameter {{ TO | = } { value | DEFAULT } | FROM CURRENT};
```

- Reset parameters for a role.

```
ALTER ROLE role_name
[IN DATABASE database_name] RESET {configuration_parameter|ALL};
```

## Parameters

- **role\_name**  
Specifies a role name.  
Value range: an existing role name. If a role name contains uppercase letters, enclose the name with double quotation marks ("").
- **IN DATABASE database\_name**  
Modifies the parameters of a role in a specified database.
- **SET configuration\_parameter {{ TO | = } { value | DEFAULT } | FROM CURRENT}**  
Sets parameters for a role. Session parameters modified by **ALTER ROLE** apply to a specified role and take effect in the next session triggered by the role.  
Value range:  
For details about the values of **configuration\_parameter** and **value**, see [SET](#).  
**DEFAULT**: clears the value of **configuration\_parameter**.  
**configuration\_parameter** will inherit the default value of the new session generated for the role.  
**FROM CURRENT**: uses the value of **configuration\_parameter** of the current session.
- **RESET {configuration\_parameter|ALL}**  
Clears the value of **configuration\_parameter**. The statement has the same effect as that of **SET configuration\_parameter TO DEFAULT**.  
Value range: **ALL** indicates that the values of all parameters are cleared.
- **ACCOUNT LOCK | ACCOUNT UNLOCK**
  - **ACCOUNT LOCK**: locks an account to forbid login to databases.
  - **ACCOUNT UNLOCK**: unlocks an account to allow login to databases.
- **PGUSER**  
In the current version, the PGUSER permission of a role cannot be modified.
- **PASSWORD/IDENTIFIED BY 'password'**  
Resets or changes the user password. Except the initial user, other administrators and common users need to enter the correct old password when changing their own passwords. Only the initial user or users with the SYSADMIN or CREATEROLE permission can reset the password of a common user without entering the old password. The initial user can reset passwords of system administrators. A system administrator cannot reset passwords of other system administrators. The user password should be enclosed in single quotation marks ('').
- **EXPIRED**  
Invalidates the password. Only the initial user or users with the SYSADMIN or CREATEROLE permission can invalidate user passwords. A system administrator can invalidate their own passwords or the passwords of other system administrators. The password of the initial user cannot be invalidated.  
The user whose password is invalid can log in to the database but cannot perform the query operation. The query operation can be performed only after the password is changed or an administrator resets the password.

For details about other parameters, see [Parameters](#) in "CREATE ROLE."

## Examples

See [Examples](#) in "CREATE ROLE."

## Helpful Links

[CREATE ROLE](#), [DROP ROLE](#), and [SET ROLE](#)

# 7.13.19 ALTER ROW LEVEL SECURITY POLICY

## Description

Alters an existing row-level security policy, including the policy name and the users and expressions affected by the policy.

## Precautions

Only the table owner or a system administrator can perform this operation.

## Syntax

```
ALTER [ROW LEVEL SECURITY] POLICY [IF EXISTS] policy_name ON table_name RENAME TO
new_policy_name;
```

```
ALTER [ROW LEVEL SECURITY] POLICY policy_name ON table_name
[TO { role_name | PUBLIC } [, ...]]
[USING (using_expression)];
```

## Parameters

- **policy\_name**  
Specifies the name of a row-level security policy.
- **table\_name**  
Specifies the name of a table to which a row-level security policy is applied.
- **new\_policy\_name**  
Specifies the new name of a row-level security policy.
- **role\_name**  
Specifies names of users affected by a row-level security policy. PUBLIC indicates that the row-level security policy will affect all users.
- **using\_expression**  
Specifies an expression defined for a row-level security policy. The return value is of the Boolean type.

## Examples

```
-- Create data table all_data.
openGauss=# CREATE TABLE all_data(id int, role varchar(100), data varchar(100));

-- Create a row-level security policy to specify that the current user can view only their own data.
openGauss=# CREATE ROW LEVEL SECURITY POLICY all_data_rls ON all_data USING(role =
CURRENT_USER);
openGauss=# \d+ all_data
 Table "public.all_data"
Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
```

```
id | integer | | plain | |
role | character varying(100) | | extended | |
data | character varying(100) | | extended | |
Row Level Security Policies:
 POLICY "all_data_rls"
 USING (((role)::name = "current_user"()))
Has OIDs: no
Location Nodes: ALL DATANODES
Options: orientation=row

-- Create users alice and bob.
openGauss=# CREATE ROLE alice WITH PASSWORD '*****';
openGauss=# CREATE ROLE bob WITH PASSWORD '*****';

-- Change the name of the all_data_rls policy.
openGauss=# ALTER ROW LEVEL SECURITY POLICY all_data_rls ON all_data RENAME TO all_data_new_rls;

-- Change the users affected by the row-level security policy.
openGauss=# ALTER ROW LEVEL SECURITY POLICY all_data_new_rls ON all_data TO alice, bob;
openGauss=# \d+ all_data
 Table "public.all_data"
 Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
id | integer | | plain | |
role | character varying(100) | | extended | |
data | character varying(100) | | extended | |
Row Level Security Policies:
 POLICY "all_data_new_rls"
 TO alice,bob
 USING (((role)::name = "current_user"()))
Has OIDs: no
Location Nodes: ALL DATANODES
Options: orientation=row, enable_rowsecurity=true

-- Modify the expression defined for the row-level security policy.
openGauss=# ALTER ROW LEVEL SECURITY POLICY all_data_new_rls ON all_data USING (id > 100 AND role
= current_user);
openGauss=# \d+ all_data
 Table "public.all_data"
 Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
id | integer | | plain | |
role | character varying(100) | | extended | |
data | character varying(100) | | extended | |
Row Level Security Policies:
 POLICY "all_data_new_rls"
 TO alice,bob
 USING (((id > 100) AND ((role)::name = "current_user"())))
Has OIDs: no
Location Nodes: ALL DATANODES
Options: orientation=row, enable_rowsecurity=true

-- Drop users alice and bob.
openGauss=# DROP ROLE alice, bob;

-- Drop the policy.
openGauss=# DROP ROW LEVEL SECURITY POLICY all_data_new_rls ON all_data;

-- Drop the all_data table.
openGauss=# DROP TABLE all_data;
```

## Helpful Links

[CREATE ROW LEVEL SECURITY POLICY](#) and [DROP ROW LEVEL SECURITY POLICY](#)

## 7.13.20 ALTER SCHEMA

### Description

Alters schema attributes.

### Precautions

- Only the owner of a schema or users granted the ALTER permission on the schema can run the **ALTER SCHEMA** command. System administrators have this permission by default. To change the owner of a schema, you must be the owner of the schema or a system administrator and a member of the new owner role.
- Only the initial user is allowed to change the owner of the pg\_catalog system schema. Changing the names of the built-in system schemas may make some functions unavailable or even affect the normal running of the database. By default, the names of the built-in system schemas cannot be changed. To ensure forward compatibility, you can change the names of the built-in system schemas only when the system is being started or upgraded or when **allow\_system\_table\_mods** is set to **on**.

### Syntax

- Rename a schema.  

```
ALTER SCHEMA schema_name
 RENAME TO new_name;
```
- Change the owner of a schema.  

```
ALTER SCHEMA schema_name
 OWNER TO new_owner;
```

### Parameters

- **schema\_name**  
Specifies the name of an existing schema.  
Value range: an existing schema name.
- **RENAME TO new\_name**  
Renames a schema. If a non-SYSADMIN user wants to change the schema name, the user must have the CREATE permission on the database.  
**new\_name**: new name of the schema.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **OWNER TO new\_owner**  
Changes the owner of a schema. To do this as a non-system administrator, you must be a direct or indirect member of the new owner role, and that role must have the CREATE permission on the database.  
**new\_owner**: new owner of the schema.  
Value range: an existing username or role name.

### Examples

```
-- Create the ds schema.
openGauss=# CREATE SCHEMA ds;
```

```
-- Rename the current schema ds to ds_new.
openGauss=# ALTER SCHEMA ds RENAME TO ds_new;

-- Create user jack.
openGauss=# CREATE USER jack PASSWORD '*****';

-- Change the owner of ds_new to jack.
openGauss=# ALTER SCHEMA ds_new OWNER TO jack;

-- Drop user jack and schema ds_new.
openGauss=# DROP SCHEMA ds_new;
openGauss=# DROP USER jack;
```

## Helpful Links

[CREATE SCHEMA](#) and [DROP SCHEMA](#)

### 7.13.21 ALTER SEQUENCE

#### Description

Alters the parameters of an existing sequence.

#### Precautions

- Only the sequence owner or a user granted with the ALTER permission can run the **ALTER SEQUENCE** command. System administrators have this permission by default. To modify a sequence owner, you must be the sequence owner or a system administrator and a member of the new owner role.
- In the current version, you can modify only the owner, owning column, and maximum value. To change other parameters, drop the sequence and create it again. Then, use the Setval function to restore parameter values.
- ALTER SEQUENCE MAXVALUE cannot be used in transactions, functions, or stored procedures.
- After the maximum value of a sequence is changed, the cache of the sequence in all sessions is cleared.
- If the LARGE identifier is used when a sequence is created, the LARGE identifier must be used when the sequence is altered.
- The ALTER SEQUENCE statement blocks the calling of nextval, setval, curval, and lastval.

#### Syntax

- Change the owning column of a sequence.  

```
ALTER [LARGE] SEQUENCE [IF EXISTS] name
 [MAXVALUE maxvalue | NO MAXVALUE | NOMAXVALUE | CACHE cache]
 [OWNED BY { table_name.column_name | NONE }] ;
```
- Change the owner of a sequence.  

```
ALTER [LARGE] SEQUENCE [IF EXISTS] name OWNER TO new_owner;
```

#### Parameters

- **name**  
Specifies the name of the sequence to be modified.

- **IF EXISTS**  
Sends a notice instead of an error when you are modifying a nonexistent sequence.
- **CACHE**  
Specifies the number of sequences stored in the memory for quick access purposes. If this parameter is not specified, the old cache value is retained.
- **OWNED BY**  
Joins a sequence with a specified column included in a table. In this way, the sequence will be dropped when you drop its joined column or the table where the column belongs to.  
  
If the sequence has been joined with another table before you use this option, the new join will overwrite the old one.  
  
The joined table and sequence must be owned by the same user and in the same schema.  
  
If **OWNED BY NONE** is used, all existing joins will be dropped.
- **new\_owner**  
Specifies the username of the new owner of the sequence. To change the owner, you must also be a direct or indirect member of the new role, and this role must have the CREATE permission on the sequence's schema.

## Examples

```
-- Create an ascending sequence named serial, starting from 101.
openGauss=# CREATE SEQUENCE serial START 101;

-- Create a table and specify default values for the sequence.
openGauss=# CREATE TABLE T1(C1 bigint default nextval('serial'));

-- Change the owning column of serial to T1.C1.
openGauss=# ALTER SEQUENCE serial OWNED BY T1.C1;

-- Drop a sequence and a table.
openGauss=# DROP SEQUENCE serial cascade;
openGauss=# DROP TABLE T1;
```

## Helpful Links

[CREATE SEQUENCE](#) and [DROP SEQUENCE](#)

## 7.13.22 ALTER SERVER

### Description

Adds, alters, or deletes the parameters of an existing server. You can query existing servers from the pg\_foreign\_server system catalog.

### Precautions

Only the server owner or a user granted the ALTER permission can run the **ALTER SERVER** command. By default, system administrators have the permission. To modify a server owner, you must be the server owner or system administrator and a member of the new owner role.



When multi-layer quotation marks are used for sensitive columns (such as **password** and **secret\_access\_key**) in **OPTIONS**, the semantics is different from that in the scenario where quotation marks are not used. Therefore, sensitive columns are not identified for anonymization.

## Syntax

- Change the parameters for a foreign server.

```
ALTER SERVER server_name [VERSION 'new_version']
[OPTIONS ({[ADD | SET | DROP] option ['value']} [, ...])];
```

In **OPTIONS**, **ADD**, **SET**, and **DROP** are operations to be performed. If these operations are not specified, ADD operations will be performed by default. **option** and **value** are the parameters of the corresponding operation.

- Change the name of a foreign server.

```
ALTER SERVER server_name
RENAME TO new_name;
```

## Parameters

- server\_name**

Specifies the name of the server to be modified.

- new\_version**

Specifies the latest version of the server.

- OPTIONS**

Changes options of the server. **ADD**, **SET**, and **DROP** are operations to be performed. If the operation is not specified explicitly, ADD operations will be performed. The option name must be unique, and the name and value are also validated with the foreign data wrapper library of the server.

In addition to the connection parameters supported by libpq, the following parameters are provided:

- **fdw\_startup\_cost**

Estimates the startup time required for a foreign table scan, including the time to establish a connection, analyze the request at the remote server, and generate a plan. The default value is **100**.

- **fdw\_tuple\_cost**

Specifies the additional consumption when each tuple is scanned on a remote server. The value specifies the extra consumption of data transmission between servers. The default value is **0.01**.

- new\_name**

Specifies the new name of the server.

## Examples

```
-- Create my_server.
openGauss=# CREATE SERVER my_server FOREIGN DATA WRAPPER log_fdw;

-- Change the name of a foreign server.
openGauss=# ALTER SERVER my_server
RENAME TO my_server_1;

-- Drop my_server_1.
openGauss=# DROP SERVER my_server_1;
```

## Helpful Links

[CREATE SERVER](#) and [DROP SERVER](#)

### 7.13.23 ALTER SESSION

#### Description

ALTER SESSION defines or modifies the conditions or parameters that affect the current session. Modified session parameters are kept until the current session is disconnected.

#### Precautions

- If the **START TRANSACTION** statement is not executed before the **SET TRANSACTION** statement, the transaction is ended instantly and the statement does not take effect.
- You can use the **transaction\_mode(s)** method declared in the **START TRANSACTION** statement to avoid using the **SET TRANSACTION** statement.

#### Syntax

- Set transaction parameters of a session.  

```
ALTER SESSION SET [SESSION CHARACTERISTICS AS] TRANSACTION
 { ISOLATION LEVEL { READ COMMITTED } | { READ ONLY | READ WRITE } } [, ...] ;
```
- Set other GUC parameters of a session.  

```
ALTER SESSION SET
 {{config_parameter { { TO | = } { value | DEFAULT }
 | FROM CURRENT }}
 | TIME ZONE time_zone

 | SCHEMA 'schema'
 | NAMES encoding_name
 | ROLE role_name PASSWORD 'password'
 | SESSION AUTHORIZATION { role_name PASSWORD 'password' | DEFAULT }
 | XML OPTION { DOCUMENT | CONTENT }
 } ;
```

#### Parameters

- **SESSION**  
Specifies that the specified parameters take effect for the current session. If neither **SESSION** nor **LOCAL** appears, **SESSION** is the default value.  
If this command is executed in a transaction, the effect of the command disappears after the transaction is rolled back. Once the surrounding transaction is committed, the effect will persist until the end of the session, unless the parameters are reset by another SET statement.
- **TIME ZONE time\_zone**  
Specifies the local time zone for the current session.  
Value range: a valid local time zone. The corresponding GUC parameter is **TimeZone**. The default value is **PRC**.
- **CURRENT\_SCHEMA schema**  
Specifies the current schema.

Value range: an existing schema name. If the schema name does not exist, the value of **CURRENT\_SCHEMA** will be empty.

- **SCHEMA schema**

Specifies the current schema. Here the schema is a string.

Example: set schema 'public';

- **NAMES encoding\_name**

Specifies the client character encoding. This statement is equivalent to **set client\_encoding to encoding\_name**.

Value range: a valid character encoding name. The GUC parameter corresponding to this option is **client\_encoding**. The default encoding is **UTF8**.

- **XML OPTION option**

Specifies the XML parsing mode.

Value range: **CONTENT** (default) and **DOCUMENT**.

- **config\_parameter**

Specifies the name of a configurable GUC parameter. You can use **SHOW ALL** to view the available GUC parameters.

## Examples

```
-- Create the ds schema.
openGauss=# CREATE SCHEMA ds;

-- Set the search path of the schema.
openGauss=# SET SEARCH_PATH TO ds, public;

-- Set the time/date type to the traditional Postgres format (date before month).
openGauss=# SET DATESTYLE TO postgres, dmy;

-- Set the character code of the current session to UTF8.
openGauss=# ALTER SESSION SET NAMES 'UTF8';

-- Set the time zone to Berkeley of California.
openGauss=# SET TIME ZONE 'PST8PDT';

-- Set the time zone to Italy.
openGauss=# SET TIME ZONE 'Europe/Rome';

-- Set the current schema.
openGauss=# ALTER SESSION SET CURRENT_SCHEMA TO tpcds;

-- Set XML OPTION to DOCUMENT.
openGauss=# ALTER SESSION SET XML OPTION DOCUMENT;

-- Create the role joe, and set the session role to joe.
openGauss=# CREATE ROLE joe WITH PASSWORD '*****!';
openGauss=# ALTER SESSION SET SESSION AUTHORIZATION joe PASSWORD '*****!';

-- Switch to the default user.
openGauss=> ALTER SESSION SET SESSION AUTHORIZATION default;

-- Drop the ds schema.
openGauss=# DROP SCHEMA ds;

-- Drop the role joe.
openGauss=# DROP ROLE joe;

-- Start a transaction and set the transaction level.
openGauss=# START TRANSACTION;
```

```
openGauss=# ALTER SESSION SET TRANSACTION READ ONLY;
openGauss=# END;
```

## Helpful Links

[SET](#)

## 7.13.24 ALTER SYNONYM

### Description

Alters the attributes of the SYNONYM object.

### Precautions

- Currently, only the owner of the SYNONYM object can be changed.
- Only system administrators have the permission to modify the owner of the SYNONYM object.
- The new owner must have the CREATE permission on the schema where the SYNONYM object resides.

### Syntax

```
ALTER SYNONYM synonym_name
OWNER TO new_owner;
```

### Parameters

- **synonym\_name**  
Specifies the name of the synonym to be modified, which can contain the schema name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **new\_owner**  
Specifies the new owner of the SYNONYM object.  
Value range: a string. It must be a valid username.

### Examples

```
-- Create synonym t1.
openGauss=# CREATE OR REPLACE SYNONYM t1 FOR ot.t1;

-- Create user u1.
openGauss=# CREATE USER u1 PASSWORD '*****';

-- Change the owner of synonym t1 to u1.
openGauss=# ALTER SYNONYM t1 OWNER TO u1;

-- Drop synonym t1.
openGauss=# DROP SYNONYM t1;

-- Drop user u1.
openGauss=# DROP USER u1;
```

## Helpful Links

[CREATE SYNONYM](#) and [DROP SYNONYM](#)

## 7.13.25 ALTER SYSTEM KILL SESSION

### Description

**ALTER SYSTEM KILL SESSION** ends a session.

### Precautions

None

### Syntax

```
ALTER SYSTEM KILL SESSION 'session_sid, serial' [IMMEDIATE];
```

### Parameters

- **session\_sid, serial**  
Specifies the SID and SERIAL of a session (To obtain the values, see the example.)  
Value range: SIDs and SERIALs of all sessions that can be queried from the system catalog dv\_sessions.
- **IMMEDIATE**  
Specifies that a session will be ended instantly after the statement is executed.

### Examples

```
-- Query session information.
openGauss=# SELECT sa.sessionid AS sid,0::integer AS serial#,ad.rolname AS username FROM
pg_stat_get_activity(NULL) AS sa LEFT JOIN pg_authid ad ON(sa.usesysid = ad.oid)WHERE
sa.application_name <> 'JobScheduler';
 sid | serial# | username
-----+-----+-----
140131075880720 | 0 | omm
140131025549072 | 0 | omm
140131073779472 | 0 | omm
140131071678224 | 0 | omm
140131125774096 | 0 |
140131127875344 | 0 |
140131113629456 | 0 |
140131094742800 | 0 |
(8 rows)

-- End the session whose SID is 140131075880720.
openGauss=# ALTER SYSTEM KILL SESSION '140131075880720,0' IMMEDIATE;
```

## 7.13.26 ALTER TABLE

### Description

Alters tables, including modifying table definitions, renaming tables, renaming specified columns in tables, renaming table constraints, setting table schemas, enabling or disabling row-level security policies, and adding or updating multiple columns.

## Precautions

- The owner of a table, users granted with the ALTER permission on the table, or users granted with the ALTER ANY TABLE permission can run the **ALTER TABLE** command. System administrators have the permission to run the command by default. To modify the owner or schema of a table, you must be the table owner or a system administrator and a member of the new owner role.
- The tablespace of a partitioned table cannot be modified, but the tablespace of a partition can be modified.
- The storage parameter **ORIENTATION** cannot be modified.
- Currently, **SET SCHEMA** can only set schemas to user schemas. It cannot set a schema to a system internal schema.
- Auto-increment columns cannot be added, or a column whose **DEFAULT** value contains the nextval() expression cannot be added.
- Row-level security cannot be enabled for foreign tables and temporary tables.
- When you delete a PRIMARY KEY constraint by constraint name, the NOT NULL constraint is not deleted. If necessary, manually delete the NOT NULL constraint.
- When JDBC is used, the **DEFAULT** value can be set through **PreparedStatement**.
- If you add a column using ADD COLUMN, all existing rows in the table are initialized to the column's default value (**NULL** if no **DEFAULT** value is specified).

If no **DEFAULT** value is specified for the new column, **NULL** is used, and no full table update is triggered.

If the new column has a **DEFAULT** value, the column must meet all the following requirements. Otherwise, the entire table is updated, leading to additional overheads and affecting online services.

1. The data type is BOOL, BYTEA, SMALLINT, BIGINT, SMALLINT, INTEGER, NUMERIC, FLOAT, DOUBLE PRECISION, CHAR, VARCHAR, TEXT, TIMESTAMPTZ, TIMESTAMP, DATE, TIME, TIMETZ, or INTERVAL.

2. The length of the **DEFAULT** value of the added column cannot exceed 128 bytes.

3. The **DEFAULT** value of the added column does not contain the volatile function.

4. The **DEFAULT** value is required and cannot be **NULL**.

If you are not sure whether condition 3 is met, check whether the **provolatile** attribute of the function in the PG\_RPOC system catalog is 'v'.

## Syntax

- Modify the definition of a table.  

```
ALTER TABLE [IF EXISTS] { table_name [*] | ONLY table_name | ONLY (table_name) }
action [, ...];
```

  - There are several clauses of **column\_clause**:  

```
ADD [COLUMN] column_name data_type [COLLATE collation] [column_constraint [...]]
| MODIFY column_name data_type
| MODIFY column_name [CONSTRAINT constraint_name] NOT NULL [ENABLE]
| MODIFY column_name [CONSTRAINT constraint_name] NULL
```

```
| DROP [COLUMN] [IF EXISTS] column_name [RESTRICT | CASCADE]
| ALTER [COLUMN] column_name [SET DATA] TYPE data_type [COLLATE collation] [USING
expression]
| ALTER [COLUMN] column_name { SET DEFAULT expression | DROP DEFAULT }
| ALTER [COLUMN] column_name { SET | DROP } NOT NULL
| ALTER [COLUMN] column_name SET STATISTICS [PERCENT] integer
| ADD STATISTICS ((column_1_name, column_2_name [, ...]))
| DELETE STATISTICS ((column_1_name, column_2_name [, ...]))
| ALTER [COLUMN] column_name SET ({attribute_option = value} [, ...])
| ALTER [COLUMN] column_name RESET (attribute_option [, ...])
| ALTER [COLUMN] column_name SET STORAGE { PLAIN | EXTERNAL | EXTENDED | MAIN }
```

 NOTE

- **ADD [ COLUMN ] column\_name data\_type [ COLLATE collation ] [ column\_constraint [ ... ] ]**

Adds a column to a table. If you add a column using ADD COLUMN, all existing rows in the table are initialized to the column's default value (**NULL** if no **DEFAULT** value is specified).

- **ADD ( { column\_name data\_type } [, ...] )**

Adds columns in the table.

- **MODIFY ( { column\_name data\_type | column\_name [ CONSTRAINT constraint\_name ] NOT NULL [ ENABLE ] | column\_name [ CONSTRAINT constraint\_name ] NULL } [, ...] )**

Modifies the data type of an existing column in the table. Running this command will clear the statistics of this column. You are advised to collect the statistics of this column again after the modification.

- **DROP [ COLUMN ] [ IF EXISTS ] column\_name [ RESTRICT | CASCADE ]**

Drops a column from a table. Indexes and constraints related to the column are automatically dropped. If an object not belonging to the table depends on the column, **CASCADE** must be specified, such as a view.

The DROP COLUMN statement does not physically remove the column, but simply makes it invisible to SQL operations. Subsequent INSERT and UPDATE operations on the table will store a **NULL** value for the column. Therefore, dropping a column takes a short period of time but does not immediately release the tablespace on the disks, because the space occupied by the dropped column is not recycled. The space will be recycled when **VACUUM** is executed.

- **ALTER [ COLUMN ] column\_name [ SET DATA ] TYPE data\_type [ COLLATE collation ] [ USING expression ]**

Alters the type of a column in a table. Indexes and simple table constraints on the column will automatically use the new data type by reparsing the originally supplied expression.

If the original data type of a column and the modified data type are binary compatible, you do not need to rewrite the entire table when running this statement. In other scenarios, the entire table is rewritten. You can check whether the original type and target type are binary compatible in the PG\_CAST system catalog. If **castmethod** is 'b', they are binary compatible. For example, if the data type of the source table is text and is converted to int, table rewriting is triggered. If it is converted to clob, table rewriting is not triggered. If table rewriting is triggered, the deleted space on the table is recycled immediately.

Running this command will clear the statistics of this column. You are advised to collect the statistics of this column again after the modification.

- **ALTER [ COLUMN ] column\_name { SET DEFAULT expression | DROP DEFAULT }**

Sets or drops the default value for a column. The default values only apply to subsequent INSERT operations; they do not cause rows already in the table to change. Defaults can also be created for views, in which case they are inserted into INSERT statements on the view before the view's ON INSERT rule is applied.

- **ALTER [ COLUMN ] column\_name { SET | DROP } NOT NULL**

Alters whether a column is marked to allow **NULL** values or to reject **NULL** values. You can only use SET NOT NULL when the column contains non-**NULL** values.

- **ALTER [ COLUMN ] column\_name SET STATISTICS [PERCENT] integer**



Specifies the per-column statistics-collection target for subsequent ANALYZE operations. The target can be set in the range from 0 to 10000. Set it to -1 to revert to using the default system statistics target.

- **ALTER [ COLUMN ] column\_name SET ( {attribute\_option = value} [, ... ] )**  
**ALTER [ COLUMN ] column\_name RESET ( attribute\_option [, ... ] )**

Sets or resets per-attribute options.

Currently, the only defined per-attribute options are **n\_distinct** and **n\_distinct\_inherited**. **n\_distinct** affects statistics of a table, while **n\_distinct\_inherited** affects the statistics of the table and its subtables. Currently, only **SET/RESET n\_distinct** is supported, and **SET/RESET n\_distinct\_inherited** is forbidden.

- **ALTER [ COLUMN ] column\_name SET STORAGE { PLAIN | EXTERNAL | EXTENDED | MAIN }**

Sets the storage mode for a column. This clause specifies whether this column is held inline or in a secondary TOAST table, and whether the data should be compressed. Only row-store tables can be set. SET STORAGE itself does not change anything in the table. It sets the strategy to be pursued during future table updates.

**column\_constraint** is as follows:

```
[CONSTRAINT constraint_name]
{ NOT NULL |
 NULL |
 CHECK (expression) |
 DEFAULT default_expr |
 GENERATED ALWAYS AS (generation_expr) STORED |
 UNIQUE index_parameters |
 PRIMARY KEY index_parameters |
 REFERENCES reftable [(refcolumn)] [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
 [ON DELETE action] [ON UPDATE action] } [DEFERRABLE | NOT DEFERRABLE |
INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- **table\_constraint\_using\_index** used to add the primary key constraint or unique constraint based on the unique index is as follows:

```
[CONSTRAINT constraint_name]
{ UNIQUE | PRIMARY KEY } USING INDEX index_name
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- **table\_constraint** is as follows:

```
[CONSTRAINT constraint_name]
{ CHECK (expression) |
 UNIQUE (column_name [, ...]) index_parameters |
 PRIMARY KEY (column_name [, ...]) index_parameters |
 PARTIAL CLUSTER KEY (column_name [, ...]) }
FOREIGN KEY (column_name [, ...]) REFERENCES reftable [(refcolumn [, ...])]
[MATCH FULL | MATCH PARTIAL | MATCH SIMPLE] [ON DELETE action] [ON UPDATE
action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

**index\_parameters** is as follows:

```
[WITH ({storage_parameter = value} [, ...])]
[USING INDEX TABLESPACE tablespace_name]
```

- Rename a table. The renaming does not affect stored data.

```
ALTER TABLE [IF EXISTS] table_name
 RENAME TO new_table_name;
```

- Rename the specified column in the table.

```
ALTER TABLE [IF EXISTS] { table_name [*] | ONLY table_name | ONLY (table_name) }
 RENAME [COLUMN] column_name TO new_column_name;
```

- Rename the constraint of the table.

```
ALTER TABLE [IF EXISTS] { table_name [*] | ONLY table_name | ONLY (table_name) }
 RENAME CONSTRAINT constraint_name TO new_constraint_name;
```

- Set the schema of the table.

```
ALTER TABLE [IF EXISTS] table_name
SET SCHEMA new_schema;
```

#### NOTE

- The schema setting moves the table into another schema. Associated indexes and constraints owned by table columns are migrated as well. Currently, the schema for sequences cannot be changed. If the table has sequences, delete the sequences, and create them again or delete the ownership between the table and sequences. In this way, the table schema can be changed.
- To change the schema of a table, you must also have the CREATE permission on the new schema. To add the table as a new child of a parent table, you must own the parent table as well. To alter the owner, you must also be a direct or indirect member of the new owning role, and that role must have the CREATE permission on the table's schema. These restrictions enforce that the user can only rebuild and delete the table. However, a system administrator can alter the ownership of any table anyway.
- All the actions except for RENAME and SET SCHEMA can be combined into a list of multiple alterations to apply in parallel. For example, it is possible to add several columns or alter the type of several columns in a single statement. This is useful with large tables, since only one pass over the tables need be made.
- Adding a CHECK or NOT NULL constraint will scan the table to validate that existing rows meet the constraint.
- Adding a column with a non-null default or changing the type of an existing column will rewrite the entire table. Rewriting a large table may take much time and temporarily needs doubled disk space.
- Add columns.  

```
ALTER TABLE [IF EXISTS] table_name
ADD ({ column_name data_type [COLLATE collation] [column_constraint [...]] } [, ...]);
```
- Update columns.  

```
ALTER TABLE [IF EXISTS] table_name
MODIFY ({ column_name data_type | column_name [CONSTRAINT constraint_name] NOT NULL
[ENABLE] | column_name [CONSTRAINT constraint_name] NULL } [, ...]);
```

## Parameters

- **IF EXISTS**  
Sends a notice instead of an error if no tables have identical names. The notice prompts that the table you are querying does not exist.
- **table\_name [\*] | ONLY table\_name | ONLY ( table\_name )**  
**table\_name** is the name of the table that you need to modify.  
If **ONLY** is specified, only the table is modified. If **ONLY** is not specified, the table and all subtables are modified. You can add the asterisk (\*) option following the table name to specify that all subtables are scanned, which is the default operation.
- **constraint\_name**  
Specifies the name of an existing constraint to drop.
- **index\_name**  
Specifies the name of an index.
- **storage\_parameter**  
Specifies the name of a storage parameter.  
The following option is added for creating an index:
  - **parallel\_workers** (int type)

Number of bgworker threads started when an index is created. For example, value **2** indicates that two bgworker threads are started to create indexes concurrently.

Value range: [0,32]. The value **0** indicates that this function is disabled.

Default value: If this parameter is not set, the concurrent index creation function is disabled.

– **hasuids** (Boolean type)

Default value: **off**

If this parameter is set to **on**, a unique table-level ID is allocated to a tuple when the tuple is updated.

The following option is added to fix optimizer statistics:

– **min\_tuples** (float8 type)

Default value: **0**

The optimizer selects the larger value of the estimated statistics and the parameter to calculate the data volume based on the estimation table of statistics.

- **new\_owner**

Specifies the name of the new table owner.

- **new\_tablespace**

Specifies the new name of the tablespace to which the table belongs.

- **column\_name, column\_1\_name, column\_2\_name**

Specifies the name of a new or existing column.

- **data\_type**

Specifies the type of a new column or a new type of an existing column.

- **collation**

Specifies the collation rule name of a column. The optional COLLATE clause specifies a collation for the new column; if omitted, the collation is the default for the new column. You can run the **select \* from pg\_collation;** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.

- **USING expression**

Specifies how to compute the new column value from the old; if omitted, the default conversion is an assignment cast from old data type to new. A USING clause must be provided if there is no implicit or assignment cast from the old to new type.

 **NOTE**

**USING** in ALTER TYPE can specify any expression involving the old values of the row; that is, it can refer to any columns other than the one being cast. This allows general casting to be done with the ALTER TYPE syntax. Because of this flexibility, the USING expression is not applied to the column's default value (if any); the result might not be a constant expression as required for a default. This means that when there is no implicit or assignment cast from old to new type, ALTER TYPE might fail to convert the default even though a USING clause is supplied. In such cases, drop the default with DROP DEFAULT, perform ALTER TYPE, and then use SET DEFAULT to add a suitable new default. Similar considerations apply to indexes and constraints involving the column.

- **NOT NULL | NULL**  
Sets whether the column allows null values.
- **ENABLE**  
Specifies that the constraint is enabled. By default, the constraint is enabled.
- **integer**  
Specifies the constant value of a signed integer. When using PERCENT, the range of **integer** is from 0 to 100.
- **attribute\_option**  
Specifies an attribute option.
- **PLAIN | EXTERNAL | EXTENDED | MAIN**  
Specifies a column-store mode.
  - **PLAIN** must be used for fixed-length values (such as integers). It must be inline and uncompressed.
  - **MAIN** is for inline, compressible data.
  - **EXTERNAL** is for external, uncompressed data. Use of **EXTERNAL** will make substring operations on **text** and **bytea** values run faster, at the penalty of increased storage space.
  - **EXTENDED** is for external, compressed data. **EXTENDED** is the default for most data types that support non-**PLAIN** storage.
- **CHECK ( expression )**  
New rows or rows to be updated must satisfy for an expression to be true. If any row produces a false result, an error is raised and the database is not modified.  
  
A CHECK constraint specified as a column constraint should reference only the column's value, while an expression in a table constraint can reference multiple columns.  
  
Currently, **CHECK ( expression )** does not include subqueries and cannot use variables apart from the current column.
- **DEFAULT default\_expr**  
Assigns a default data value to a column.  
  
The data type of the default expression must match that of the column.  
  
The default expression will be used in any INSERT operation that does not specify a value for the column. If there is no default value for a column, then the default value is **NULL**.
- **GENERATED ALWAYS AS ( generation\_expr ) STORED**  
This clause creates a column as a generated column. The value of the generated column is calculated by **generation\_expr** when data is written (inserted or updated). **STORED** indicates that the value of the generated column is stored as an ordinary column.

 **NOTE**

- The generation expression cannot refer to data other than the current row in any way. The generation expression cannot reference other generation columns or system columns. The generation expression cannot return a result set. No subquery, aggregate function, or window function can be used. The function called by the generation expression can only be an immutable function.
- Default values cannot be specified for generated columns.
- The generated column cannot be used as a part of the partition key.
- Do not specify the generated column and the CASCADE, SET NULL, and SET DEFAULT actions of the ON UPDATE constraint at the same time. Do not specify the generated column and the SET NULL, and SET DEFAULT actions of the ON DELETE constraint at the same time.
- The method of modifying and deleting generated columns is the same as that of ordinary columns. Delete the ordinary column that the generated column depends on. The generated column is automatically deleted. The type of the column on which the generated column depends cannot be changed.
- The generated column cannot be directly written. In the INSERT or UPDATE statement, values cannot be specified for generated columns, but the keyword DEFAULT can be specified.
- The permission control for generated columns is the same as that for ordinary columns.
- Columns cannot be generated for column-store tables and memory-optimized tables (MOTs). In foreign tables, only **postgres\_fdw** supports generated columns.

- **UNIQUE index\_parameters**

**UNIQUE ( column\_name [, ... ] ) index\_parameters**

Specifies that a group of one or more columns of a table can contain only unique values.

- **PRIMARY KEY index\_parameters**

**PRIMARY KEY ( column\_name [, ... ] ) index\_parameters**

Specifies that a column or columns of a table can contain only unique (non-duplicate) and non-NULL values.

- **REFERENCES reftable [ ( refcolumn ) ] [ MATCH matchtype ] [ ON DELETE action ] [ ON UPDATE action ] (column constraint)**

**FOREIGN KEY ( column\_name [, ... ] ) REFERENCES reftable [ ( refcolumn [, ... ] ) ] [ MATCH matchtype ] [ ON DELETE action ] [ ON UPDATE action ] (table constraint)**

The foreign key constraint requires that the group consisting of one or more columns in the new table should contain and match only the referenced column values in the referenced table. If **refcolumn** is omitted, the primary key of **reftable** is used. The referenced column should be the only column or primary key in the referenced table. A foreign key constraint cannot be defined between a temporary table and a permanent table.

There are three types of matching between a reference column and a referenced column:

- **MATCH FULL:** A column with multiple foreign keys cannot be **NULL** unless all foreign key columns are **NULL**.
- **MATCH SIMPLE** (default): Any unexpected foreign key column can be **NULL**.
- **MATCH PARTIAL:** This option is not supported currently.

In addition, when certain operations are performed on the data in the referenced table, the operations are performed on the corresponding columns in the new table. The ON DELETE clause specifies the operations to be executed after a referenced row in the referenced table is deleted. The ON UPDATE clause specifies the operation to be performed when the referenced column data in the referenced table is updated. Possible responses to the ON DELETE and ON UPDATE clauses are as follows:

- **NO ACTION** (default): When a foreign key is deleted or updated, an error indicating that the foreign key constraint is violated is reported. If the constraint is deferrable and there are still any referenced rows, this error will occur when the constraint is checked.
- **RESTRICT**: When a foreign key is deleted or updated, an error indicating that the foreign key constraint is violated is created. It is the same as **NO ACTION** except that the constraint is not deferrable.
- **CASCADE**: deletes any row that references the deleted row from the new table, or update the column value of the referenced row in the new table to the new value of the referenced column.
- **SET NULL**: sets the referenced columns to **NULL**.
- **SET DEFAULT**: sets the referenced columns to their default values.
- **DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE**

Sets whether the constraint can be deferrable.

- **DEFERRABLE**: deferrable to the end of the transaction and checked using SET CONSTRAINTS.
- **NOT DEFERRABLE**: checks immediately after the execution of each command.
- **INITIALLY IMMEDIATE**: checks immediately after the execution of each statement.
- **INITIALLY DEFERRED**: checks when the transaction ends.

#### NOTE

Ustore tables do not support the **DEFERRABLE** and **INITIALLY DEFERRED** constraints.

- **WITH ( {storage\_parameter = value} [, ... ] )**  
Specifies an optional storage parameter for a table or an index.
- **tablespace\_name**  
Specifies the name of the tablespace where the index locates.
- **new\_table\_name**  
Specifies the new table name.
- **new\_column\_name**  
Specifies the new name of a specific column in a table.
- **new\_constraint\_name**  
Specifies the new name of a table constraint.
- **new\_schema**  
Specifies the new schema name.

- **CASCADE**  
Automatically drops objects that depend on the dropped column or constraint (for example, views referencing the column).
- **RESTRICT**  
Refuses to drop the column or constraint if there are any dependent objects. This is the default behavior.
- **schema\_name**  
Specifies the schema name of a table.

## Examples of Altering a Table

- Rename a table.  

```
openGauss=# CREATE TABLE aa(c1 int, c2 int);
openGauss=# ALTER TABLE IF EXISTS aa RENAME TO test_alt1;
```
- Alter the schema of a table.  

```
-- Create test_schema.
openGauss=# CREATE SCHEMA test_schema;
-- Change the schema of the test_alt1 table to test_schema.
openGauss=# ALTER TABLE test_alt1 SET SCHEMA test_schema;
-- Query table information.
openGauss=# SELECT schemaname,tablename FROM pg_tables WHERE tablename = 'test_alt1';
schemaname | tablename
-----+-----
test_schema | test_alt1
(1 row)
```
- Change the owner of a table.  

```
-- Create test_user.
openGauss=# CREATE USER test_user PASSWORD 'XXXXXXXXXX';
-- Change the owner of the test_alt1 table to test_user.
openGauss=# ALTER TABLE IF EXISTS test_schema.test_alt1 OWNER TO test_user;
-- Query.
openGauss=# SELECT tablename, schemaname, tableowner FROM pg_tables WHERE tablename = 'test_alt1';
tablename | schemaname | tableowner
-----+-----+-----
test_alt1 | test_schema | test_user
(1 row)
```
- Alter the tablespace of a table.  

```
-- Create the tbs_data1 tablespace.
openGauss=# CREATE TABLESPACE tbs_data1 RELATIVE LOCATION 'tablespace1/tbs_data1';
-- Change the tablespace of the test_alt1 table to tbs_data1.
openGauss=# ALTER TABLE test_schema.test_alt1 SET TABLESPACE tbs_data1;
-- Query.
openGauss=# SELECT tablename, tablespace FROM pg_tables WHERE tablename = 'test_alt1';
tablename | tablespace
-----+-----
test_alt1 | tbs_data1
(1 row)

-- Drop.
openGauss=# DROP TABLE test_schema.test_alt1;
openGauss=# DROP TABLESPACE tbs_data1;
openGauss=# DROP SCHEMA test_schema;
openGauss=# DROP USER test_user;
```

## Examples of Changing a Column

- Change column names.  

```
-- Create a table.
openGauss=# CREATE TABLE test_alt2(c1 INT,c2 INT);
-- Change column names.
```

```
openGauss=# ALTER TABLE test_alt2 RENAME c1 TO id;
openGauss=# ALTER TABLE test_alt2 RENAME COLUMN c2 to areaid;
-- Query.
openGauss=# \d test_alt1
 Table "public.test_alt1"
 Column | Type | Modifiers
-----+-----+-----
 id | integer |
 areaid | integer |
```

- Add a column.

```
-- Add a column to the test_alt1 table.
openGauss=# ALTER TABLE IF EXISTS test_alt2 ADD COLUMN name VARCHAR(20);
-- Query.
openGauss=# \d test_alt2
 Table "public.test_alt1"
 Column | Type | Modifiers
-----+-----+-----
 id | integer |
 areacode | integer |
 name | character varying(20) |
```

- Alter the data type of a column.

```
-- Change the type of the name column in the test_alt1 table.
openGauss=# ALTER TABLE test_alt1 MODIFY name VARCHAR(50);
-- Query.
openGauss=# \d test_alt1
 Table "public.test_alt2"
 Column | Type | Modifiers
-----+-----+-----
 id | integer |
 areaid | integer |
 name | character varying(50) |
-- Change the type of the name column in the test_alt1 table.
openGauss=# ALTER TABLE test_alt2 ALTER COLUMN name TYPE VARCHAR(25);
-- Query.
openGauss=# \d test_alt2
 Table "public.test_alt2"
 Column | Type | Modifiers
-----+-----+-----
 id | integer |
 areaid | integer |
 name | character varying(25) |
```

- Drop a column.

```
-- Drop the areaid column from test_alt1.
openGauss=# ALTER TABLE test_alt2 DROP COLUMN areaid;
-- Query.
openGauss=# \d test_alt2
 Table "public.test_alt2"
 Column | Type | Modifiers
-----+-----+-----
 id | integer |
 name | character varying(25) |
```

- Modify the column-store mode.

```
-- View table details.
openGauss=# \d+ test_alt2
 Table "public.test_alt2"
 Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
 id | integer | | plain | |
 name | character varying(25) | | extended | |
Has OIDs: no
Options: orientation=row, storage_type=USTORE

-- Change the storage mode of the name column in the test_alt2 table.
openGauss=# ALTER TABLE test_alt2 ALTER COLUMN name SET STORAGE PLAIN;

-- Query.
openGauss=# \d+ test_alt2
```



```
Table "public.test_alt2"
Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
id | integer | | plain | |
name | character varying(25) | | plain | |
Has OIDs: no
Options: orientation=row, storage_type=USTORE

-- Drop.
openGauss=# DROP TABLE test_alt2;
```

## Examples of Changing a Constraint

- Add a NOT NULL constraint to a column.

```
-- Create a table.
openGauss=# CREATE TABLE test_alt3(pid INT, areaid CHAR(5), name VARCHAR(20));
-- Add a NOT NULL constraint to pid.
openGauss=# ALTER TABLE test_alt3 MODIFY pid NOT NULL;
-- Query.
openGauss=# \d test_alt3
Table "public.test_alt3"
Column | Type | Modifiers
-----+-----+-----
pid | integer | not null
areaid | character(5) |
name | character varying(20) |
```

- Cancel the NOT NULL constraint on a column.

```
openGauss=# ALTER TABLE test_alt3 MODIFY pid NULL;
-- Query.
openGauss=# \d test_alt3
Table "public.test_alt3"
Column | Type | Modifiers
-----+-----+-----
pid | integer |
areaid | character(5) |
name | character varying(20) |
```

- Alter the default value of a column.

```
-- Alter the default value of id in the test_alt1 table.
openGauss=# ALTER TABLE test_alt3 ALTER COLUMN areaid SET DEFAULT '00000';
-- Query.
openGauss=# \d test_alt3
Table "public.test_alt3"
Column | Type | Modifiers
-----+-----+-----
pid | integer |
areaid | character(5) | default '00000':bpchar
name | character varying(20) |
-- Drop the default value of id.
openGauss=# ALTER TABLE test_alt3 ALTER COLUMN areaid DROP DEFAULT;
-- Query.
openGauss=# \d test_alt3
Table "public.test_alt3"
Column | Type | Modifiers
-----+-----+-----
pid | integer |
areaid | character(5) |
name | character varying(20) |
```

- Add a table-level constraint.

- Directly add a constraint.

```
-- Add a primary key constraint to the table.
openGauss=# ALTER TABLE test_alt3 ADD CONSTRAINT pk_test3_pid PRIMARY KEY (pid);
-- Query.
openGauss=# \d test_alt3
Table "public.test_alt3"
Column | Type | Modifiers
-----+-----+-----
```

```
pid | integer | not null
areaid | integer |
name | character varying(20) |
Indexes:
 "pk_test3_pid" PRIMARY KEY, btree (pid) TABLESPACE pg_default
```

- Create an index and then add a constraint.

```
-- Create a table.
openGauss=# CREATE TABLE test_alt4(c1 INT, c2 INT);
-- Create an index.
openGauss=# CREATE UNIQUE INDEX pk_test4_c1 ON test_alt4(c1);
-- Associate the created index when adding a constraint.
openGauss=# ALTER TABLE test_alt4 ADD CONSTRAINT pk_test4_c1 PRIMARY KEY USING
INDEX pk_test4_c1;
-- Query.
openGauss=# \d test_alt4
 Table "public.test_alt4"
 Column | Type | Modifiers
-----+-----+-----
c1 | integer | not null
c2 | integer |
Indexes:
 "pk_test4_c1" PRIMARY KEY, btree (c1) TABLESPACE pg_default
-- Drop.
openGauss=# DROP TABLE test_alt4;
```

- Drop a table-level constraint.

```
-- Drop a constraint.
openGauss=# ALTER TABLE test_alt3 DROP CONSTRAINT IF EXISTS pk_test3_pid;
-- Query.
openGauss=# \d test_alt3
 Table "public.test_alt3"
 Column | Type | Modifiers
-----+-----+-----
pid | integer | not null
areaid | integer |
name | character varying(20) |
Indexes:
 "pk_test3_pid" PRIMARY KEY, btree (pid) TABLESPACE pg_default
-- Drop.
openGauss=# DROP TABLE test_alt3;
```

## Helpful Links

[CREATE TABLE](#) and [DROP TABLE](#)

## 7.13.27 ALTER TABLE PARTITION

### Description

Alters table partitions, including adding, deleting, splitting, merging, clearing, exchanging, and renaming partitions, moving partition tablespaces, and modifying partition attributes.

### Precautions

- Only the owner of a partitioned table or users granted the ALTER permission on the partitioned table can run the **ALTER TABLE PARTITION** command. System administrators have the permission to run the command by default.
- The tablespace of the added partition cannot be PG\_GLOBAL.
- The name of the added partition must be different from names of existing partitions in the partitioned table.
- The key value of the added partition must be consistent with the type of partition keys in the partitioned table.

- If a range partition is added, the key value of the added partition must be greater than the upper limit of the last range partition in the partitioned table.
- If a list partition is added, the key value of the added partition cannot be the same as that of an existing partition.
- Hash partitions cannot be added.
- If the number of partitions in the target partitioned table has reached the maximum (**1048575**), partitions cannot be added.
- If a partitioned table has only one partition, the partition cannot be deleted.
- Use PARTITION FOR() to choose partitions. The number of specified values in the brackets should be the same as the number of columns specified when you define a partition, and they must be consistent.
- The **Value** partitioned table does not support the ALTER PARTITION operation.
- Partitions cannot be added to an interval partitioned table.
- Hash partitioned tables do not support splitting, combination, addition, and deletion of partitions.
- List partitioned tables do not support partition splitting or partition combination.
- If **enable\_gpi\_auto\_update** is set to **on**, the global index is automatically updated even if the UPDATE GLOBAL INDEX clause is not declared.

## Syntax

- Modify the syntax of the table partition.  

```
ALTER TABLE [IF EXISTS] { table_name [*] | ONLY table_name | ONLY (table_name) }
 action [, ...];
```

**action** indicates the following clauses for maintaining partitions. For the partition continuity when multiple clauses are used for partition maintenance, GaussDB runs **DROP PARTITION** and then **ADD PARTITION**, and finally runs the rest clauses in sequence.

```
move_clause |
exchange_clause |
row_clause |
merge_clause |
modify_clause |
split_clause |
add_clause |
drop_clause |
truncate_clause
```

  - The move\_clause syntax is used to move the partition to a new tablespace.  

```
MOVE PARTITION { partition_name | FOR (partition_value [, ...]) } TABLESPACE tablespacename
```
  - The exchange\_clause syntax is used to move the data from an ordinary table to a specified partition.  

```
EXCHANGE PARTITION { (partition_name) | FOR (partition_value [, ...]) }
 WITH TABLE {[ONLY] ordinary_table_name | ordinary_table_name * | ONLY
 (ordinary_table_name)}
 [{ WITH | WITHOUT } VALIDATION] [VERBOSE] [UPDATE GLOBAL INDEX]
```

The ordinary table and partition whose data is to be exchanged must meet the following requirements:

    - The number of columns of the ordinary table is the same as that of the partition, and their information should be consistent, including

column name, data type, constraint, collation information, and storage parameter.

- The number and information of indexes of the ordinary table and partition should be consistent.
- The number and information of constraints of the ordinary table and partition should be consistent.
- An ordinary table cannot be a temporary table. A partitioned table can only be a range partitioned table, list partitioned table, hash partitioned table, or interval partitioned table.
- Ordinary tables and partitioned tables do not support dynamic data masking and row-level security constraints.

---

**NOTICE**

- When the exchange is done, the data and tablespace of the ordinary table and partition are exchanged. The statistics about ordinary tables and partitions become unreliable, and they should be analyzed again.
- A non-partition key cannot be used to create a local unique index. Therefore, if an ordinary table contains a unique index, data cannot be exchanged.
- If the DROP COLUMN operation is performed on an ordinary or partitioned table, the dropped column still exists physically. Therefore, you need to ensure that the dropped column of the ordinary table is strictly aligned with that of the partition.

- 
- The row\_clause syntax is used to set row movement of a partitioned table.

```
{ ENABLE | DISABLE } ROW MOVEMENT
```

- The merge\_clause syntax is used to merge partitions into one. Currently, only range partitions can be merged.

```
MERGE PARTITIONS { partition_name } [, ...] INTO PARTITION partition_name
[TABLESPACE tablespacename] [UPDATE GLOBAL INDEX]
```

---

**NOTICE**

Ustore tables do not support ALTER TABLE MERGE PARTITIONS in transaction blocks and stored procedures.

- 
- The modify\_clause syntax is used to set whether a partitioned index is usable.

```
MODIFY PARTITION partition_name { UNUSABLE LOCAL INDEXES | REBUILD UNUSABLE LOCAL INDEXES }
```

- The split\_clause syntax is used to split one partition into partitions. Currently, only range partitions can be split.

```
SPLIT PARTITION { partition_name | FOR (partition_value [, ...]) } { split_point_clause | no_split_point_clause } [UPDATE GLOBAL INDEX]
```

- The syntax of split\_point\_clause is as follows:

```
AT (partition_value) INTO (PARTITION partition_name [TABLESPACE
tablespacename] , PARTITION partition_name [TABLESPACE tablespacename])
```

#### NOTICE

- The size of the split point should be in the range of partition keys of the partition to be split. The split point can only split one partition into two new partitions.

- The syntax of **no\_split\_point\_clause** is as follows:  
INTO { ( partition\_less\_than\_item [, ...] ) | ( partition\_start\_end\_item [, ...] ) }

#### NOTICE

- The first new partition key specified by **partition\_less\_than\_item** should be greater than that of the previously split partition (if any), and the last partition key specified by **partition\_less\_than\_item** should equal that of the partition being split.
- The first new partition key specified by **partition\_start\_end\_item** should equal that of the former partition (if any), and the last partition key specified by **partition\_start\_end\_item** should equal that of the partition being split.
- **partition\_less\_than\_item** supports a maximum of 4 partition keys, while **partition\_start\_end\_item** supports only one partition key. For details about the supported data types, see • [PARTITION BY RANGE\(partition\\_key\)](#).
- **partition\_less\_than\_item** and **partition\_start\_end\_item** cannot be used in the same statement.

- The syntax of **partition\_less\_than\_item** is as follows:  
PARTITION partition\_name VALUES LESS THAN ( { partition\_value | MAXVALUE } [, ...] )  
[ TABLESPACE tablespacename ]

- The **partition\_start\_end\_item** syntax is as follows. For details about the constraints, see [START END](#).

```
PARTITION partition_name {
 {START(partition_value) END (partition_value) EVERY (interval_value)} |
 {START(partition_value) END ({partition_value | MAXVALUE})} |
 {START(partition_value)} |
 {END({partition_value | MAXVALUE})}
} [TABLESPACE tablespacename]
```

- The **add\_clause** syntax is used to add one or more partitions to a specified partitioned table.

```
ADD PARTITION (partition_col1_name = partition_col1_value [, partition_col2_name =
partition_col2_value] [, ...])
[LOCATION 'location1']
[PARTITION (partition_colA_name = partition_colA_value [, partition_colB_name =
partition_colB_value] [, ...])]
[LOCATION 'location2']
ADD {partition_less_than_item | partition_start_end_item| partition_list_item }
```

The **partition\_list\_item** syntax is as follows:

```
PARTITION partition_name VALUES (list_values_clause)
[TABLESPACE tablespacename]
```

#### NOTICE

- **partition\_list\_item** supports only one partition key. For details about the data types supported by **partition\_list\_item**, see • [PARTITION BY LIST\(partition\\_key\)](#).
- Interval and hash partitioned tables do not support partition addition.

- The drop\_clause syntax is used to remove a partition from a specified partitioned table.

```
DROP PARTITION { partition_name | FOR (partition_value [, ...]) } [UPDATE GLOBAL INDEX]
```

#### NOTICE

- Hash partitioned table does not support partition deletion.
- If a partitioned table has only one partition, the partition cannot be deleted.

- The truncate\_clause syntax is used to remove a specified partition from a partitioned table.

```
TRUNCATE PARTITION { partition_name | FOR (partition_value [, ...]) } [UPDATE GLOBAL INDEX]
```

- Alter the name of a partition.

```
ALTER TABLE [IF EXISTS] { table_name [*] | ONLY table_name | ONLY (table_name) }
```

```
RENAME PARTITION { partition_name | FOR (partition_value [, ...]) } TO partition_new_name;
```

## Parameters

- **table\_name**  
Specifies the name of a partitioned table.  
Value range: an existing partitioned table name.
- **partition\_name**  
Specifies the name of a partition.  
Value range: an existing partition name.
- **tablespacename**  
Specifies which tablespace the partition moves to.  
Value range: an existing tablespace name.
- **partition\_value**  
Specifies the key value of a partition.  
The value specified by **PARTITION FOR ( partition\_value [, ...] )** can uniquely identify a partition.  
Value range: partition keys for the partition to be renamed.
- **UNUSABLE LOCAL INDEXES**  
Sets all the indexes unusable in the partition.
- **REBUILD UNUSABLE LOCAL INDEXES**  
Rebuilds all the indexes in the partition.
- **{ ENABLE | DISABLE } ROW MOVEMENT**  
Specifies whether to enable row movement.

If the tuple value is updated on the partition key during the UPDATE operation, the partition where the tuple is located is altered. Setting this parameter enables error messages to be reported or movement of the tuple between partitions.

Value range:

- **ENABLE**: Row movement is enabled.
- **DISABLE**: Row movement is disabled.

It is enabled by default.

- **ordinary\_table\_name**

Specifies the name of the ordinary table whose data is to be migrated.

Value range: an existing ordinary table name.

- **{ WITH | WITHOUT } VALIDATION**

Checks whether the ordinary table data meets the specified partition key range of the partition to be migrated.

Value range:

- **WITH**: checks whether the ordinary table data meets the partition key range of the partition to be migrated. If any data does not meet the required range, an error is reported.
- **WITHOUT**: does not check whether the ordinary table data meets the partition key range of the partition to be migrated.

The default value is **WITH**.

The check is time consuming, especially when the data volume is large. Therefore, use **WITHOUT** when you are sure that the current ordinary table data meets the partition key range of the partition to be migrated.

- **VERBOSE**

When **VALIDATION** is **WITH**, if the ordinary table contains data that is out of the partition key range, insert the data into the correct partition. If there is no correct partition where the data can be inserted, an error is reported.

---

**NOTICE**

Only when **VALIDATION** is **WITH**, **VERBOSE** can be specified.

---

- **partition\_new\_name**

Specifies the new name of a partition.

Value range: a string that complies with the [Identifier Naming Conventions](#).

## Examples

See [Examples](#) in "CREATE TABLE PARTITION."

## Helpful Links

[CREATE TABLE PARTITION](#) and [DROP TABLE](#)

## 7.13.28 ALTER TABLE SUBPARTITION

### Description

Alters partitions from a level-2 partitioned table, including adding, deleting, clearing, and splitting partitions.

### Precautions

- Currently, partitions from the level-2 partitioned table can be added, deleted, cleared, or split only.
- The tablespace of the added partition cannot be PG\_GLOBAL.
- The name of the added partition must be different from the names of the existing level-1 and level-2 partitions in the partitioned table.
- The key value of the added partition must be consistent with the type of partition keys in the partitioned table.
- If a range partition is added, the key value of the added partition must be greater than the upper limit of the last range partition in the partitioned table. To add a partition to a table with the **MAXVALUE** partition, you are advised to use the SPLIT syntax.
- If a list partition is added, the key value of the added partition cannot be the same as that of an existing partition. To add a partition to a table with the **DEFAULT** partition, you are advised to use the SPLIT syntax.
- Hash partitions cannot be added. However, if the level-2 partition mode of an level-2 partitioned table is hash but the level-1 partition mode is not hash, you can add a level-1 partition and create the corresponding level-2 partition.
- If the number of partitions in the target partitioned table has reached the maximum (**1048575**), partitions cannot be added.
- If the partitioned table contains only one level-1 or level-2 partition, the partition cannot be deleted.
- Hash partitions cannot be deleted.
- Use PARTITION FOR() to choose partitions. The number of specified values in the brackets should be the same as the number of columns specified when you define a partition, and they must be consistent.
- Only level-2 partitions (leaf nodes) can be split. Only range and list partitioning policies can be used and hash partitioning policies are not supported. The list partitioning policy can be used only when the default partition is used.
- Only the owner of a partitioned table or users granted the ALTER permission on the partitioned table can run the **ALTER TABLE PARTITION** command. System administrators have the permission to run the command by default.
- If the ALTER statement does not contain UPDATE GLOBAL INDEX, the original GLOBAL index is invalid. In this case, other indexes are used for query. If the ALTER statement contains UPDATEGLOBAL INDEX, the original GLOBAL index is still valid and the index function is correct.
- If **enable\_gpi\_auto\_update** is set to **on**, the global index is automatically updated even if the UPDATE GLOBAL INDEX clause is not declared.



## Syntax

- Modify the syntax of the table partition.  
`ALTER TABLE [ IF EXISTS ] { table_name [*] | ONLY table_name | ONLY ( table_name ) }  
action [, ... ];`

**action** indicates the following clauses for maintaining partitions.

```
add_clause |
drop_clause |
split_clause |
truncate_clause
```

- The `add_clause` syntax is used to add one or more partitions to a specified partitioned table. The syntax can be used in level-1 partitions.  
`ADD {partition_less_than_item | partition_list_item } ( ( subpartition_definition_list ) )`

It can also be used in level-2 partitions.

```
MODIFY PARTITION partition_name ADD subpartition_definition
```

**partition\_less\_than\_item** defines a range partition. The syntax is as follows:

```
PARTITION partition_name VALUES LESS THAN (partition_value | MAXVALUE) [TABLESPACE
tablespacename]
```

**partition\_list\_item** defines a list partition. The syntax is as follows:

```
PARTITION partition_name VALUES (partition_value [, ...] | DEFAULT) [TABLESPACE
tablespacename]
```

**subpartition\_definition\_list** contains the **subpartition\_definition** object of one or more level-2 partitions. The syntax is as follows:

```
SUBPARTITION subpartition_name [VALUES LESS THAN (partition_value | MAXVALUE) | VALUES
(partition_value [, ...] | DEFAULT)] [TABLESPACE tablespacename]
```

---

### NOTICE

If the level-1 partition is a hash partition, you cannot use ADD to add a level-1 partition. If the level-2 partition is a hash partition, you cannot use MODIFY to add a level-2 partition.

- The `drop_clause` syntax is used to remove a partition from a specified partitioned table. The syntax can be used in level-1 partitions.

```
DROP PARTITION { partition_name | FOR (partition_value) } [UPDATE GLOBAL INDEX]
```

It can also be used in level-2 partitions.

```
DROP SUBPARTITION { subpartition_name | FOR (partition_value, subpartition_value) } [UPDATE
GLOBAL INDEX]
```

---

### NOTICE

- If the level-1 partition is a hash partition, the level-1 partition cannot be deleted. If the level-2 partition is a hash partition, the level-2 partition cannot be deleted.
- At least one sub-partition must be retained.

- The `split_clause` syntax is used to split one partition into partitions.  
`SPLIT SUBPARTITION { subpartition_name } { split_point_clause } [ UPDATE GLOBAL INDEX ]`

The `split_point_clause` syntax used to specify a split point in the range partitioning policy is as follows:

```
AT (subpartition_value) INTO (SUBPARTITION subpartition_name [TABLESPACE tablespacename] ,
SUBPARTITION subpartition_name [TABLESPACE tablespacename])
```

The `split_point_clause` syntax used to specify a split point in the list partitioning policy is as follows:

```
VALUES (subpartition_value) INTO (SUBPARTITION subpartition_name [TABLESPACE
tablespacename] , SUBPARTITION subpartition_name [TABLESPACE tablespacename])
```

#### NOTICE

- The size of the split point should be in the range of the splitting partition key.
  - One partition can be split into only two new partitions.
  - In the range partitioning policy, the current partition is split based on the split point into one partition with the range of smaller values and the other with the range of larger values. Therefore, only one split point can be used in this policy. In the list partitioning policy, there can be multiple but no more than 64 split points. These split points are extracted from the boundary values of the current partition as a new partition, and the remaining boundary values of the current partition are used as another new partition.
  - Only the default list partition can be split.
- 
- The `truncate_clause` syntax is used to remove a specified partition from a partitioned table.  
`TRUNCATE SUBPARTITION { subpartition_name } [ UPDATE GLOBAL INDEX ]`

#### NOTICE

When `truncate_clause` is executed in a level-2 partition, `AccessExclusiveLock` of the main table is obtained, blocking operations on the table.

## Parameters

- **table\_name**  
Specifies the name of a partitioned table.  
Value range: an existing partitioned table name.
- **subpartition\_name**  
Specifies the name of a level-2 partition name.  
Value range: an existing level-2 partition name.
- **tablespacename**  
Specifies which tablespace the partition moves to.  
Value range: an existing tablespace name.

## Examples

See [Examples](#) in "CREATE TABLE SUBPARTITION."

## Helpful Links

[CREATE TABLE SUBPARTITION](#) and [DROP TABLE](#)

## 7.13.29 ALTER TABLESPACE

### Description

Alters the attributes of a tablespace.

### Precautions

- Only the tablespace owner or a user granted the ALTER permission can run the **ALTER TABLESPACE** command. System administrators have this permission by default. To modify a tablespace owner, you must be the tablespace owner or a system administrator and a member of the new owner role.
- To change the owner, you must also be a direct or indirect member of the new owning role.

#### NOTE

If **new\_owner** is the same as **old\_owner**, the current user will not be verified. A message indicating successful **ALTER** execution is displayed.

### Syntax

- The syntax of renaming a tablespace is as follows:  

```
ALTER TABLESPACE tablespace_name
 RENAME TO new_tablespace_name;
```
- The syntax of setting the owner of a tablespace is as follows:  

```
ALTER TABLESPACE tablespace_name
 OWNER TO new_owner;
```
- The syntax of setting the attributes of a tablespace is as follows:  

```
ALTER TABLESPACE tablespace_name
 SET ({tablespace_option = value} [, ...]);
```
- The syntax of resetting the attributes of a tablespace is as follows:  

```
ALTER TABLESPACE tablespace_name
 RESET ({ tablespace_option } [, ...]);
```
- The syntax of setting the quota of a tablespace is as follows:  

```
ALTER TABLESPACE tablespace_name
 RESIZE MAXSIZE { UNLIMITED | 'space_size'};
```

### Parameters

- **tablespace\_name**  
Specifies the tablespace to be modified.  
Value range: an existing tablespace name.
- **new\_tablespace\_name**  
Specifies the new name of a tablespace.  
The new name cannot start with **PG\_**.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **new\_owner**  
Specifies the new owner of the tablespace.  
Value range: an existing username.
- **tablespace\_option**

Sets or resets the parameters of a tablespace.

Value range:

- **seq\_page\_cost**: sets the optimizer to calculate the cost of obtaining disk pages in sequence. The default value is **1.0**.
- **random\_page\_cost**: sets the optimizer to calculate the cost of obtaining disk pages in a non-sequential manner. The default value is **4.0**.

 **NOTE**

- The value of **random\_page\_cost** is relative to that of **seq\_page\_cost**. It is meaningless when the value is equal to or less than the value of **seq\_page\_cost**.
- The prerequisite for the default value **4.0** is that the optimizer uses indexes to scan table data and the hit ratio of table data in the cache is about 90%.
- If the size of the table data space is smaller than that of the physical memory, decrease the value to a proper level. On the contrary, if the hit ratio of table data in the cache is lower than 90%, increase the value.
- If random-access memory like SSD is adopted, the value can be decreased to a certain degree to reflect the cost of true random scan.

Value range: Positive number of the floating-point type.

- **RESIZE MAXSIZE**

Resets the maximum size of tablespace.

Value range:

- **UNLIMITED**: No limit is set for the tablespace.
- Determined by **space\_size**. For details about the format, see [CREATE TABLESPACE](#).

 **NOTE**

- If the adjusted quota is smaller than the current tablespace usage, the adjustment is successful. You need to decrease the tablespace usage to a value less than the new quota before writing data to the tablespace.
- You can also use the following statement to change the value of **MAXSIZE**:

```
ALTER TABLESPACE tablespace_name RESIZE MAXSIZE
{ 'UNLIMITED' | 'space_size'};
```

## Examples

See [Examples](#) in "CREATE TABLESPACE."

## Helpful Links

[CREATE TABLESPACE](#) and [DROP TABLESPACE](#)

## 7.13.30 ALTER TRIGGER

### Function

**ALTER TRIGGER** renames a trigger.

 **NOTE**

Currently, only the name can be modified.

## Precautions

Only the owner of a table where the trigger is created and a system administrator can run the **ALTER TRIGGER** statement.

## Syntax

```
ALTER TRIGGER trigger_name ON table_name RENAME TO new_name;
```

## Parameter Description

- **trigger\_name**  
Specifies the name of the trigger to be modified.  
Value range: an existing trigger
- **table\_name**  
Specifies the name of the table where the trigger to be modified is located.  
Value range: an existing table having a trigger
- **new\_name**  
Specifies the new name after modification.  
Value range: a string, which complies with the identifier naming convention. A value contains a maximum of 63 characters and cannot be the same as other triggers on the same table.

## Examples

See examples in [CREATE TRIGGER](#).

## Helpful Links

[CREATE TRIGGER](#), [DROP TRIGGER](#), and [ALTER TABLE](#)

## 7.13.31 ALTER TYPE

### Description

Alters the definition of a type.

### Precautions

Only the type owner or a user granted with the ALTER permission can run the **ALTER TYPE** command. System administrators have this permission by default. To modify the owner or schema of a type, you must be a type owner or system administrator and a member of the new owner role.

### Syntax

- Modify a type.  

```
ALTER TYPE name action [, ...];
```

The clauses corresponding to **action** are as follows:

  - Add a new attribute to a composite type.  

```
ADD ATTRIBUTE attribute_name data_type [COLLATE collation] [CASCADE | RESTRICT]
```

- Drop an attribute from a composite type.  
DROP ATTRIBUTE [ IF EXISTS ] attribute\_name [ CASCADE | RESTRICT ]
- Alter the type of an attribute in a composite type.  
ALTER ATTRIBUTE attribute\_name [ SET DATA ] TYPE data\_type [ COLLATE collation ] [ CASCADE | RESTRICT ]
- Add a new attribute to a composite type.  
ALTER TYPE name ADD ATTRIBUTE attribute\_name data\_type [ COLLATE collation ] [ CASCADE | RESTRICT ];
- Drop an attribute from a composite type.  
ALTER TYPE name DROP ATTRIBUTE [ IF EXISTS ] attribute\_name [ CASCADE | RESTRICT ];
- Alter the type of an attribute in a composite type.  
ALTER TYPE name ALTER ATTRIBUTE attribute\_name [ SET DATA ] TYPE data\_type [ COLLATE collation ] [ CASCADE | RESTRICT ];
- Alter the owner of a type.  
ALTER TYPE name OWNER TO { new\_owner | CURRENT\_USER | SESSION\_USER };
- Alter the name of a type.  
ALTER TYPE name RENAME TO new\_name;
- Alter the name of an attribute in a composite type.  
ALTER TYPE name RENAME ATTRIBUTE attribute\_name TO new\_attribute\_name [ CASCADE | RESTRICT ];
- Move a type to a new schema.  
ALTER TYPE name SET SCHEMA new\_schema;
- Add a new value to an enumerated type.  
ALTER TYPE name ADD VALUE [ IF NOT EXISTS ] new\_enum\_value [ { BEFORE | AFTER } neighbor\_enum\_value ];
- Alter an enumerated value in the value list.  
ALTER TYPE name RENAME VALUE existing\_enum\_value TO new\_enum\_value;

## Parameters

- **name**  
Specifies the name of an existing type that needs to be modified (optionally schema-qualified).
- **new\_name**  
Specifies the new name of the type.
- **new\_owner**  
Specifies the new owner of the type.
- **new\_schema**  
Specifies the new schema of the type.
- **attribute\_name**  
Specifies the name of the attribute to be added, modified, or deleted.
- **new\_attribute\_name**  
Specifies the new name of the attribute to be renamed.
- **data\_type**  
Specifies the data type of the attribute to be added, or the new type of the attribute to be modified.
- **new\_enum\_value**  
Specifies a new enumerated value. It is a non-null string with a maximum length of 63 bytes.

- **neighbor\_enum\_value**  
Specifies an existing enumerated value before or after which a new enumerated value will be added.
- **existing\_enum\_value**  
Specifies an enumerated value to be changed. It is a non-null string with a maximum length of 63 bytes.
- **CASCADE**  
Determines that the type to be modified, its associated records, and subtables that inherit the type will all be updated.
- **RESTRICT**  
Refuses to update the associated records of the modified type. This is the default action.

---

#### NOTICE

- **ADD ATTRIBUTE, DROP ATTRIBUTE, and ALTER ATTRIBUTE** can be combined for processing. For example, it is possible to add several attributes or change the types of several attributes at the same time in one command.
  - To change a schema of a type, you must have the CREATE permission on the new schema. To change the owner, you must be a direct or indirect member of the new owning role, and the member must have the CREATE permission on the schema of this type. (These restrictions enforce that the user can only rebuild and drop the type. However, system administrators can change ownership of any type in any way.) To add an attribute or change the type of an attribute, you must also have the USAGE permission of this type.
- 

## Examples

See [Examples](#) in "CREATE TYPE."

## Helpful Links

[CREATE TYPE](#) and [DROP TYPE](#)

## 7.13.32 ALTER USER

### Description

Alters the attributes of a database user.

### Precautions

Session parameters modified by ALTER USER apply to a specified user and take effect in the next session.

## Syntax

- Modify user permissions or other information.

```
ALTER USER user_name [[WITH] option [...]];
```

The option clause is as follows:

```
{ CREATEDB | NOCREATEDB }
| { CREATEROLE | NOCREATEROLE }
| { INHERIT | NOINHERIT }
| { AUDITADMIN | NOAUDITADMIN }
| { SYSADMIN | NOSYSADMIN }
| { MONADMIN | NOMONADMIN }
| { OPRADMIN | NOOPRADMIN }
| { POLADMIN | NOPOLADMIN }
| { USEFT | NOUSEFT }
| { LOGIN | NOLOGIN }
| { REPLICATION | NOREPLICATION }
| { VCADMIN | NOVCADMIN }
| { PERSISTENCE | NOPERSISTENCE }
| CONNECTION LIMIT connlimit
| [ENCRYPTED | UNENCRYPTED] PASSWORD { 'password' [EXPIRED] | DISABLE | EXPIRED }
| [ENCRYPTED | UNENCRYPTED] IDENTIFIED BY { 'password' [REPLACE 'old_password' |
EXPIRED] | DISABLE }
| VALID BEGIN 'timestamp'
| VALID UNTIL 'timestamp'
| PERM SPACE 'spacelimit'
| PGUSER
```

- Change the username.

```
ALTER USER user_name
 RENAME TO new_name;
```

- Lock or unlock.

```
ALTER USER user_name
 ACCOUNT { LOCK | UNLOCK };
```

- Change the value of a specified parameter associated with the user.

```
ALTER USER user_name
 SET configuration_parameter { { TO | = } { value | DEFAULT } | FROM CURRENT };
```

- Reset the value of a specified parameter associated with the user.

```
ALTER USER user_name
 RESET { configuration_parameter | ALL };
```

## Parameters

- **user\_name**

Specifies the current username.

Value range: an existing username. If a username contains uppercase letters, enclose the name with double quotation marks ("").

- **new\_password**

Specifies a new password.

The new password must:

- Differ from the old password.
- Contain at least eight characters. This is the default length.
- Differ from the username or the username spelled backward.
- Contain at least three of the following character types: uppercase characters, lowercase characters, digits, and special characters (limited to ~!@#\$%^&\*()-\_+=\|[]{};,:<.>/?). If the password contains characters other than the preceding characters, an error will be reported during statement execution.



- The user password should be enclosed in single quotation marks (').
- Value range: a string.
- **old\_password**  
Specifies the old password.
- **ACCOUNT { LOCK | UNLOCK }**
  - **ACCOUNT LOCK**: locks an account to forbid login to databases.
  - **ACCOUNT UNLOCK**: unlocks an account to allow login to databases.
- **PGUSER**  
In the current version, the **PGUSER** attribute of a user cannot be modified.

For details about other parameters, see "Parameters" in [CREATE ROLE](#) and [ALTER ROLE](#).

## Examples

See [Examples](#) in "CREATE USER."

## Helpful Links

[CREATE ROLE](#), [CREATE USER](#), and [DROP USER](#)

## 7.13.33 ALTER USER MAPPING

### Description

ALTER USER MAPPING is used to change the definition of the mapping from a user to a foreign server.

### Precautions

- If the **password** option is displayed, ensure that the **usermapping.key.cipher** and **usermapping.key.rand** files exist in the *\$GAUSSHOME/bin* directory of each node in GaussDB. If the two files do not exist, use the `gs\_guc` tool to generate them and use the `gs\_ssh` tool to release them to the *\$GAUSSHOME/bin* directory on each node.
- When multi-layer quotation marks are used for sensitive columns (such as **password**) in **OPTIONS**, the semantics is different from that in the scenario where quotation marks are not used. Therefore, sensitive columns are not identified for anonymization.

### Syntax

```
ALTER USER MAPPING FOR { user_name | USER | CURRENT_USER | PUBLIC }
 SERVER server_name
 OPTIONS ([ADD | SET | DROP] option ['value'] [, ...])
```

In **OPTIONS**, **ADD**, **SET**, and **DROP** are operations to be performed. If these operations are not specified, **ADD** operations will be performed by default. **option** and **value** are the parameters and values of the corresponding operation.

## Parameters

- **user\_name**  
Specifies username of the mapping.  
**CURRENT\_USER** and **USER** match the name of the current user. **PUBLIC** is used to match all current and future usernames in the system.
- **server\_name**  
Specifies name of the server to which the user is mapped.
- **OPTIONS**  
Changes an option for the user mapping. The new option overwrites any previously specified option. **ADD**, **SET**, and **DROP** are operations to be performed. If the operation is not specified explicitly, ADD operations will be performed. The option name must be unique and will be validated with the foreign data wrapper of the server.

## Examples

For details, see [Examples](#) in "CREATE USER MAPPING."

## Helpful Links

[CREATE USER MAPPING](#) and [DROP USER MAPPING](#)

## 7.13.34 ALTER VIEW

### Description

ALTER VIEW alters the auxiliary attributes of a view. If you want to change the query definition of a view, use CREATE OR REPLACE VIEW.

### Precautions

Only the view owner or a user granted the ALTER permission can run the **ALTER VIEW** command. System administrators have this permission by default. The following are permission restrictions depending on attributes to be modified:

- To modify the schema of a view, you must be the owner of the view or a system administrator and have the CREATE permission on the new schema.
- To modify the owner of a view, you must be the owner of the view or a system administrator and a member of the new owner role, with the CREATE permission on the schema of the view.
- Do not change the type of a column in a view.

### Syntax

- Set the default value of a view column.  

```
ALTER VIEW [IF EXISTS] view_name
 ALTER [COLUMN] column_name SET DEFAULT expression;
```
- Remove the default value of a view column.  

```
ALTER VIEW [IF EXISTS] view_name
 ALTER [COLUMN] column_name DROP DEFAULT;
```

- Change the owner of a view.  

```
ALTER VIEW [IF EXISTS] view_name
OWNER TO new_owner;
```
- Rename a view.  

```
ALTER VIEW [IF EXISTS] view_name
RENAME TO new_name;
```
- Set the schema of a view.  

```
ALTER VIEW [IF EXISTS] view_name
SET SCHEMA new_schema;
```
- Set the options of a view.  

```
ALTER VIEW [IF EXISTS] view_name
SET ({ view_option_name [= view_option_value] } [, ...]);
```
- Reset the options of a view.  

```
ALTER VIEW [IF EXISTS] view_name
RESET (view_option_name [, ...]);
```

## Parameters

- **IF EXISTS**  
If this option is used, no error is generated when the view does not exist, and only a message is displayed.
- **view\_name**  
Specifies the view name, which can be schema-qualified.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Specifies an optional list of names to be used for columns of the view. If not given, the column names are deduced from the query.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **SET/DROP DEFAULT**  
Sets or drops the default value of a column. This parameter does not take effect.
- **new\_owner**  
Specifies the new owner of a view.
- **new\_name**  
Specifies the new view name.
- **new\_schema**  
Specifies the new schema of the view.
- **view\_option\_name [ = view\_option\_value ]**  
Specifies an optional parameter for a view.  
Currently, **view\_option\_name** supports only the **security\_barrier** parameter.  
Use this parameter when the view provides row-level security.  
Value range: Boolean type, **TRUE** or **FALSE**.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer table.
openGauss=# CREATE TABLE tpcds.customer
```

```
(
c_customer_sk INTEGER NOT NULL,
c_customer_id CHARACTER(16) NOT NULL
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.customer VALUES (1, 'AAAAAAAAABAAAAAAA'),(100,
'AAAAAAAAACAAAAAAA'),(150, 'AAAAAAAADAAAAAAA');

-- Create a view consisting of rows with c_customer_sk less than 150.
openGauss=# CREATE VIEW tpcds.customer_details_view_v1 AS
 SELECT * FROM tpcds.customer
 WHERE c_customer_sk < 150;

-- Rename a view.
openGauss=# ALTER VIEW tpcds.customer_details_view_v1 RENAME TO customer_details_view_v2;

-- Change the schema of a view.
openGauss=# ALTER VIEW tpcds.customer_details_view_v2 SET schema public;

-- Drop a view.
openGauss=# DROP VIEW public.customer_details_view_v2;

-- Drop the tpcds.customer table.
openGauss=# DROP TABLE tpcds.customer;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[CREATE VIEW](#) and [DROP VIEW](#)

## 7.13.35 ANALYZE | ANALYSE

### Description

- ANALYZE collects statistics about ordinary tables in a database and stores the results in the PG\_STATISTIC and PG\_STATISTIC\_EXT system catalogs. After you run the **ANALYZE** command, you can query the collected statistics in the preceding system catalogs or system views PG\_STATS and PG\_EXT\_STATS. The execution plan generator uses these statistics to determine which one is the most effective execution plan.
- If no parameters are specified, ANALYZE analyzes each table and partitioned table in the database. You can also specify **table\_name**, **column\_name**, and **partition\_name** to limit the analysis to a specified table, column, or partitioned table.
- **ANALYZE | ANALYZE VERIFY** is used to check whether data files of ordinary tables (row-store tables) in a database are damaged.

### Precautions

- Non-temporary tables cannot be analyzed in an anonymous block, transaction block, function, or stored procedure. Temporary tables in a stored procedure can be analyzed but their statistics updates cannot be rolled back.
- Remote read is not involved in the ANALYZE VERIFY scenario. Therefore, the remote read parameter does not take effect. If the system detects that a page is damaged due to an error in a key system catalog, the system directly reports an error and does not continue the detection.

- With no parameter specified, ANALYZE processes all the tables on which the current user has permission to analyze in the current database. With tables specified, ANALYZE processes only the specified tables.
- To perform ANALYZE operation to a table, you must be a table owner or a user granted the VACUUM permission on the table. By default, system administrators have this permission. However, database owners are allowed to analyze all tables in their databases, except shared catalogs. (The restriction for shared catalogs means that a true database-wide ANALYZE operation can only be executed by system administrators). ANALYZE skips tables on which users do not have permissions.
- ANALYZE does not collect columns for which comparison or equivalent operations cannot be performed, for example, columns of the cursor type.

## Syntax

- Collect statistics information about a table.  

```
{ ANALYZE | ANALYSE } [VERBOSE]
[table_name [(column_name [, ...])]] ;
```
- Collect partition statistics information about a partitioned table. This syntax is not supported currently.  

```
{ ANALYZE | ANALYSE } [VERBOSE]
table_name [(column_name [, ...])] PARTITION (partition_name) ;
```

### NOTE

An ordinary partitioned table supports the syntax but not the function of collecting statistics about specified partitions.

- Check the data files in the current database.  

```
{ANALYZE | ANALYSE} VERIFY {FAST|COMPLETE};
```

### NOTE

- In fast mode, DML operations need to be performed on the tables to be verified concurrently. As a result, an error is reported during the verification. In the current fast mode, data is directly read from the disk. When other threads modify files concurrently, the obtained data is incorrect. Therefore, you are advised to perform the verification offline.
- You can perform operations on the entire database. Because a large number of tables are involved, you are advised to save the results in redirection mode.  

```
gsql -d database -p port -f sqlfile> sqllog.txt 2>&1
```
- NOTICE is used to check only tables that are visible to external systems. The detection of internal tables is included in the external tables on which NOTICE depends and is not displayed externally.
- This statement can be executed with error tolerance.
- If a key system catalog is damaged during a full database operation, an error is reported and the operation stops.
- Check data files of tables and indexes.  

```
{ANALYZE | ANALYSE} VERIFY {FAST|COMPLETE} {table_name|index_name} [CASCADE];
```

 **NOTE**

- Operations on ordinary tables and index tables are supported, but **CASCADE** operations on indexes of index tables are not supported. The **CASCADE** mode is used to process all index tables of the main table. When the index tables are checked separately, the **CASCADE** mode is not required.
- When the main table is checked, the internal tables of the main table, such as the TOAST table and cudesc table, are also checked.
- When the system displays a message indicating that the index table is damaged, you are advised to run the **reindex** command to rebuild the index.
- Check the data files of the table partition.  
`{ANALYZE | ANALYSE} VERIFY {FAST|COMPLETE} table_name PARTITION (partition_name) [CASCADE];`

 **NOTE**

You can check a single partition of a table, but cannot perform the **CASCADE** operation on the indexes of an index table.

## Parameters

- **VERBOSE**

Enables the display of progress messages.

 **NOTE**

If **VERBOSE** is specified, **ANALYZE** displays the progress information, indicating the table that is being processed. Statistics about tables are also displayed.

- **table\_name**

Specifies the name (possibly schema-qualified) of a specific table to analyze. If omitted, all regular tables (but not foreign tables) in the current database are analyzed.

Currently, you can use **ANALYZE** to collect statistics only from row-store tables.

Value range: an existing table name.

- **column\_name, column\_1\_name, column\_2\_name**

Specifies the name of a specific column to analyze. All columns are analyzed by default.

Value range: an existing column name.

- **partition\_name**

For a partitioned table, you can specify **partition\_name** following the keyword **PARTITION** to analyze the statistics of this table. Currently, **ANALYZE** can be performed on partitioned tables, but statistics of specified partitions cannot be analyzed.

Value range: a partition name of a table.

- **index\_name**

Specifies the name of the specific index table to be analyzed (possibly schema-qualified).

Value range: an existing table name.

- **FAST|COMPLETE**

For a row-store table, the **FAST** mode verifies the CRC and page header of the row-store table. If the verification fails, an alarm is generated. In **COMPLETE** mode, the pointer and tuple of the row-store table are parsed and verified.

- **CASCADE**

In **CASCADE** mode, all indexes of the current table are verified.

## Examples

-- Create a table.

```
openGauss=# CREATE TABLE customer_info
(
 WR_RETURNED_DATE_SK INTEGER
 WR_RETURNED_TIME_SK INTEGER
 WR_ITEM_SK INTEGER NOT NULL,
 WR_REFUNDED_CUSTOMER_SK INTEGER
)
;
```

-- Create a partitioned table.

```
openGauss=# CREATE TABLE customer_par
(
 WR_RETURNED_DATE_SK INTEGER
 WR_RETURNED_TIME_SK INTEGER
 WR_ITEM_SK INTEGER NOT NULL,
 WR_REFUNDED_CUSTOMER_SK INTEGER
)
PARTITION BY RANGE(WR_RETURNED_DATE_SK)
(
 PARTITION P1 VALUES LESS THAN(2452275),
 PARTITION P2 VALUES LESS THAN(2452640),
 PARTITION P3 VALUES LESS THAN(2453000),
 PARTITION P4 VALUES LESS THAN(MAXVALUE)
)
ENABLE ROW MOVEMENT;
```

-- Use ANALYZE to update statistics.

```
openGauss=# ANALYZE customer_info;
openGauss=# ANALYZE customer_par;
```

-- Use ANALYZE VERBOSE to update statistics and display table information.

```
openGauss=# ANALYZE VERBOSE customer_info;
INFO: analyzing "public.customer_info"(datanode pid=38661)
INFO: ANALYZE INFO : estimate total rows of "customer_info": scanned 0 pages of total 0 pages with 1
retry times, containing 0 live rows and 0 dead rows, estimated 0 total rows(datanode pid=38661)
INFO: ANALYZE INFO : "customer_info": scanned 0 of 0 pages, containing 0 live rows and 0 dead rows; 0
rows in sample, 0 estimated total rows(datanode pid=38661)
ANALYZE
```

### NOTE

If any environment-related fault occurs, check the logs of the primary node of the database.

-- Drop the table.

```
openGauss=# DROP TABLE customer_info;
openGauss=# DROP TABLE customer_par;
```

## 7.13.36 BEGIN

### Description

BEGIN may be used to initiate an anonymous block or a single transaction. This section describes the syntax of BEGIN used to initiate an anonymous block. For details about the BEGIN syntax that initiates transactions, see [START TRANSACTION](#).

An anonymous block is a structure that can dynamically create and execute stored procedure code instead of permanently storing code as a database object in the database.

### Precautions

None

### Syntax

- Enable an anonymous block.  

```
[DECLARE [declare_statements]]
BEGIN
 execution_statements
END;
/
```
- Start a transaction.  

```
BEGIN [WORK | TRANSACTION]
[
 {
 ISOLATION LEVEL { READ COMMITTED | SERIALIZABLE | REPEATABLE READ }
 | { READ WRITE | READ ONLY }
 } [, ...]
];
```

### Parameters

- **declare\_statements**  
Declares a variable, including its name and type, for example, **sales\_cnt int**.
- **execution\_statements**  
Specifies the statement to be executed in an anonymous block.  
Value range: DML operations (such as select, insert, delete, and update) or registered functions in the system catalog.

### Examples

```
-- Generate a string using an anonymous block.
openGauss=# BEGIN
dbe_output.print_line('Hello');
END;
/
```

### Helpful Links

[START TRANSACTION](#)



## 7.13.37 CALL

### Description

**CALL** calls defined functions and stored procedures.

### Precautions

The owner of a function or stored procedure, users granted the EXECUTE permission on the function or stored procedure, or users granted the EXECUTE ANY FUNCTION permission can call the function or stored procedure. System administrators have the permission to call the function or stored procedure by default.

### Syntax

```
CALL [schema.|package.] {func_name| procedure_name} (param_expr);
```

### Parameters

- **schema**  
Specifies the name of the schema where a function or stored procedure is located.
- **package**  
Specifies the name of the package where a function or stored procedure is located.
- **func\_name**  
Specifies the name of the function or stored procedure to be called.  
Value range: an existing function name.
- **param\_expr**  
Specifies a list of parameters. Use := or => to separate a parameter name and its value. This method allows parameters to be placed in any order. If only parameter values are in the list, the value order must be the same as that defined in the function or stored procedure.  
Value range: an existing function parameter name or stored procedure parameter name.

#### NOTE

The parameters include input parameters (whose name and type are separated by **IN**) and output parameters (whose name and type are separated by **OUT**). When you run the **CALL** command to call a function or stored procedure, the parameter list must contain an output parameter for non-overloaded functions. You can set the output parameter to a variable or any constant. For details, see [Examples](#). For an overloaded package function, the parameter list can have no output parameter, but the function may not be found. If an output parameter is contained, it must be a constant.

### Examples

```
-- Create a function named func_add_sql, calculate the sum of two integers, and return the result.
openGauss=# CREATE FUNCTION func_add_sql(num1 integer, num2 integer) RETURN integer
AS
BEGIN
RETURN num1 + num2;
```

```
END;
/

-- Pass by parameter value.
openGauss=# CALL func_add_sql(1, 3);

-- Pass by naming tag method.
openGauss=# CALL func_add_sql(num1 => 1,num2 => 3);
openGauss=# CALL func_add_sql(num2 := 2, num1 := 3);

-- Drop the function.
openGauss=# DROP FUNCTION func_add_sql;

-- Create a function with output parameters.
openGauss=# CREATE FUNCTION func_increment_sql(num1 IN integer, num2 IN integer, res OUT integer)
RETURN integer
AS
BEGIN
res := num1 + num2;
END;
/

-- Specify a constant as an output parameter.
openGauss=# CALL func_increment_sql(1,2,1);

-- Drop the function.
openGauss=# DROP FUNCTION func_increment_sql;
```

## Helpful Links

[CREATE FUNCTION](#) and [CALL](#)

## 7.13.38 CHECKPOINT

### Function

A checkpoint is a point in the transaction log sequence at which all data files have been updated to reflect the information in the log. All data files will be flushed to a disk.

**CHECKPOINT** forces a transaction log checkpoint. By default, WALs periodically specify checkpoints in a transaction log. You may use **gs\_guc** to specify run-time parameters **checkpoint\_segments**, **checkpoint\_timeout**, and **incremental\_checkpoint\_timeout** to adjust the atomized checkpoint intervals.

### Precautions

- Only the system administrator and O&M administrator can invoke **CHECKPOINT**.
- **CHECKPOINT** forces an immediate checkpoint when the related command is issued, without waiting for a regular checkpoint scheduled by the system.

### Syntax

```
CHECKPOINT;
```

### Parameter Description

None

## Examples

```
-- Set a checkpoint.
openGauss=# CHECKPOINT;
```

## 7.13.39 CLEAN CONNECTION

### Description

Clears database connections. You may use this statement to delete a specific user's connections to a specified database.

### Precautions

- GaussDB does not support specified nodes and supports only TO ALL.
- This function can be used to clear the normal connections that are being used only in force mode.

### Syntax

```
CLEAN CONNECTION
 TO { COORDINATOR (nodename [, ...]) | NODE (nodename [, ...]) | ALL [CHECK] [FORCE] }
 [FOR DATABASE dbname]
 [TO USER username];
```

### Parameters

- **CHECK**  
This parameter can be specified only when the node list is specified as **TO ALL**. Setting this parameter will check whether a database is accessed by other sessions before its connections are cleared. If any sessions are detected before **DROP DATABASE** is executed, an error will be reported and the database will not be dropped.
- **FORCE**  
This parameter can be specified only when the node list is specified as **TO ALL**. Setting this parameter will send **SIGTERM** signals to all the threads related to the specified **dbname** and **username** and forcibly shut them down.
- **COORDINATOR ( nodename [, ... ] ) | NODE ( nodename [, ... ] ) | ALL**  
Only **TO ALL** is supported. This parameter must be specified. All specified connections on the node will be deleted.
- **dbname**  
Deletes connections to a specified database. If this parameter is not specified, connections to all databases will be deleted.  
Value range: an existing database name.
- **username**  
Deletes connections of a specific user. If this parameter is not specified, connections of all users will be deleted.  
Value range: an existing username.

## Examples

```
-- Create the test_clean_connection database.
openGauss=# CREATE DATABASE test_clean_connection;
```

```
-- Create user jack.
openGauss=# CREATE USER jack PASSWORD '*****';

-- Clean the user jack's connections to the template1 database.
openGauss=# CLEAN CONNECTION TO ALL FOR DATABASE template1 TO USER jack;

-- Clean all connections of user jack.
openGauss=# CLEAN CONNECTION TO ALL TO USER jack;

-- Clean all the connections to the test_clean_connection database.
openGauss=# CLEAN CONNECTION TO ALL FORCE FOR DATABASE test_clean_connection;

-- Drop user jack.
openGauss=# DROP USER jack;

-- Drop the test_clean_connection database.
openGauss=# DROP DATABASE test_clean_connection;
```

## 7.13.40 CLOSE

### Function

**CLOSE** frees the resources associated with an open cursor.

### Precautions

- After a cursor is closed, no subsequent operations are allowed on it.
- A cursor should be closed when it is no longer needed.
- Every non-holdable open cursor is implicitly closed when a transaction is terminated by **COMMIT** or **ROLLBACK**.
- A holdable cursor is implicitly closed if the transaction that created it aborts by **ROLLBACK**.
- If the cursor creation transaction is successfully committed, the holdable cursor remains open until an explicit **CLOSE** operation is executed, or the client disconnects.
- GaussDB does not have an explicit **OPEN** cursor statement. A cursor is considered open when it is declared. You can view all available cursors by querying the **pg\_cursors** system view.

### Syntax

```
CLOSE { cursor_name | ALL } ;
```

### Parameter Description

- **cursor\_name**  
Specifies the name of a cursor to be closed.
- **ALL**  
Closes all open cursors.

### Examples

See [Examples](#) in **FETCH**.

## Helpful Links

[FETCH](#) and [MOVE](#)

### 7.13.41 CLUSTER

#### Description

- Clusters a table based on an index.
- CLUSTER instructs GaussDB to cluster the table specified by **table\_name** based on the index specified by **index\_name**. The index must have been defined by **table\_name**.
- When a table is clustered, it is physically reordered based on the index information. Clustering is a one-time operation. When the table is subsequently updated, the changes are not clustered. That is, no attempt is made to store new or updated rows according to their index order.
- When a table is clustered, GaussDB records which index the table was clustered by. CLUSTER table\_name reclusters the clustered index that was previously recorded in the table. You can also use ALTER TABLE table\_name CLUSTER on index\_name to set the index of a specified table for subsequent CLUSTER operations, or use ALTER TABLE table\_name SET WITHOUT CLUSTER to clear the previously clustered index of a specified table.
- If CLUSTER does not contain parameters, all tables that have been clustered in the database owned by the current user will be reprocessed. If a system administrator uses this command, all clustered tables are reclustered.
- When a table is being clustered, an ACCESS EXCLUSIVE lock is acquired on it. This prevents any other database operations (both read and write) from being performed on the table until the CLUSTER operation is finished.

#### Precautions

- Only row-store B-tree indexes support CLUSTER.
- In the case where you are accessing single rows randomly within a table, the actual order of the data in the table is unimportant. However, if there are many accesses to some data and an index groups the data, using the CLUSTER index improves performance.
- If you use an index to query a table for a range or multiple rows, CLUSTER will also help because once the index identifies the table page for the first row that matches, all other rows that match are probably already on the same table page. This saves disk accesses and speeds up the query.
- During clustering, the system creates a temporary backup of the table created in the index sequence and a temporary backup of each index in the table. Therefore, ensure that the disk has sufficient free space during clustering, which is at least the sum of the table size and all index sizes.
- CLUSTER records which indexes have been used for clustering. Therefore, you can manually specify indexes for the first time, cluster specified tables, and set a maintenance script that will be executed periodically. You only need to run the **CLUSTER** command without parameters. In this way, tables that you want to periodically cluster can be automatically updated.
- The optimizer records table clustering statistics. After clustering a table, you need to execute the ANALYZE operation to ensure that the optimizer has the

latest clustering information. Otherwise, the optimizer may select a non-optimal query plan.

- **CLUSTER** cannot be executed in transactions.
- If the GUC parameter **xc\_maintenance\_mode** is not set to **on**, the **CLUSTER** operation skips all system catalogs.

## Syntax

- Cluster a table.  
`CLUSTER [ VERBOSE ] table_name [ USING index_name ];`
- Cluster a partition.  
`CLUSTER [ VERBOSE ] table_name PARTITION ( partition_name ) [ USING index_name ];`
- Recluster a table.  
`CLUSTER [ VERBOSE ];`

## Parameters

- **VERBOSE**  
Enables the display of progress messages.
- **table\_name**  
Specifies the table name.  
Value range: an existing table name.
- **index\_name**  
Specifies the index name.  
Value range: name of an existing index.
- **partition\_name**  
Specifies the partition name.  
Value range: an existing partition name.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create a partitioned table.
openGauss=# CREATE TABLE tpcds.inventory_p1
(
 INV_DATE_SK INTEGER NOT NULL,
 INV_ITEM_SK INTEGER NOT NULL,
 INV_WAREHOUSE_SK INTEGER NOT NULL,
 INV_QUANTITY_ON_HAND INTEGER
)
PARTITION BY RANGE(INV_DATE_SK)
(
 PARTITION P1 VALUES LESS THAN(2451179),
 PARTITION P2 VALUES LESS THAN(2451544),
 PARTITION P3 VALUES LESS THAN(2451910),
 PARTITION P4 VALUES LESS THAN(2452275),
 PARTITION P5 VALUES LESS THAN(2452640),
 PARTITION P6 VALUES LESS THAN(2453005),
 PARTITION P7 VALUES LESS THAN(MAXVALUE)
);

-- Create an index named ds_inventory_p1_index1.
openGauss=# CREATE INDEX ds_inventory_p1_index1 ON tpcds.inventory_p1 (INV_ITEM_SK) LOCAL;
```

```
-- Cluster the tpcds.inventory_p1 table.
openGauss=# CLUSTER tpcds.inventory_p1 USING ds_inventory_p1_index1;

-- Cluster the p3 partition.
openGauss=# CLUSTER tpcds.inventory_p1 PARTITION (p3) USING ds_inventory_p1_index1;

-- Recluster a table.
openGauss=# CLUSTER;

-- Drop the index.
openGauss=# DROP INDEX tpcds.ds_inventory_p1_index1;

-- Drop the partitioned table.
openGauss=# DROP TABLE tpcds.inventory_p1;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.42 COMMENT

### Description

Defines or changes the comment of an object.

### Precautions

- Each object stores only one comment. Therefore, you need to modify a comment and issue a new **COMMENT** command to the same object. To delete the comment, write **NULL** at the position of the text string. When an object is deleted, the comment is automatically deleted.
- Currently, there is no security protection for viewing comments. Any user connected to a database can view all the comments for objects in the database. For shared objects such as databases, roles, and tablespaces, comments are stored globally so any user connected to any database in the cluster can see all the comments for shared objects. Therefore, do not put security-critical information in comments.
- To comment objects, you must be an object owner or user granted the COMMENT permission. System administrators have this permission by default.
- Roles do not have owners, so the rule for COMMENT ON ROLE is that you must be a system administrator to comment on a SYSADMIN role, or have the CREATEROLE permission to comment on non-SYSADMIN roles. A system administrator can comment on all objects.

### Syntax

```
COMMENT ON
{
 AGGREGATE agg_name (agg_type [, ...]) |
 CAST (source_type AS target_type) |
 COLLATION object_name |
 COLUMN { table_name.column_name | view_name.column_name } |
 CONSTRAINT constraint_name ON table_name |
 CONVERSION object_name |
 DATABASE object_name |
 DOMAIN object_name |
 EXTENSION object_name |
 FOREIGN DATA WRAPPER object_name |
 FOREIGN TABLE object_name |
 FUNCTION function_name ([{[argname] [argmode] argtype} [, ...]]) |
 INDEX object_name |
```

```
LARGE OBJECT large_object_oid |
OPERATOR operator_name (left_type, right_type) |
OPERATOR CLASS object_name USING index_method |
OPERATOR FAMILY object_name USING index_method |
[PROCEDURAL] LANGUAGE object_name |
ROLE object_name |
SCHEMA object_name |
SERVER object_name |
TABLE object_name |
TABLESPACE object_name |
TEXT SEARCH CONFIGURATION object_name |
TEXT SEARCH DICTIONARY object_name |
TEXT SEARCH PARSER object_name |
TEXT SEARCH TEMPLATE object_name |
TYPE object_name |
VIEW object_name |
TRIGGER trigger_name ON table_name
}
IS 'text';
```

## Parameters

- **agg\_name**  
Name of an aggregate function.
- **agg\_type**  
Data type of the aggregate function parameters.
- **source\_type**  
Source data type in the type conversion.
- **target\_type**  
Target data type in the type conversion.
- **object\_name**  
Object name.
- **table\_name.column\_name**  
**view\_name.column\_name**  
Column name You can add the table name or view name as the prefix.
- **constraint\_name**  
Table constraint whose comment is defined or modified.
- **table\_name**  
Table name.
- **function\_name**  
Function name.
- **argname,argmode,argtype**  
Name, schema, and type of the function parameters.
- **large\_object\_oid**  
OID of a large object.
- **operator\_name**  
Operator name.
- **left\_type,right\_type**  
Data type of the operator parameters (optionally schema-qualified). If the prefix or suffix operator does not exist, the **NONE** option can be added.



- **trigger\_name**  
Trigger name.
- **text**  
Comments.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer table.
openGauss=# CREATE TABLE tpcds.customer
(
 c_customer_sk INTEGER NOT NULL,
 c_customer_id CHAR(16) NOT NULL
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.customer VALUES (50, 'AAAAAAAAABAAAAAA'),(100,
'AAAAAAAAACAAAAAA'),(150, 'AAAAAADAAAAAA');

-- Create the tpcds.customer_demographics_t2 table.
openGauss=# CREATE TABLE tpcds.customer_demographics_t2
(
 CD_DEMO_SK INTEGER NOT NULL,
 CD_GENDER CHAR(1) ,
 CD_MARITAL_STATUS CHAR(1) ,
 CD_EDUCATION_STATUS CHAR(20) ,
 CD_PURCHASE_ESTIMATE INTEGER ,
 CD_CREDIT_RATING CHAR(10) ,
 CD_DEP_COUNT INTEGER ,
 CD_DEP_EMPLOYED_COUNT INTEGER ,
 CD_DEP_COLLEGE_COUNT INTEGER
)
;

-- Comment out the tpcds.customer_demographics_t2.cd_demo_sk column.
openGauss=# COMMENT ON COLUMN tpcds.customer_demographics_t2.cd_demo_sk IS 'Primary key of
customer demographics table.';

-- Create a view consisting of rows with c_customer_sk less than 150.
openGauss=# CREATE VIEW tpcds.customer_details_view_v2 AS
 SELECT *
 FROM tpcds.customer
 WHERE c_customer_sk < 150;

-- Comment out the tpcds.customer_details_view_v2 view.
openGauss=# COMMENT ON VIEW tpcds.customer_details_view_v2 IS 'View of customer detail';

-- Drop the view.
openGauss=# DROP VIEW tpcds.customer_details_view_v2;

-- Drop the tpcds.customer_demographics_t2 table.
openGauss=# DROP TABLE tpcds.customer_demographics_t2;

-- Drop the tpcds.customer table.
openGauss=# DROP TABLE tpcds.customer;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.43 COMMIT | END

### Description

COMMIT or END commits all operations of a transaction.

### Precautions

Only the creator of a transaction or system administrators can run the **COMMIT** command. The creation and commit operations do not need to be in different sessions.

### Syntax

```
{ COMMIT | END } [WORK | TRANSACTION] ;
```

### Parameters

- **COMMIT | END**  
Commits the current transaction and makes all changes made by the transaction become visible to others.
- **WORK | TRANSACTION**  
Specifies an optional keyword, which has no effect except increasing readability.

### Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create a table.
openGauss=# CREATE TABLE tpcds.customer_demographics_t2
(
 CD_DEMO_SK INTEGER NOT NULL,
 CD_GENDER CHAR(1) ,
 CD_MARITAL_STATUS CHAR(1) ,
 CD_EDUCATION_STATUS CHAR(20) ,
 CD_PURCHASE_ESTIMATE INTEGER ,
 CD_CREDIT_RATING CHAR(10) ,
 CD_DEP_COUNT INTEGER ,
 CD_DEP_EMPLOYED_COUNT INTEGER ,
 CD_DEP_COLLEGE_COUNT INTEGER
);

-- Start a transaction.
openGauss=# START TRANSACTION;

-- Insert data.
openGauss=# INSERT INTO tpcds.customer_demographics_t2 VALUES(1,'M', 'U', 'DOCTOR DEGREE', 1200,
'GOOD', 1, 0, 0);
openGauss=# INSERT INTO tpcds.customer_demographics_t2 VALUES(2,'F', 'U', 'MASTER DEGREE', 300,
'BAD', 1, 0, 0);

-- Commit the transaction to make all changes permanent.
openGauss=# COMMIT;

-- Query data.
openGauss=# SELECT * FROM tpcds.customer_demographics_t2;

-- Drop the tpcds.customer_demographics_t2 table.
```

```
openGauss=# DROP TABLE tpceds.customer_demographics_t2;
-- Drop the schema.
openGauss=# DROP SCHEMA tpceds CASCADE;
```

## Helpful Links

[ROLLBACK](#)

## 7.13.44 COMMIT PREPARED

### Description

Commits a prepared two-phase transaction.

### Precautions

- The function is only available in maintenance mode (when the GUC parameter **xc\_maintenance\_mode** is **on**). Exercise caution when enabling the mode. It is used by maintenance engineers for troubleshooting. Common users should not use the mode.
- Only the creator of a transaction or system administrators can run the command. The creation and commit operations do not need to be in different sessions.
- The transaction function is maintained automatically by the database, and should be not visible to users.

### Syntax

```
COMMIT PREPARED transaction_id [WITH commit_sequence_number];
```

### Parameters

- **transaction\_id**  
Specifies the identifier of the transaction to be committed. The identifier must be different from those for current prepared transactions.
- **commit\_sequence\_number**  
Specifies the sequence number of the transaction to be committed. It is a 64-bit, incremental, unsigned number.

### Examples

```
-- Start.
openGauss=# begin;

-- Prepare a transaction whose identifier is trans_test.
openGauss=# PREPARE TRANSACTION 'trans_test';

-- Create a table.
openGauss=# CREATE TABLE item1(id int);

-- Commit the transaction whose identifier is trans_test.
openGauss=# COMMIT PREPARED 'trans_test';

-- Drop the table.
openGauss=# DROP TABLE item1;
```

## Helpful Links

[PREPARE TRANSACTION](#) and [ROLLBACK PREPARED](#)

## 7.13.45 COPY

### Description

COPY copies data between tables and files.

COPY FROM copies data from a file to a table, and COPY TO copies data from a table to a file.

### Precautions

- When the **enable\_copy\_server\_files** parameter is disabled, only the initial user is allowed to run the **COPY FROM FILENAME** or **COPY TO FILENAME** command. When the **enable\_copy\_server\_files** parameter is enabled, users with the SYSADMIN permission or users who inherit the permissions of the built-in role `gs_role_copy_files` are allowed to run the **COPY FROM FILENAME** or **COPY TO FILENAME** command. By default, database configuration files and key files are not allowed, and you can run the **COPY FROM FILENAME** or **COPY TO FILENAME** command for certificate files and audit logs to prevent unauthorized users from viewing or modifying sensitive files. When **enable\_copy\_server\_files** is set to **on**, administrators can use the GUC parameter **safe\_data\_path** to set the path for common users to import and export to the subpath of the set path. If this GUC parameter is not set (by default), the path used by common users is not blocked.
- **COPY** applies only to tables but not views.
- COPY TO requires the SELECT permission on the table to be read, and COPY FROM requires the INSERT permission on the table to be inserted.
- If a list of columns is specified, COPY copies only the data of the specified columns between the file and the table. If a table has any columns that are not in the column list, COPY FROM inserts default values into those columns.
- If a data source file is specified, the server must be able to access the file. If **STDIN** is specified, data flows between the client and the server. When entering data, use the **TAB** key to separate the columns of the table and use a backslash and a period (\.) in a new row to indicate the end of the input.
- COPY FROM throws an error if any row in the data file contains more or fewer columns than expected.
- The end of the data can be represented by a line that contains only backslashes and periods (\.). If data is read from a file, the end flag is unnecessary. If data is copied between client applications, an end tag must be provided.
- In COPY FROM, **\N** is an empty string. To enter the actual value **\N**, use **\\N**.
- COPY FROM can preprocess data using column expressions, but column expressions do not support subqueries.
- When a data format error occurs during **COPY FROM** execution, the transaction is rolled back. However, the error information is insufficient, making it difficult to locate the error data from a large amount of raw data.

- COPY FROM and COPY TO apply to low concurrency and local import and export of a small amount of data.
- If the target table has triggers, COPY is supported.
- Ensure that the generated column is not in the list of the specified column in the COPY statement. When COPY... TO is used to export data, if no column list is specified, all columns except the generated columns in the table are exported. When COPY... FROM is used to import data, the generated columns are automatically updated and saved as ordinary columns.
- During the export using COPY TO, if the column data in the table contains the '\0' character, the column data will be truncated during the export. Only the data before '\0' will be exported.

## Syntax

- Copy data from a file to a table.

```
COPY table_name [(column_name [, ...])]
 FROM { 'filename' | STDIN }
 [[USING] DELIMITERS 'delimiters']
 [WITHOUT ESCAPING]
 [LOG ERRORS]
 [REJECT LIMIT 'limit']
 [WITH (option [, ...])]
 | copy_option
 | TRANSFORM ({ column_name [data_type] [AS transform_expr] } [, ...])
 | FIXED FORMATTER ({ column_name(offset, length) } [, ...]) [(option [, ...]) | copy_option
 [...]]];
```

### NOTE

In the syntax, **FIXED FORMATTER ( { column\_name( offset, length ) } [, ...] )** and non-conflicting items of **[ copy\_option [ ...] ]** can be in any sequence.

- Copy data from a table to a file.

```
COPY table_name [(column_name [, ...])]
 TO { 'filename' | STDOUT }
 [[USING] DELIMITERS 'delimiters']
 [WITHOUT ESCAPING]
 [WITH (option [, ...])]
 | copy_option
 | FIXED FORMATTER ({ column_name(offset, length) } [, ...]) [(option [, ...]) | copy_option
 [...]]];
```

```
COPY query {(SELECT) | (VALUES)}
 TO { 'filename' | STDOUT }
 [WITHOUT ESCAPING]
 [WITH (option [, ...])]
 | copy_option
 | FIXED FORMATTER ({ column_name(offset, length) } [, ...]) [(option [, ...]) | copy_option
 [...]]];
```

### NOTE

1. The restrictions on the COPY TO syntax are as follows:  
(**query**) is incompatible with **[USING] DELIMITERS**. If the data comes from a query result, COPY TO cannot specify **[USING] DELIMITERS**.
2. Use spaces to separate **copy\_option** following **FIXED FORMATTER**.
3. **copy\_option** is the native parameter, while **option** is the parameter imported by a compatible foreign table.
4. In the syntax, **FIXED FORMATTER ( { column\_name( offset, length ) } [, ...] )** and non-conflicting items of **[ copy\_option [ ...] ]** can be in any sequence.

The optional parameters of the option clause are as follows:

```
FORMAT 'format_name'
| OIDS [boolean]
| DELIMITER 'delimiter_character'
| NULL 'null_string'
| HEADER [boolean]
| USEEOF [boolean]
| FILEHEADER 'header_file_string'
| FREEZE [boolean]
| QUOTE 'quote_character'
| ESCAPE 'escape_character'
| EOL 'newline_character'
| NOESCAPING [boolean]
| FORCE_QUOTE { (column_name [, ...]) | * }
| FORCE_NOT_NULL (column_name [, ...])
| ENCODING 'encoding_name'
| IGNORE_EXTRA_DATA [boolean]
| FILL_MISSING_FIELDS [boolean]
| COMPATIBLE_ILLEGAL_CHARS [boolean]
| DATE_FORMAT 'date_format_string'
| TIME_FORMAT 'time_format_string'
| TIMESTAMP_FORMAT 'timestamp_format_string'
| SMALLDATETIME_FORMAT 'smalldatetime_format_string'
```

The optional parameters of the `copy_option` clause are as follows:

```
OIDS
| NULL 'null_string'
| HEADER
| USEEOF
| FILEHEADER 'header_file_string'
| FREEZE
| FORCE_NOT_NULL column_name [, ...]
| FORCE_QUOTE { column_name [, ...] | * }
| BINARY
| CSV
| QUOTE [AS] 'quote_character'
| ESCAPE [AS] 'escape_character'
| EOL 'newline_character'
| ENCODING 'encoding_name'
| IGNORE_EXTRA_DATA
| FILL_MISSING_FIELDS [{ 'one' | 'multi' }]
| COMPATIBLE_ILLEGAL_CHARS
| DATE_FORMAT 'date_format_string'
| TIME_FORMAT 'time_format_string'
| TIMESTAMP_FORMAT 'timestamp_format_string'
| SMALLDATETIME_FORMAT 'smalldatetime_format_string'
| SKIP int_number
| WHEN { (start - end) | column_name } { = | != } 'string'
| SEQUENCE ({ column_name (integer [, incr]) [, ...] })
| FILLER ({ column_name [, ...] })
| CONSTANT ({ column_name 'constant_string' [, ...] })
```

## Parameters

- **query**  
Specifies that the results are to be copied.  
Value range: Only one `SELECT` or `VALUES` command is supported. A semicolon (;) is not required at the end of the command.
- **table\_name**  
Specifies the name (possibly schema-qualified) of an existing table.  
Value range: an existing table name.
- **column\_name**  
Specifies an optional list of columns to be copied.  
Value range: any columns. All columns will be copied if no column list is specified.

- **STDIN**  
Specifies that input comes from the standard input.
- **STDOUT**  
Specifies that output goes to the standard output.
- **FIXED**  
Fixes column length. When the column length is fixed, **DELIMITER**, **NULL**, and **CSV** cannot be specified. When **FIXED** is specified, **BINARY**, **CSV**, and **TEXT** cannot be specified by **option** or **copy\_option**.

 **NOTE**

The definition of fixed length is as follows:

1. The column length of each record is the same.
2. Spaces are used for column padding. Columns of the numeric type are left-aligned and columns of the string type are right-aligned.
3. No delimiters are used between columns.

- **[USING] DELIMITERS 'delimiters'**

The string that separates columns within each row (line) of the file, and it cannot be larger than 10 bytes.

Value range: The delimiter cannot include any of the following characters:  
\\abcdefghijklmnopqrstuvwxyz0123456789

Value range: The default value is a tab character in text format and a comma in CSV format.

- **WITHOUT ESCAPING**

Specifies, in text format, whether to escape the backslash (\\) and its following characters.

Value range: text only.

- **LOG ERRORS**

If this parameter is specified, the error tolerance mechanism for data type errors in the COPY FROM statement is enabled.

Value range: a value set while data is imported using COPY FROM.

 **NOTE**

The restrictions of this error tolerance parameter are as follows:

- This error tolerance mechanism captures only the data type errors (DATA\_EXCEPTION) that occur during data parsing of COPY FROM on the primary node of the database.
- If existing error tolerance parameters (for example, **IGNORE\_EXTRA\_DATA**) of the COPY statement are enabled, the error of the corresponding type will be processed as specified by the parameters and no error will be reported. Therefore, the error table does not contain such error data.

- **LOG ERRORS DATA**

The differences between **LOG ERRORS DATA** and **LOG ERRORS** are as follows:

- a. **LOG ERRORS DATA** fills the **rawrecord** column in the error tolerance table.
- b. Only users with the super permission can use the **LOG ERRORS DATA** parameter.

#### NOTICE

If error content is too complex, it may fail to be written to the error tolerance table by using **LOG ERRORS DATA**, causing the task failure.

For errors that cannot be read in certain code, the error codes are `ERRCODE_CHARACTER_NOT_IN_REPERTOIRE` and `ERRCODE_UNTRANSLATABLE_CHARACTER`. The **rawrecord** column is not recorded.

- **REJECT LIMIT 'limit'**

Used with the **LOG ERROR** option to set the upper limit of the tolerated errors in the COPY FROM statement. If the number of errors exceeds the limit, later errors will be reported based on the original mechanism.

Value range: a positive integer (1 to `INTMAX`) or **'unlimited'**.

Default value: If **LOG ERRORS** is not specified, an error will be reported. If **LOG ERRORS** is specified, the default value is **0**.

#### NOTE

In the error tolerance mechanism described in the description of **LOG ERRORS**, the count of **REJECT LIMIT** is calculated based on the number of data parsing errors on the primary node of the database where the **COPY FROM** statement is executed, not based on the number of all errors on the primary node.

- **FORMATTER**

Defines the place of each column in the data file in fixed length mode.

Defines the place of each column in the data file in the **column(offset,length)** format.

Value range:

- The value of **offset** must be larger than 0. The unit is byte.
- The value of **length** must be larger than 0. The unit is byte.

The total length of all columns must be less than 1 GB.

Replace columns that are not in the file with null.

- **OPTION { option\_name ' value ' }**

Specifies all types of parameters of a compatible foreign table.

- **FORMAT**

Specifies the format of the source data file in the foreign table.

Value range: **CSV**, **TEXT**, **FIXED**, and **BINARY**.

- The CSV file can process newline characters efficiently, but cannot process certain special characters well.
- The TEXT file can process certain special characters efficiently, but cannot process newline characters well.
- In FIXED files, the column length of each record is the same. Spaces are used for padding, and the excessive part will be truncated.
- All data in the BINARY file is stored/read as binary format rather than as text. It is faster than the text and CSV formats, but a binary-format file is less portable.



Default value: **TEXT**

– DELIMITER

Specifies the character that separates columns within each row (line) of the file.

 NOTE

- The value of **delimiter** cannot be `\r` or `\n`.
- A delimiter cannot be the same as the value of **null**. The delimiter for the CSV format cannot be same as the value of **quote**.
- The delimiter for the TEXT format data cannot contain lowercase letters, digits, or special characters (.,\).
- The data length of a single row should be less than 1 GB. A row that has many columns using long delimiters cannot contain much valid data.
- You are advised to use multi-character delimiters or invisible delimiters. For example, you can use multi-characters (such as `$^&`) and invisible characters (such as `0x07`, `0x08`, and `0x1b`).

Value range: a multi-character delimiter within 10 bytes.

Default value:

- A tab character in text format.
- A comma (,) in CSV format.
- No delimiter in FIXED format.

– NULL

Specifies the string that represents a null value.

Value range:

- A null value cannot be `\r` or `\n`. The maximum length is 100 characters.
- A null value cannot be the same as the **delimiter** or **quote** value.

Default value:

- The default value for the CSV format is an empty string without quotation marks.
- The default value for the TEXT format is `\N`.

– HEADER

Specifies whether a file contains a header with the names of each column in the file. **header** is available only for CSV and FIXED files.

When data is imported, if **header** is **on**, the first row of the data file will be identified as the header and ignored. If **header** is **off**, the first row will be identified as a data row.

When data is exported, if header is **on**, **fileheader** must be specified. If **header** is **off**, an exported file does not contain a header.

Value range: **true**, **on**, **false**, and **off**.

Default value: **false**

– HEADER

Does not report an error for \. in the imported data.

Value range: **true**, **on**, **false**, and **off**.

Default value: **false**

– QUOTE

Specifies a quoted character string for a CSV file.

Default value: double quotation marks ("").

 NOTE

- The value of **quote** cannot be the same as that of the **delimiter** or **null** parameter.
- The value of **quote** must be a single-byte character.
- You are advised to set **quote** to an invisible character, such as **0x07**, **0x08**, or **0x1b**.

– ESCAPE

Specifies an escape character for a CSV file. The value must be a single-byte character.

Default value: double quotation marks (""). If the value is the same as that of **quote**, it will be replaced by \0.

– EOL 'newline\_character'

Specifies the newline character style of the imported or exported data file.

Value range: multi-character newline characters within 10 bytes. Common newline characters include \r (0x0D), \n (0x0A), and \r\n (0x0D0A). Special newline characters include \$ and #.

 NOTE

- The **EOL** parameter supports only the TEXT format for data import and export and does not support the CSV or FIXED format for data import. For forward compatibility, the EOL parameter can be set to **0x0D** or **0x0D0A** for data export in the CSV or FIXED format.
- The value of **EOL** cannot be the same as that of the **delimiter** or **null** parameter.
- The EOL parameter value cannot contain the following characters: .abcdefghijklmnopqrstuvwxyz0123456789.

– FORCE\_QUOTE { ( column\_name [, ...] ) | \* }

In **CSV COPY TO** mode, forces quotation marks to be used for all non-null values in each specified column. **NULL** values are not quoted.

Value range: name of an existing column.

– FORCE\_NOT\_NULL ( column\_name [, ...] )

In **CSV COPY FROM** mode, the value for a specified column cannot be null.

Value range: name of an existing column.

– ENCODING

Specifies the encoding of data files. If this option is omitted, the current client-side encoding is used.

– IGNORE\_EXTRA\_DATA

Specifies whether to ignore excessive columns when the number of data source files exceeds the number of foreign table columns. This parameter is used only during data import.

Value range: **true**, **on**, **false**, and **off**.

- If this parameter is set to **true** or **on** and the number of source data files exceeds the number of foreign table columns, excessive columns will be ignored.
- When the parameter is **false** or **off**, and the number of data source files is more than the number of foreign table columns, the following error information will be displayed:  
extra data after last expected column

Default value: **false**

#### NOTICE

If a newline character at the end of a row is missing and the row and another row are integrated into one, data in another row is ignored after the parameter is set to **true**.

#### – COMPATIBLE\_ILLEGAL\_CHARS

Specifies whether to tolerate invalid characters during data import. The parameter is valid only for data import using COPY FROM.

Value range: **true**, **on**, **false**, and **off**.

- If this parameter is set to **true** or **on**, invalid characters are tolerated and imported to the database after conversion.
- If this parameter is set to **false** or **off** and an error occurs when there are invalid characters, the import will be interrupted.

Default value: **false** or **off**.

#### NOTE

The rules for converting invalid characters are as follows:

1. **\0** is converted to a space.
2. Other invalid characters are converted to question marks.
3. When **compatible\_illegal\_chars** is set to **true** or **on**, after invalid characters such as **NULL**, **DELIMITER**, **QUOTE**, and **ESCAPE** are converted to spaces or question marks, an error message like "illegal chars conversion may confuse COPY escape 0x20" will be displayed to remind you of possible parameter confusion caused by the conversion.

#### – FILL\_MISSING\_FIELDS

Specifies how to handle the problem that the last column of a row in a source data file is lost during data import.

Value range: **true**, **on**, **false**, and **off**.

Default value: **false** or **off**.

#### – DATE\_FORMAT

Specifies the DATE format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify

bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.

Value range: a valid DATE value. For details, see [Date and Time Processing Functions and Operators](#).

#### NOTE

You can use the **TIMESTAMP\_FORMAT** parameter to set the DATE format to **TIMESTAMP** for data import. For details, see **TIMESTAMP\_FORMAT** below.

- **TIME\_FORMAT**  
Specifies the TIME format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid TIME value. Time zones cannot be used. For details, see [Date and Time Processing Functions and Operators](#).
- **TIMESTAMP\_FORMAT**  
Specifies the TIMESTAMP format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid TIMESTAMP value. Time zones cannot be used. For details, see [Date and Time Processing Functions and Operators](#).
- **SMALLDATETIME\_FORMAT**  
Specifies the SMALLDATETIME format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid SMALLDATETIME value. For details, see [Date and Time Processing Functions and Operators](#).
- **COPY\_OPTION { option\_name ' value ' }**  
Specifies all types of native parameters of COPY.
  - **NULL null\_string**  
Specifies the string that represents a null value.

---

#### NOTICE

When using COPY FROM, any data item that matches this string will be stored as a null value, so make sure that you use the same string as you used with COPY TO.

---

Value range:

- A null value cannot be **\r** or **\n**. The maximum length is 100 characters.
- A null value cannot be the same as the **delimiter** or **quote** value.

Default value:

- The default value for the TEXT format is `\N`.
- The default value for the CSV format is an empty string without quotation marks.
- HEADER  
Specifies whether a file contains a header with the names of each column in the file. **header** is available only for CSV and FIXED files.  
When data is imported, if **header** is **on**, the first row of the data file will be identified as the header and ignored. If **header** is **off**, the first row will be identified as a data row.  
When data is exported, if header is **on**, **fileheader** must be specified. If **header** is **off**, an exported file does not contain a header.
- HEADER  
Does not report an error for `\.` in the imported data.
- FILEHEADER  
Specifies a file that defines the content in the header for exported data. The file contains data description of each column.

---

**NOTICE**

- This parameter is available only when **header** is **on** or **true**.
  - **fileheader** specifies an absolute path.
  - The file can contain only one row of header information, and ends with a newline character. Excess rows will be discarded. (Header information cannot contain newline characters.)
  - The length of the file including the newline character cannot exceed 1 MB.
- 
- FREEZE  
Sets the **COPY** loaded data row as **frozen**, like these data have executed **VACUUM FREEZE**.  
This is a performance option of initial data loading. The data will be frozen only when the following three requirements are met:
    - The table being loaded has been created or truncated in the same transaction before copying.
    - There are no cursors open in the current transaction.
    - There are no original snapshots in the current transaction.

 **NOTE**

When COPY is completed, all the other sessions will see the data immediately. However, this violates the general principle of MVCC visibility, and users should understand that this may cause potential risks.

- FORCE NOT NULL column\_name [, ...]  
In **CSV COPY FROM** mode, the specified column is not null. If the column is null, its value is regarded as a string of 0 characters.

Value range: name of an existing column.

- FORCE QUOTE { column\_name [, ...] | \* }

In **CSV COPY TO** mode, forces quotation marks to be used for all non-null values in each specified column. **NULL** values are not quoted.

Value range: name of an existing column.

- BINARY

Specifies that data is stored and read in binary mode instead of text mode. In binary mode, you cannot declare **DELIMITER**, **NULL**, or **CSV**. When **BINARY** is specified, **CSV**, **FIXED**, and **TEXT** cannot be specified through **option** or **copy\_option**.

- CSV

Enables the CSV mode. When **CSV** is specified, **BINARY**, **FIXED**, and **TEXT** cannot be specified through **option** or **copy\_option**.

- QUOTE [AS] 'quote\_character'

Specifies a quoted character string for a CSV file.

Default value: double quotation marks ("").

#### NOTE

- The value of **quote** cannot be the same as that of the **delimiter** or **null** parameter.
- The value of **quote** must be a single-byte character.
- You are advised to set **quote** to an invisible character, such as **0x07**, **0x08**, or **0x1b**.

- ESCAPE [AS] 'escape\_character'

Specifies an escape character for a CSV file. The value must be a single-byte character.

Default value: double quotation marks (""). If the value is the same as that of **quote**, it will be replaced by **\0**.

- EOL 'newline\_character'

Specifies the newline character style of the imported or exported data file.

Value range: multi-character newline characters within 10 bytes. Common newline characters include **\r** (0x0D), **\n** (0x0A), and **\r\n** (0x0D0A). Special newline characters include **\$** and **#**.

#### NOTE

- The **EOL** parameter supports only the TEXT format for data import and export and does not support the CSV or FIXED format. For forward compatibility, the **EOL** parameter can be set to **0x0D** or **0x0D0A** for data export in the CSV or FIXED format.
- The value of **EOL** cannot be the same as that of the **delimiter** or **null** parameter.
- The EOL parameter value cannot contain the following characters: .abcdefghijklmnopqrstuvwxyz0123456789.

- ENCODING 'encoding\_name'

Specifies the name of a file encoding format.

Value range: a valid encoding format.

Default value: current encoding format.

– IGNORE\_EXTRA\_DATA

Specifies that when the number of data source files exceeds the number of foreign table columns, excess columns at the end of the row are ignored. This parameter is used only during data import.

If this parameter is not used and the number of columns in the data source file is greater than that defined in the foreign table, the following error information is displayed:

extra data after last expected column

– COMPATIBLE\_ILLEGAL\_CHARS

Specifies that invalid characters are tolerated during data import. Invalid characters are converted and then imported to the database. No error is reported and the import is not interrupted. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.

If this parameter is not used, an error is reported when invalid characters are encountered during the import, and the import is interrupted.

 NOTE

The rules for converting invalid characters are as follows:

1. \0 is converted to a space.
2. Other invalid characters are converted to question marks.
3. When **compatible\_illegal\_chars** is set to **true** or **on**, after invalid characters such as **NULL**, **DELIMITER**, **QUOTE**, and **ESCAPE** are converted to spaces or question marks, an error message like "illegal chars conversion may confuse COPY escape 0x20" will be displayed to remind you of possible parameter confusion caused by the conversion.

– FILL\_MISSING\_FIELDS [ { 'one' | 'multi' } ]

Specifies how to handle the problem that the last columns of a row in a source data file are lost during data import. If **one** or **multi** is not specified or **one** is specified, the missing of the last column is handled in the default mode. If **multi** is specified, the missing of the last multiple columns are handled in the default mode.

Value range: **true**, **on**, **false**, and **off**.

Default value: **false** or **off**.

---

NOTICE

Do not specify this option. Currently, it does not enable error tolerance, but will make the parser ignore the said errors during data parsing on the primary node of the database. Such errors will not be recorded in the COPY error table (enabled by using **LOG ERRORS REJECT LIMIT**) but will be reported later by database node. Therefore, do not specify this option.

---

– DATE\_FORMAT 'date\_format\_string'

Specifies the DATE format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify

bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.

Value range: a valid DATE value. For details, see [Date and Time Processing Functions and Operators](#).

#### NOTE

You can use the **TIMESTAMP\_FORMAT** parameter to set the DATE format to **TIMESTAMP** for data import. For details, see **TIMESTAMP\_FORMAT** below.

- **TIME\_FORMAT** 'time\_format\_string'  
Specifies the TIME format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid TIME value. Time zones cannot be used. For details, see [Date and Time Processing Functions and Operators](#).
- **TIMESTAMP\_FORMAT** 'timestamp\_format\_string'  
Specifies the TIMESTAMP format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid TIMESTAMP value. Time zones cannot be used. For details, see [Date and Time Processing Functions and Operators](#).
- **SMALLDATETIME\_FORMAT** 'smalldatetime\_format\_string'  
Specifies the SMALLDATETIME format for data import. The BINARY format is not supported. When data of such format is imported, error "cannot specify bulkload compatibility options in BINARY mode" will occur. The parameter is valid only for data import using COPY FROM.  
Value range: a valid SMALLDATETIME value. For details, see [Date and Time Processing Functions and Operators](#).
- **TRANSFORM** ( { column\_name [ data\_type ] [ AS transform\_expr ] } [, ...] )  
Specify the conversion expression of each column in the table. **data\_type** specifies the data type of the column in the expression parameter. **transform\_expr** is the target expression that returns the result value whose data type is the same as that of the target column in the table. For details about the expression, see [Expressions](#).
- **SKIP** int\_number  
Specifies that the first *int\_number* rows of the data file are skipped during data import.
- **WHEN** { ( start - end ) | column\_name } { = | != } 'string'  
When data is imported, each row of data is checked. Only the rows that meet the WHEN condition are imported to the table.
- **SEQUENCE** ( { column\_name ( integer [, incr] ) [, ...] } )  
During data import, columns modified by SEQUENCE do not read data from the data file. The values are incremented by the value of **incr** based on the specified integer. If **incr** is not specified, the values are incremented from 1 by default.



- FILLER ( { column\_name [, ...] } )

When data is imported, the column modified by FILLER is discarded after being read from the data file.

#### NOTE

To use FILLER, you need to specify the list of columns to be copied. During data processing, data is processed based on the position of the **filler** column in the column list.

- CONSTANT ( { column\_name 'constant\_string' [, ...] } )

When data is imported, the column modified by CONSTANT is not read from the data file, and **constant\_string** is used to assign a value to the column.

The following special backslash sequences are recognized by COPY FROM:

- **\b**: Backslash (ASCII 8)
- **\f**: Form feed (ASCII 12)
- **\n**: Newline character (ASCII 10)
- **\r**: Carriage return character (ASCII 13)
- **\t**: Tab (ASCII 9)
- **\v**: Vertical tab (ASCII 11)
- **\digits**: Backslash followed by one to three octal digits specifies that the ASCII value is the character with that numeric code.
- **\xdigits**: Backslash followed by an x and one or two hex digits specifies the character with that numeric code.

## Permission Control Examples

```
openGauss=> copy t1 from '/home/xy/t1.csv';
ERROR: COPY to or from a file is prohibited for security concerns
HINT: Anyone can COPY to stdout or from stdin. gsql's \copy command also works for anyone.
openGauss=> grant gs_role_copy_files to xxx;
```

This error occurs because a non-initial user does not have the COPY permission. To solve this problem, set the **enable\_copy\_server\_files** parameter to **on**. Then, administrators can use the COPY function. To perform the COPY operation, common users must also join the **gs\_role\_copy\_files** group.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.ship_mode table.
openGauss=# CREATE TABLE tpcds.ship_mode
(
 SM_SHIP_MODE_SK INTEGER NOT NULL,
 SM_SHIP_MODE_ID CHAR(16) NOT NULL,
 SM_TYPE CHAR(30) ,
 SM_CODE CHAR(10) ,
 SM_CARRIER CHAR(20) ,
 SM_CONTRACT CHAR(20)
);

-- Insert a single data record into the tpcds.ship_mode table.
openGauss=# INSERT INTO tpcds.ship_mode VALUES (1,'a','b','c','d','e');
```

```
-- Copy data from the tpcds.ship_mode file to the /home/omm/ds_ship_mode.dat file.
openGauss=# COPY tpcds.ship_mode TO '/home/omm/ds_ship_mode.dat';

-- Output tpcds.ship_mode to STDOUT.
openGauss=# COPY tpcds.ship_mode TO STDOUT;

-- Output the data of tpcds.ship_mode to STDOUT. The parameters are as follows: The delimiter is ','
(delimiter ',') and the encoding format is UTF8 (encoding 'utf8').
openGauss=# COPY tpcds.ship_mode TO STDOUT WITH (delimiter ',', encoding 'utf8');

-- Output the data of tpcds.ship_mode to STDOUT. The parameters are as follows: The import format is
CSV (format 'CSV'), and the exported content of the SM_SHIP_MODE_SK column is enclosed in quotation
marks (force_quote(SM_SHIP_MODE_SK)).
openGauss=# COPY tpcds.ship_mode TO STDOUT WITH (format 'CSV', force_quote(SM_SHIP_MODE_SK));

-- Create the tpcds.ship_mode_t1 table.
openGauss=# CREATE TABLE tpcds.ship_mode_t1
(
 SM_SHIP_MODE_SK INTEGER NOT NULL,
 SM_SHIP_MODE_ID CHAR(16) NOT NULL,
 SM_TYPE CHAR(30) ,
 SM_CODE CHAR(10) ,
 SM_CARRIER CHAR(20) ,
 SM_CONTRACT CHAR(20)
)
;

-- Copy data from STDIN to the tpcds.ship_mode_t1 table.
openGauss=# COPY tpcds.ship_mode_t1 FROM STDIN;

-- Copy data from the /home/omm/ds_ship_mode.dat file to the tpcds.ship_mode_t1 table.
openGauss=# COPY tpcds.ship_mode_t1 FROM '/home/omm/ds_ship_mode.dat';

-- Copy data from the /home/omm/ds_ship_mode.dat file to the tpcds.ship_mode_t1 table, convert the
data using the TRANSFORM expression, and insert the 10 characters on the left of the SM_TYPE column
into the table.
openGauss=# COPY tpcds.ship_mode_t1 FROM '/home/omm/ds_ship_mode.dat' TRANSFORM (SM_TYPE AS
LEFT(SM_TYPE, 10));

-- Copy data from the /home/omm/ds_ship_mode.dat file to the tpcds.ship_mode_t1 table, with the
import format set to TEXT (format 'text'), the delimiter set to '\t' (delimiter E'\t'), excessive columns
ignored (ignore_extra_data 'true'), and characters not escaped (noescaping 'true').
openGauss=# COPY tpcds.ship_mode_t1 FROM '/home/omm/ds_ship_mode.dat' WITH(format 'text',
delimiter E'\t', ignore_extra_data 'true', noescaping 'true');

-- Copy data from the /home/omm/ds_ship_mode.dat file to the tpcds.ship_mode_t1 table, with the
import format set to FIXED, fixed-length format specified (FORMATTER(SM_SHIP_MODE_SK(0, 2),
SM_SHIP_MODE_ID(2,16), SM_TYPE(18,30), SM_CODE(50,10), SM_CARRIER(61,20),
SM_CONTRACT(82,20))), excessive columns ignored (ignore_extra_data), and headers included (header).
openGauss=# COPY tpcds.ship_mode_t1 FROM '/home/omm/ds_ship_mode.dat' FIXED
FORMATTER(SM_SHIP_MODE_SK(0, 2), SM_SHIP_MODE_ID(2,16), SM_TYPE(18,30), SM_CODE(50,10),
SM_CARRIER(61,20), SM_CONTRACT(82,20)) header ignore_extra_data;

-- Drop tables and the schema.
openGauss=# DROP TABLE tpcds.ship_mode;
openGauss=# DROP TABLE tpcds.ship_mode_t1;
openGauss=# DROP SCHEMA tpcds;
```

## 7.13.46 CREATE AGGREGATE

### Description

Defines a new aggregate function.

### Syntax

```
CREATE AGGREGATE name (input_data_type [, ...]) (
 SFUNC = sfunc,
```

```
SType = state_data_type
[, FINALFUNC = ffunc]
[, INITCOND = initial_condition]
[, SORTOP = sort_operator]
)

or the old syntax

CREATE AGGREGATE name (
 BASETYPE = base_type,
 SFUNC = sfunc,
 STYPE = state_data_type
 [, FINALFUNC = ffunc]
 [, INITCOND = initial_condition]
 [, SORTOP = sort_operator]
)
```

## Parameters

- **name**  
Name (optionally schema-qualified) of the aggregate function to be created.
- **input\_data\_type**  
Data type of the input to be processed by the aggregate function. To create a zero-parameter aggregate function, you can use an asterisk (\*) instead of a list of input data types. (count(\*) is an instance of this aggregate function.)
- **base\_type**  
In the CREATE AGGREGATE syntax, the input data type is specified by the **basetype** parameter instead of following **name**. Note that the previous syntax allows only one input parameter. To create a zero-parameter aggregate function, you can set **basetype** to **ANY** instead of \*.
- **sfunc**  
Name of the state conversion function that will be called on each input line. For an aggregate function with *N* parameters, **sfunc** must have *N*+1 parameters. The first parameter is of the **state\_data\_type** type, and the other parameters match the declared input data types. The function must return a value of the **state\_data\_type** type. This function accepts the current state value and the current input data, and returns the next state value.
- **state\_data\_type**  
Data type of the aggregation status value.
- **ffunc**  
Final processing function called after all the input lines have been converted, which calculates the result of aggregation. This function must accept a parameter of **state\_data\_type**. The output data type of the aggregation is defined as the return type of this function. If **ffunc** is not specified, the state value of the aggregation result is used as the aggregation result, and the output type is **state\_data\_type**.
- **initial\_condition**  
Initial setting (value) of a state value. It must be a text constant value acceptable to **state\_data\_type**. If not specified, the initial state value is **NULL**.
- **sort\_operator**  
Sort operator used for MIN or MAX aggregation. This is just an operator name (optionally schema-qualified). This operator assumes that the input data type is the same as that of aggregation.

## Examples

```
-- Create a user-defined function.
openGauss=# CREATE OR REPLACE FUNCTION int_add(int,int)
returns int as $BODY$
declare
begin
return $1 + $2;
end;
$BODY$ language plpgsql;

-- Create an aggregate function.
openGauss=# CREATE AGGREGATE sum_add(int)
(
sfunc = int_add,

stype = int,

initcond = '0'
);

-- Create a test table and add data.
openGauss=# CREATE TABLE test_sum(a int,b int,c int);
openGauss=# INSERT INTO test_sum VALUES(1,2),(2,3),(3,4),(4,5);

-- Execute the aggregate function.
openGauss=# SELECT sum_add(a) FROM test_sum;
sum_add

10

-- Drop the aggregate function.
openGauss=# DROP AGGREGATE sum_add(int);

-- Drop the user-defined function.
openGauss=# DROP FUNCTION int_add(int,int);

-- Drop the test table.
openGauss=# DROP TABLE test_sum;
```

## 7.13.47 CREATE AUDIT POLICY

### Description

Creates a unified audit policy.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

The masking policy takes effect only after the security policy is enabled, that is, **enable\_security\_policy** is set to **on**.

### Syntax

```
CREATE AUDIT POLICY [IF NOT EXISTS] policy_name { privilege_audit_clause | access_audit_clause } [, ...]
[filter_group_clause] [ENABLE | DISABLE];
```

- **privilege\_audit\_clause**  
PRIVILEGES { DDL | ALL } [ ON LABEL ( resource\_label\_name [, ... ] ) ]
- **access\_audit\_clause**  
ACCESS { DML | ALL } [ ON LABEL ( resource\_label\_name [, ... ] ) ]

- **filter\_group\_clause**  
FILTER ON { FILTER\_TYPE ( filter\_value [, ... ] ) } [, ... ]

## Parameters

- **policy\_name**  
Specifies the audit policy name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **resource\_label\_name**  
Specifies the resource label name.
- **DDL**  
Specifies the operations that are audited in the database: **CREATE**, **ALTER**, **DROP**, **ANALYZE**, **COMMENT**, **GRANT**, **REVOKE**, **SET**, **SHOW**, **LOGIN\_ANY**, **LOGIN\_FAILURE**, **LOGIN\_SUCCESS**, and **LOGOUT**.
- **DML**  
Specifies the operations that are audited in the database: **SELECT**, **COPY**, **DEALLOCATE**, **DELETE**, **EXECUTE**, **INSERT**, **PREPARE**, **REINDEX**, **TRUNCATE**, and **UPDATE**.
- **ALL**  
Specifies all operations supported by the specified DDL or DML statements in the database. When the form is { DDL | ALL }, **ALL** indicates all DDL operations. When the form is { DML | ALL }, **ALL** indicates all DML operations.
- **FILTER\_TYPE**  
Specifies the types of information to be filtered by the policy, including **APP**, **ROLES**, and **IP**.
- **filter\_value**  
Specifies the detailed information to be filtered.
- **ENABLE|DISABLE**  
Enables or disables the unified audit policy. If **ENABLE|DISABLE** is not specified, **ENABLE** is used by default.

## Examples

```
-- Create users dev_audit and bob_audit.
openGauss=# CREATE USER dev_audit PASSWORD '*****';
openGauss=# CREATE USER bob_audit PASSWORD '*****';

-- Create table tb_for_audit.
openGauss=# CREATE TABLE tb_for_audit(col1 text, col2 text, col3 text);

-- Create a resource label.
openGauss=# CREATE RESOURCE LABEL adt_lb0 ADD TABLE(tb_for_audit);

-- Perform the CREATE operation on the database to create an audit policy.
openGauss=# CREATE AUDIT POLICY adt1 PRIVILEGES CREATE;

-- Perform the SELECT operation on the database to create an audit policy.
openGauss=# CREATE AUDIT POLICY adt2 ACCESS SELECT;

-- Create an audit policy to audit only the CREATE operations performed on the adt_lb0 resource by users dev_audit and bob_audit.
openGauss=# CREATE AUDIT POLICY adt3 PRIVILEGES CREATE ON LABEL(adt_lb0) FILTER ON
ROLES(dev_audit, bob_audit);
```

```
-- Create an audit policy to audit only the SELECT, INSERT, and DELETE operations performed on the
adt_lb0 resource by users dev_audit and bob_audit using client tool gsql on the servers whose IP
addresses are 10.20.30.40 and 127.0.0.0/24.
openGauss=# CREATE AUDIT POLICY adt4 ACCESS SELECT ON LABEL(adt_lb0), INSERT ON LABEL(adt_lb0),
DELETE FILTER ON ROLES(dev_audit, bob_audit), APP(gsql), IP('10.20.30.40', '127.0.0.0/24');

-- Drop audit policies.
openGauss=# DROP AUDIT POLICY adt1, adt2, adt3, adt4;

-- Drop a resource label.
openGauss=# DROP RESOURCE LABEL adt_lb0;

-- Drop the tb_for_audit table.
openGauss=# DROP TABLE tb_for_audit;

-- Drop the dev_audit and bob_audit users.
openGauss=# DROP USER dev_audit, bob_audit;
```

## Helpful Links

[ALTER AUDIT POLICY](#) and [DROP AUDIT POLICY](#)

## 7.13.48 CREATE CAST

### Description

Defines a conversion.

### Syntax

```
CREATE CAST (source_type AS target_type)
 WITH FUNCTION function_name (argument_type [, ...])
 [AS ASSIGNMENT | AS IMPLICIT]

CREATE CAST (source_type AS target_type)
 WITHOUT FUNCTION
 [AS ASSIGNMENT | AS IMPLICIT]

CREATE CAST (source_type AS target_type)
 WITH INOUT
 [AS ASSIGNMENT | AS IMPLICIT]
```

### Parameters

- **source\_type**  
Type of the source data to be converted.
- **target\_type**  
Type of the target data to be converted.
- **function\_name(argument\_type [, ...])**  
Function used for conversion. The function name can be modified with a schema name. If it is not modified with a schema name, the function will be found in the schema search path. The result data type of the function must match the target type of the conversion. Its parameters are discussed below.
- **WITHOUT FUNCTION**  
Indicates that the source type is a binary castable to the target type, so no function is needed to perform this conversion.
- **WITH INOUT**

Indicates that the conversion is an I/O conversion, which is performed by calling the output function of the source data type and transferring the result to the input function of the target data type.

- **AS ASSIGNMENT**

Indicates that the conversion can be implicitly called in assignment mode.

- **AS IMPLICIT**

Indicates that the transformation can be implicitly called in any environment.

A conversion implementation function can have one to three parameters. The type of the first parameter must be the same as that of the source type to be converted, or can be forcibly converted from the binary of the source type to be converted. If the second parameter exists, it must be of the integer type. It receives these type modifiers associated with the target type, or **-1** if nothing is present. If the third parameter exists, it must be of the Boolean type. If the conversion is an explicit type conversion, **true** is received. Otherwise, **false** is received.

The return type of a conversion function must be the same as the target type of the conversion, or the two types can be binary coercible.

Typically, a transformation must have different source and target data types. However, if there is a conversion implementation function with more than one parameter, it is allowed to declare a conversion with the same source and target types. This is used to represent a length enforcement function of a specific type in the system catalog. The named function is used to force a value of this type to be the type modifier value given by the second parameter.

If the source type and target type of a type conversion are different and more than one parameter is received, it indicates that only one step is required to convert one type to another and the length conversion is performed at the same time. If no such conversion is available, converting to a type that uses a type modifier involves two steps, one to convert between data types, and the other to apply a conversion specified by the modifier.

Currently, domain type conversion does not take effect. Transformations are typically targeted to the domain-related data types to which they belong.

#### NOTE

Cast is executed depending on the permission of the user who calls it. When calling a cast created by others, check the execution content of the cast function to prevent the cast creator from performing unauthorized operations with the permission of the executor.

## Examples

To create an assignment mapping from type bigint to type int4, use the int4(bigint) function:

```
CREATE CAST (bigint AS int4) WITH FUNCTION int4(bigint) AS ASSIGNMENT;
```

(The conversion has been predefined in the system.)

## Compatibility

The CREATE CAST instruction complies with the SQL standard. Except that the SQL does not have extra parameters that can be forcibly converted to binary types or implement functions.

## Helpful Links

[DROP CAST](#)

## 7.13.49 CREATE CONVERSION

### Description

Defines a new conversion between two character set encodings.

### Precautions

- The **DEFAULT** parameter indicates that the conversion between the source encoding and the target encoding is executed by default between the client and the server. To support this usage, bidirectional conversion must be defined, that is, both conversion from A to B and conversion from B to A are supported.
- To perform conversion, you must have the EXECUTE permission on function and the CREATE permission on the target schema.
- SQL\_ASCII cannot be used for either source encoding or target encoding because the server behavior is hardwired when SQL\_ASCII "encoding" is involved.
- You can remove user-defined conversions using DROP CONVERSION.

### Syntax

```
CREATE [DEFAULT] CONVERSION name
FOR source_encoding TO dest_encoding FROM function_name
```

### Parameters

- **DEFAULT**  
Specifies that the conversion is the default conversion from the source encoding to the target encoding. There should be only one default conversion for each encoding pair in a schema.
- **name**  
Specifies the name of the conversion, which can be restricted by the schema. If not restricted by a schema, the conversion is defined in the current schema. The conversion name must be unique in a schema.
- **source\_encoding**  
Source encoding name.
- **dest\_encoding**  
Target encoding name.
- **function\_name**



Function used for conversion. A function name can be restricted by a schema. If not, the function is found in the path.

```
conv_proc(
 integer, -- Source encoding ID
 integer, -- Target encoding ID
 cstring, -- Source character string (C character string ending with a null value)
 internal, -- Target (filled with a null-terminated C character string)
 integer -- Length of the source string
) RETURNS void;
```

## Examples

```
-- Use myfunc to create an encoding conversion from UTF8 to LATIN1.
CREATE CONVERSION myconv FOR 'UTF8' TO 'LATIN1' FROM myfunc;
```

## 7.13.50 CREATE DATABASE

### Description

Creates a database. By default, the new database will be created only by cloning the standard system database **template0**.

### Precautions

- A user who has the CREATEDB permission or a system administrator can create a database.
- It cannot be executed inside a transaction block.
- During the database creation, an error message indicating that permission denied is displayed, possibly because the permission on the data directory in the file system is insufficient. If an error message, indicating no space left on device is displayed, the possible cause is that the disk space is used up.

### Syntax

```
CREATE DATABASE database_name
[[WITH] { [OWNER [=] user_name] |
 [TEMPLATE [=] template] |
 [ENCODING [=] 'encoding'] |
 [LC_COLLATE [=] 'lc_collate'] |
 [LC_CTYPE [=] 'lc_ctype'] |
 [DBCOMPATIBILITY [=] 'compatibilty_type'] |
 [TABLESPACE [=] tablespace_name] |
 [CONNECTION LIMIT [=] connlimit] }[...] ;
```

### Parameters

- **database\_name**  
Specifies the database name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **OWNER [=] user\_name**  
Specifies the owner of the new database. If omitted, the default owner is the current user.  
Value range: an existing username.
- **TEMPLATE [=] template**  
Specifies a template name. That is, the template from which the database is created. GaussDB creates a database by copying data from a template

database. GaussDB has two default template databases **template0** and **template1** and a default user database **postgres**.

Value range: **template0**.

- **ENCODING [ = ] 'encoding'**

Specifies the character encoding used by the database. The value can be a string (for example, **SQL\_ASCII**) or an integer.

By default, the encoding format of the template database is used, and the codes of the template databases **template0** and **template1** are related to the OS environment. The character encoding of **template1** cannot be changed. To change the encoding, use **template0** to create a database.

Common values are **GBK**, **UTF8**, **Latin1**, and **GB18030**. The supported character sets are as follows:

**Table 7-104** GaussDB character set

| Name         | Description                       | Language            | Server-side Encoding | ICU Support | Number of Bytes/Characters | Alias              |
|--------------|-----------------------------------|---------------------|----------------------|-------------|----------------------------|--------------------|
| BIG5         | Big Five                          | Traditional Chinese | No                   | No          | 1-2                        | WIN950, Windows950 |
| EUC_CN       | Extended Unix Code-CN             | Simplified Chinese  | Yes                  | Yes         | 1-3                        | -                  |
| EUC_JP       | Extended Unix Code-JP             | Japanese            | Yes                  | Yes         | 1-3                        | -                  |
| EUC_JIS_2004 | Extended Unix Code-JP, JIS X 0213 | Japanese            | Yes                  | No          | 1-3                        | -                  |
| EUC_KR       | Extended Unix Code-KR             | Korean              | Yes                  | Yes         | 1-3                        | -                  |
| EUC_TW       | Extended Unix Code-Taiwan, China  | Traditional Chinese | Yes                  | Yes         | 1-3                        | -                  |
| GB18030      | National standards                | Chinese             | Yes                  | No          | 1-4                        | -                  |

| Name       | Description                 | Language                   | Server-side Encoding | ICU Support | Number of Bytes/Characters | Alias              |
|------------|-----------------------------|----------------------------|----------------------|-------------|----------------------------|--------------------|
| GBK        | Extended national standards | Simplified Chinese         | Yes                  | No          | 1-2                        | WIN936, Windows936 |
| ISO_8859_5 | ISO 8859-5, ECMA 113        | Latin/Cyrillic             | Yes                  | Yes         | 1                          | -                  |
| ISO_8859_6 | ISO 8859-6, ECMA 114        | Latin/Arabic               | Yes                  | Yes         | 1                          | -                  |
| ISO_8859_7 | ISO 8859-7, ECMA 118        | Latin/Greek                | Yes                  | Yes         | 1                          | -                  |
| ISO_8859_8 | ISO 8859-8, ECMA 121        | Latin/Hebrew               | Yes                  | Yes         | 1                          | -                  |
| JOHAB      | JOHAB                       | Korean                     | No                   | No          | 1-3                        | -                  |
| KOI8R      | KOI8-R                      | Cyrillic (Russian)         | Yes                  | Yes         | 1                          | KOI8               |
| KOI8U      | KOI8-U                      | Cyrillic (Ukrainian)       | Yes                  | Yes         | 1                          | -                  |
| LATIN1     | ISO 8859-1, ECMA 94         | Western European languages | Yes                  | Yes         | 1                          | ISO88591           |
| LATIN2     | ISO 8859-2, ECMA 94         | Central European languages | Yes                  | Yes         | 1                          | ISO88592           |

| Name          | Description                | Language                     | Server-side Encoding | ICU Support | Number of Bytes/Characters | Alias                                 |
|---------------|----------------------------|------------------------------|----------------------|-------------|----------------------------|---------------------------------------|
| LATIN3        | ISO 8859-3, ECMA 94        | South European languages     | Yes                  | Yes         | 1                          | ISO88593                              |
| LATIN4        | ISO 8859-4, ECMA 94        | North European languages     | Yes                  | Yes         | 1                          | ISO88594                              |
| LATIN5        | ISO 8859-9, ECMA 128       | Turkish                      | Yes                  | Yes         | 1                          | ISO88599                              |
| LATIN6        | ISO 8859-10, ECMA 144      | Germanic languages           | Yes                  | Yes         | 1                          | ISO885910                             |
| LATIN7        | ISO 8859-13                | Baltic languages             | Yes                  | Yes         | 1                          | ISO885913                             |
| LATIN8        | ISO 8859-14                | Celtic languages             | Yes                  | Yes         | 1                          | ISO885914                             |
| LATIN9        | ISO 8859-15                | LATIN1 with Euro and accents | Yes                  | Yes         | 1                          | ISO885915                             |
| LATIN10       | ISO 8859-16, ASRO SR 14111 | Romanian                     | Yes                  | No          | 1                          | ISO885916                             |
| MULE_INTERNAL | Mule internal code         | Multilingual Emacs           | Yes                  | No          | 1-4                        | -                                     |
| SJIS          | Shift JIS                  | Japanese                     | No                   | No          | 1-2                        | Mskanji, ShiftJIS, WIN932, Windows932 |

| Name           | Description                | Language                   | Server-side Encoding | ICU Support | Number of Bytes/Characters | Alias              |
|----------------|----------------------------|----------------------------|----------------------|-------------|----------------------------|--------------------|
| SHIFT_JIS_2004 | Shift JIS, JIS X 0213      | Japanese                   | No                   | No          | 1-2                        | -                  |
| SQL_ASCII      | Unspecified (see the text) | Any                        | Yes                  | No          | 1                          | -                  |
| UHC            | Unified Hangul Code        | Korean                     | No                   | No          | 1-2                        | WIN949, Windows949 |
| UTF8           | Unicode, 8-bit             | All                        | Yes                  | Yes         | 1-4                        | Unicode            |
| WIN866         | Windows CP866              | Cyrillic                   | Yes                  | Yes         | 1                          | ALT                |
| WIN874         | Windows CP874              | Thai                       | Yes                  | No          | 1                          | -                  |
| WIN1250        | Windows CP1250             | Central European languages | Yes                  | Yes         | 1                          | -                  |
| WIN1251        | Windows CP1251             | Cyrillic                   | Yes                  | Yes         | 1                          | WIN                |
| WIN1252        | Windows CP1252             | Western European languages | Yes                  | Yes         | 1                          | -                  |
| WIN1253        | Windows CP1253             | Greek                      | Yes                  | Yes         | 1                          | -                  |
| WIN1254        | Windows CP1254             | Turkish                    | Yes                  | Yes         | 1                          | -                  |
| WIN1255        | Windows CP1255             | Hebrew                     | Yes                  | Yes         | 1                          | -                  |

| Name    | Description    | Language         | Server-side Encoding | ICU Support | Number of Bytes/Characters | Alias                      |
|---------|----------------|------------------|----------------------|-------------|----------------------------|----------------------------|
| WIN1256 | Windows CP1256 | Arabic           | Yes                  | Yes         | 1                          | -                          |
| WIN1257 | Windows CP1257 | Baltic languages | Yes                  | Yes         | 1                          | -                          |
| WIN1258 | Windows CP1258 | Vietnamese       | Yes                  | Yes         | 1                          | ABC, TCVN, TCVN5712, VSCII |

---

**NOTICE**

Note that not all client APIs support the preceding character sets.

The SQL\_ASCII setting is different from other settings. If the character set of the server is SQL\_ASCII, the server interprets the byte values 0 to 127 according to the ASCII standard. The byte values 128 to 255 are regarded as the characters that cannot be parsed. If this parameter is set to SQL\_ASCII, no code conversion occurs. Therefore, this setting is not basically used to declare the specified encoding used, because this declaration ignores the encoding. In most cases, if you use any non-ASCII data, do not use the SQL\_ASCII setting because the database will not be able to convert or verify non-ASCII characters.

---

**NOTICE**

- The character set encoding of the new database must be compatible with the local settings (**LC\_COLLATE** and **LC\_CTYPE**).
- When the specified character encoding set is **GBK**, some uncommon Chinese characters cannot be directly used as object names. This is because the byte encoding overlaps with the ASCII characters @A-Z[\]^\_`a-z{[] when the second byte of the GBK ranges from 0x40 to 0x7E. @[\\]^\_{'[] is an operator in the database. If it is directly used as an object name, a syntax error will be reported. For example, the GBK hexadecimal code is **0x8240**, and the second byte is **0x40**, which is the same as the ASCII character @. Therefore, the character cannot be used as an object name. If you need to use this function, you can add double quotation marks ("" ) to avoid this problem when creating and accessing objects.
- If the client code is A and the server code is B, the conversion between encoding formats A and B must exist in the database. For details about encoding format conversion supported by the database, see the system catalog **PG\_CONVERSION**. (If the encoding format cannot be converted, it is recommended that the encoding format on the client be the same as that on the server. You can change the encoding format on the client by setting the GUC parameter **client\_encoding**.)

- **LC\_COLLATE [ = ] 'lc\_collate'**

Specifies the character set used by the new database. For example, set this parameter by using **lc\_collate = 'zh\_CN.gbk'**.

The use of this parameter affects the sort order of strings (for example, the order of using **ORDER BY** for execution and the order of using indexes on text columns). By default, the sorting order of the template database is used.

Value range: character sets supported by the OS.

- **LC\_CTYPE [ = ] 'lc\_ctype'**

Specifies the character class used by the new database. For example, set this parameter by using **lc\_ctype = 'zh\_CN.gbk'**. The use of this parameter affects the classification of characters, such as uppercase letters, lowercase letters, and digits. By default, the character classification of the template database is used.

Value range: character classes supported by the OS.

 **NOTE**

The value ranges of **lc\_collate** and **lc\_ctype** depend on the character sets supported by the local environment. For example, in the Linux OS, you can run the **locale -a** command to obtain the list of character sets supported by the OS. When using the **lc\_collate** and **lc\_ctype** parameters, you can select the required character sets and character categories.

- **DBCOMPATIBILITY [ = ] 'compatibility\_type'**

Specifies the compatible database type. The default compatible database is the **O** database.

Value range: **A**, **B**, **C**, and **PG**, indicating the **O**, **MY**, **TD**, and **POSTGRES** databases, respectively.

 NOTE

- For A compatibility, the database treats empty strings as **NULL** and replaces **DATE** with **TIMESTAMP(0) WITHOUT TIME ZONE**.
  - When a character string is converted to an integer, if the input is invalid, the input will be converted to 0 due to B compatibility, and an error will be reported due to other compatibility issues.
  - For PG compatibility, CHAR and VARCHAR are counted by character. For other compatibility types, they are counted by byte. For example, for the UTF-8 character set, CHAR(3) can store three Chinese characters in PG compatibility scenarios, but can store only one Chinese character in other compatibility scenarios.
- **TABLESPACE [ = ] tablespace\_name**  
Specifies the tablespace of the database.  
Value range: an existing tablespace name.
  - **CONNECTION LIMIT [ = ] connlimit**  
Specifies the maximum number of concurrent connections that can be made to the new database.

---

**NOTICE**

- System administrators are not restricted by this parameter.
- The number of concurrent connections of each primary database node is calculated separately (which is the value of **connlimit**). Number of all connections of the database = Value of **connlimit** x Number of normal primary database nodes.

---

Value range: an integer in the range  $[-1, 2^{31} - 1]$ . The default value is **-1**, indicating that there is no limit.

The restrictions on character encoding are as follows:

- If the locale is set to **C** (or **POSIX**), all encoding types are allowed. For other locale settings, the character encoding must be the same as that of the locale.
- If the character encoding mode is SQL\_ASCII and the modifier is an administrator, the character encoding mode can be different from the locale setting.
- The encoding and region settings must match the template database, except that **template0** is used as a template. This is because other databases may contain data that does not match the specified encoding, or may contain indexes whose sorting order is affected by **LC\_COLLATE** and **LC\_CTYPE**. Copying this data will invalidate the indexes in the new database. **template0** does not contain any data or indexes that may be affected.

## Examples

```
-- Create users jim and tom.
openGauss=# CREATE USER jim PASSWORD '*****';
openGauss=# CREATE USER tom PASSWORD '*****';

-- Create database music using GBK (the local encoding type is also GBK).
openGauss=# CREATE DATABASE music ENCODING 'GBK' template = template0;

-- Create database music2 and specify user jim as its owner.
```



```
openGauss=# CREATE DATABASE music2 OWNER jim;

-- Create database music3 using template template0 and specify user jim as its owner.
openGauss=# CREATE DATABASE music3 OWNER jim TEMPLATE template0;

-- Set the maximum number of connections to database music to 10.
openGauss=# ALTER DATABASE music CONNECTION LIMIT= 10;

-- Rename database music to music4.
openGauss=# ALTER DATABASE music RENAME TO music4;

-- Change the owner of database music2 to user tom.
openGauss=# ALTER DATABASE music2 OWNER TO tom;

-- Set the tablespace of database music3 to PG_DEFAULT.
openGauss=# ALTER DATABASE music3 SET TABLESPACE PG_DEFAULT;

-- Disable the default index scan on database music3.
openGauss=# ALTER DATABASE music3 SET enable_indexscan TO off;

-- Reset the enable_indexscan parameter.
openGauss=# ALTER DATABASE music3 RESET enable_indexscan;

-- Drop databases.
openGauss=# DROP DATABASE music2;
openGauss=# DROP DATABASE music3;
openGauss=# DROP DATABASE music4;

-- Drop the jim and tom users.
openGauss=# DROP USER jim;
openGauss=# DROP USER tom;

-- Create a database compatible with the TD format.
openGauss=# CREATE DATABASE td_compatible_db DBCOMPATIBILITY 'C';

-- Create a database compatible with the A format.
openGauss=# CREATE DATABASE ora_compatible_db DBCOMPATIBILITY 'A';

-- Drop the databases that are compatible with the TD and A formats.
openGauss=# DROP DATABASE td_compatible_db;
openGauss=# DROP DATABASE ora_compatible_db;
```

## Helpful Links

[ALTER DATABASE](#) and [DROP DATABASE](#)

## Suggestions

- **create database**  
Database cannot be created in a transaction.
- **ENCODING LC\_COLLATE LC\_CTYPE**  
If the setting of **Encoding** of the new database does not match that of the template database (SQL\_ASCII) ('GBK', 'UTF8', 'LATIN1', or 'GB18030'), **template [=] template0** must be specified.

## 7.13.51 CREATE DIRECTORY

### Description

**CREATE DIRECTORY** creates a directory. The directory defines an alias for a path in the server file system and is used to store data files used by users.

## Precautions

- When **enable\_access\_server\_directory** is set to **off**, only the initial user is allowed to create directory objects. When **enable\_access\_server\_directory** is set to **on**, the user with the SYSADMIN permission and the user who inherits permissions of the built-in role **gs\_role\_directory\_create** can create directory objects.
- By default, the user who creates a directory has the read and write permissions on the directory.
- The default owner of a directory is the user who creates the directory.
- A directory cannot be created for the following paths:
  - The path contains special characters.
  - The path is a relative path.
  - The path is a symbolic link.
- The following validity check is performed during directory creation:
  - Check whether the path exists in the OS. If it does not exist, a message is displayed, indicating the potential risks.
  - Check whether the database initial user **omm** has the R/W/X permissions for the OS path. If the user does not have all the permissions, a message is displayed, indicating the potential risks.
- In the database, ensure that the path is the same on all the nodes. Otherwise, the path may fail to be found on some nodes when the directory is used.

## Syntax

```
CREATE [OR REPLACE] DIRECTORY directory_name
AS 'path_name';
```

## Parameters

- **directory\_name**  
Specifies the name of a directory.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **path\_name**  
Specifies the OS path for which a directory is to be created.  
Value range: a valid OS path

## Examples

```
-- Create a directory.
openGauss=# CREATE OR REPLACE DIRECTORY dir as '/tmp/';
```

## Helpful Links

[ALTER DIRECTORY](#) and [DROP DIRECTORY](#)

## 7.13.52 CREATE FUNCTION

### Description

Creates a function.

## Precautions

- If the parameters or return values of a function have precision, the precision is not checked.
- When creating a function, you are advised to explicitly specify the schemas of tables in the function definition. Otherwise, the function may fail to be executed.
- **current\_schema** and **search\_path** specified by **SET** during function creation are invalid. **search\_path** and **current\_schema** before and after function execution should be the same.
- If a function has output parameters, the SELECT statement uses the default values of the output parameters when calling the function. When the CALL statement calls the function, it requires that the output parameters must be specified. When the CALL statement calls an overloaded PACKAGE function, it can use the default values of the output parameters. For details, see examples in [CALL](#).
- Only the functions compatible with PostgreSQL or those with the PACKAGE attribute can be overloaded. After **REPLACE** is specified, a new function is created instead of replacing a function if the number of parameters, parameter type, or return value is different.
- You can use the SELECT statement to specify different parameters using identical functions. The syntax does not support calling identical functions without the PACKAGE attribute.
- When you create a function, you cannot insert other agg functions out of the avg function or other functions.
- By default, the permissions to execute new functions are granted to PUBLIC users. For details, see [GRANT](#). By default, a user inherits the permissions of the PUBLIC role. Therefore, the user has the permission to execute a function and view the definition of the function. In addition, to execute the function, the user must have the USAGE permission on the schema to which the function locates. When creating a function, the user can revoke the default execution permissions from **PUBLIC** and grant them to other users as needed. To avoid the time window during which new functions can be accessed by all users, create functions and set function execution permissions in a transaction. After the database object isolation attribute is enabled, common users can view only the definitions of functions that they have permission to execute.
- When calling functions without parameters inside another function, you can omit brackets and call functions using their names directly.
- When functions with output parameters are called inside another function which is an assignment expression, you can omit the output parameters of the called functions.
- Functions which are compatible to A-style database support viewing, exporting, and importing parameter comments.
- Functions which are compatible to A-style database support viewing, exporting, and importing comments between IS/AS and plsql\_body.
- Users granted with the CREATE ANY FUNCTION permission can create or replace functions in the user schemas.

- The default permission on a function is SECURITY INVOKER. To change the permission to SECURITY DEFINER, set the GUC parameter **behavior\_compat\_options** to '**plsql\_security\_definer**'.
- For PL/pgSQL functions, after **behavior\_compat\_options** is set to '**proc\_outparam\_override**', the behaviors of **out** and **inout** change. In the functions, **return**, **out**, and **inout** can be returned at the same time. Before the parameter is set to '**proc\_outparam\_override**', only **return** is returned. For details, see [Examples](#).
- Formal parameters cannot be overloaded if only the custom ref cursor type is different from the sys\_refcursor type.
- When an overloaded function is called, the variable type must be specified.
- Functions with the **OUT** parameter do not support nested calling. For example, you are advised to change **b := func(a,func(c,1));** to **tmp := func(c,1); b := func(a,tmp).**
- For PL/pgSQL functions, of **behavior\_compat\_options** is set to '**proc\_outparam\_override**', the restrictions are as follows:
  - a. If a function with the **out/inout** parameter already exists in the same schema or package, you cannot create another function with the same name with the **out/inout** parameter.
  - b. An **out** parameter must be added no matter whether the SELECT or CALL statement is used to call a stored procedure.
  - c. In some scenarios, functions cannot be used in expressions (compared with those before the parameter is enabled), for example, left assignment in a stored procedure and **call function**. For details, see [Examples](#).
  - d. Functions without return cannot be called. The perform function can be used to call functions.
  - e. When a function is called in a stored procedure, **out** or **inout** parameters cannot be set to constants. For details, see [Examples](#).
- If a function with the definer specified is created in a schema of another user, the function will be executed by another user, which may cause unauthorized operations. Therefore, exercise caution when performing this operation.

## Syntax

- Syntax (compatible with PostgreSQL) for creating a user-defined function:

```
CREATE [OR REPLACE] FUNCTION function_name
[([{ argname [argmode] argtype [{ DEFAULT | := | = } expression] } [, ...]])]
[RETURNS rettype [DETERMINISTIC] | RETURNS TABLE ({ column_name column_type }
[, ...])]
LANGUAGE lang_name
[
 { IMMUTABLE | STABLE | VOLATILE }
 | { SHIPPABLE | NOT SHIPPABLE }
 | WINDOW
 | [NOT] LEAKPROOF
 | { CALLED ON NULL INPUT | RETURNS NULL ON NULL INPUT | STRICT }
 | [{ EXTERNAL } SECURITY INVOKER | [EXTERNAL] SECURITY DEFINER | AUTHID DEFINER |
AUTHID CURRENT_USER]
 | { fenced | not fenced }
 | { PACKAGE }
 | COST execution_cost
 | ROWS result_rows
 | SET configuration_parameter { { TO | = } value | FROM CURRENT }
][...]
{
```

```
 AS 'definition'
 }
}

• A-style database syntax of creating a user-defined function:
CREATE [OR REPLACE] FUNCTION function_name
([{ argname [argmode] argtype [{ DEFAULT | := | = } expression] } [, ...]])
RETURN rettype [DETERMINISTIC]
[
 {IMMUTABLE | STABLE | VOLATILE }
 | {SHIPPABLE | NOT SHIPPABLE}
 | {PACKAGE}
 | {FENCED | NOT FENCED}
 | [NOT] LEAKPROOF
 | {CALLED ON NULL INPUT | RETURNS NULL ON NULL INPUT | STRICT }
 | [{ EXTERNAL] SECURITY INVOKER | [EXTERNAL] SECURITY DEFINER |
AUTHID DEFINER | AUTHID CURRENT_USER
}
]
[COST execution_cost
| ROWS result_rows
| SET configuration_parameter { {TO | =} value | FROM CURRENT
| LANGUAGE lang_name
][...]

{
 IS | AS
} plsql_body
/
```

## Parameters

- **function\_name**  
Specifies the name of the function to create (optionally schema-qualified).  
Value range: a string that complies with the [Identifier Naming Conventions](#).  
The value contains a maximum of 63 characters. If the value contains more than 63 characters, the database truncates it and retains the first 63 characters as the function name.
- **argname**  
Specifies the parameter name of the function.  
Value range: a string that complies with the [Identifier Naming Conventions](#).  
The value contains a maximum of 63 characters. If the value contains more than 63 characters, the database truncates it and retains the first 63 characters as the function parameter name.
- **argmode**  
Specifies the parameter mode of the function.  
Value range: **IN**, **OUT**, **INOUT**, and **VARIADIC**. The default value is **IN**. Only the parameters of **OUT** can be followed by **VARIADIC**. The parameters of **OUT** and **INOUT** cannot be used in the function definition of **RETURNS TABLE**.  
 **NOTE**  
**VARIADIC** specifies parameters of the array type.
- **argtype**  
Specifies the parameter type of the function. **%TYPE** or **%ROWTYPE** can be used to indirectly reference a variable or table type. For details, see [Variable Definition Statements](#).
- **expression**

Specifies the default expression of a parameter.

 **NOTE**

It is recommended that you define all default parameters after all non-default parameters.

- **rettype**

Specifies the return data type. Same as **argtype**, **%TYPE** or **%ROWTYPE** can also be used to indirectly reference types.

When there is **OUT** or **INOUT** parameter, the **RETURNS** clause can be omitted. If the clause exists, the result type of the clause must be the same as that of the output parameter. If there are multiple output parameters, the result type of the clause is **RECORD**. Otherwise, the result type of the clause is the same as that of a single output parameter.

The **SETOF** modifier indicates that the function will return a set of items, rather than a single item.

 **NOTE**

In **FUNCTION argtype** and **rettype** outside **PACKAGE**, **%TYPE** cannot reference the type of the *PACKAGE* variable.

- **column\_name**

Specifies the column name.

- **column\_type**

Specifies the column type.

- **definition**

Specifies a string constant defining a function. Its meaning depends on the language. It can be an internal function name, a path pointing to a target file, an SQL query, or text in a procedural language.

- **DETERMINISTIC**

Specifies an API compatible with the SQL syntax. You are advised not to use it.

- **LANGUAGE lang\_name**

Specifies the name of the language that is used to implement the function. It can be **SQL**, **internal**, or the name of a customized process language. To ensure downward compatibility, the name can use single quotation marks. Contents in single quotation marks must be capitalized.

Due to compatibility issues, no matter which language is specified when an A-style database is created, the language used is PL/pgSQL.

- **WINDOW**

Indicates that this function is a window function. The **WINDOW** attribute cannot be changed when replacing an existing function definition.

---

**NOTICE**

For a customized window function, the value of **LANGUAGE** can only be **internal**, and the referenced internal function must be a window function.

---

- **IMMUTABLE**  
Specifies that the function always returns the same result if the parameter values are the same.
- **STABLE**  
Specifies that the function cannot modify the database, and that within a single table scan it will consistently return the same result for the same parameter value, but its result varies by SQL statements.
- **VOLATILE**  
Specifies that the function value can change in a single table scan and no optimization is performed.
- **SHIPPABLE|NOT SHIPPABLE**  
Specifies whether the function can be pushed down for execution. This port is reserved and is not recommended.
- **FENCED|NOT FENCED**  
**FENCED** indicates that a function is called in worker threads to prevent the server from crashing due to incorrect C code implementation. The **NOT FENCED** mode is not recommended.
- **PACKAGE**  
Specifies whether the function can be overloaded. PostgreSQL-style functions can be overloaded, and this parameter is designed for functions of other styles.
  - All PACKAGE and non-PACKAGE functions cannot be overloaded or replaced.
  - PACKAGE functions do not support parameters of the VARIADIC type.
  - The PACKAGE attribute of functions cannot be modified.
- **LEAKPROOF**  
Specifies that the function has no side effects. **LEAKPROOF** can be set only by system administrators.
- **CALLED ON NULL INPUT**  
Declares that some parameters of the function can be called in normal mode if the parameter values are **NULL**. This parameter can be omitted.
- **RETURNS NULL ON NULL INPUT**  
**STRICT**  
Specifies that a function always returns **NULL** when the value of any of its parameters is **NULL**. If this parameter is specified, the function is not executed when there are null parameters; instead a null result is returned automatically.  
**RETURNS NULL ON NULL INPUT** and **STRICT** have the same functions.
- **EXTERNAL**  
The purpose is to be compatible with SQL statements and it is optional. This feature applies to all functions, not only external functions.
- **SECURITY INVOKER**  
**AUTHID CURRENT\_USER**  
Specifies that the function will be executed with the permissions of the user who calls it. This parameter can be omitted.

**SECURITY INVOKER** and **AUTHID CURRENT\_USER** have the same functions.

- **SECURITY DEFINER**

**AUTHID DEFINER**

Specifies that the function will be executed with the permissions of the user who created it.

**AUTHID DEFINER** and **SECURITY DEFINER** have the same functions.

- **COST execution\_cost**

Estimates the execution cost of a function.

The unit of **execution\_cost** is **cpu\_operator\_cost**.

Value range: a positive integer.

- **ROWS result\_rows**

Estimates the number of rows returned by the function. This is only allowed when the function is declared to return a set.

Value range: a positive number. The default value is **1000**.

- **configuration\_parameter**

- **value**

Sets a specified database session parameter to a specified value. If the value is **DEFAULT** or **RESET**, the default setting is used in the new session. **OFF** disables the setting.

Value range: a string.

- **DEFAULT**
    - **OFF**
    - **RESET**
    - Specified default value

- **FROM CURRENT**

Uses the value of **configuration\_parameter** of the current session.

- **plsql\_body**

Specifies the PL/SQL stored procedure body.

---

**NOTICE**

When you perform password-related operations, such as user creation, password change, and encryption/decryption, in a function body, the password will be recorded in the system catalogs and logs in plaintext. To prevent sensitive information leakage, you are advised not to perform operations on the function body related to sensitive information, such as passwords or keys.

---

## Examples

```
-- Define a function as SQL query.
openGauss=# CREATE FUNCTION func_add_sql(integer, integer) RETURNS integer
AS 'select $1 + $2;'
```



```
LANGUAGE SQL
IMMUTABLE
RETURNS NULL ON NULL INPUT;

-- Add an integer by parameter name using PL/pgSQL.
openGauss=# CREATE OR REPLACE FUNCTION func_increment_plsql(i integer) RETURNS integer AS $$
BEGIN
 RETURN i + 1;
END;
$$ LANGUAGE plpgsql;

-- Return the RECORD type.
openGauss=# CREATE OR REPLACE FUNCTION func_increment_sql(i int, out result_1 bigint, out result_2
bigint)
RETURNS SETOF RECORD
AS $$
BEGIN
 result_1 = i + 1;
 result_2 = i * 10;
RETURN next;
END;
$$LANGUAGE PL/SQL;

-- Return a record containing multiple output parameters.
openGauss=# CREATE FUNCTION func_dup_sql(in int, out f1 int, out f2 text)
AS $$ SELECT $1, CAST($1 AS text) || ' is text' $$
LANGUAGE SQL;

openGauss=# SELECT * FROM func_dup_sql(42);

-- Compute the sum of two integers and return the result (if the input is null, the returned result is null).
openGauss=# CREATE FUNCTION func_add_sql2(num1 integer, num2 integer) RETURN integer
AS
BEGIN
RETURN num1 + num2;
END;
/
-- Alter the execution rule of function func_add_sql2 to IMMUTABLE (that is, the same result is returned if
the parameter remains unchanged).
openGauss=# ALTER FUNCTION func_add_sql2(INTEGER, INTEGER) IMMUTABLE;

-- Rename the func_add_sql2 function as add_two_number.
openGauss=# ALTER FUNCTION func_add_sql2(INTEGER, INTEGER) RENAME TO add_two_number;

-- Change the owner of function add_two_number to omm.
openGauss=# ALTER FUNCTION omm(INTEGER, INTEGER) OWNER TO omm;

-- Drop functions.
openGauss=# DROP FUNCTION add_two_number;
openGauss=# DROP FUNCTION func_increment_sql;
openGauss=# DROP FUNCTION func_dup_sql;
openGauss=# DROP FUNCTION func_increment_plsql;
openGauss=# DROP FUNCTION func_add_sql;

-- Set parameters.
openGauss=# SET behavior_compat_options='proc_outparam_override';
-- Create functions.
openGauss=# CREATE OR REPLACE FUNCTION func1(in a integer, out b integer)
RETURNS int
AS $$
DECLARE
 c int;
BEGIN
 c := 1;
 b := a + c;
 return c;
END; $$
LANGUAGE 'PL/SQL' NOT FENCED;
-- Return return and output parameters at the same time.
```

```
openGauss=# DECLARE
 result integer;
 a integer := 2;
 b integer := NULL;
BEGIN
 result := func1(a => a, b => b);
 raise info 'b is: %', b;
 raise info 'result is: %', result;
END;
/
INFO: b is: 3
INFO: result is: 1
ANONYMOUS BLOCK EXECUTE
-- Left assignment expressions are not supported.
openGauss=# DECLARE
 result integer;
 a integer := 2;
 b integer := NULL;
BEGIN
 result := func1(a => a, b => b) + 1;
 raise info 'b is: %', b;
 raise info 'result is: %', result;
END;
/
ERROR: when invoking function func1, maybe input something superfluous.
CONTEXT: compilation of PL/SQL function "inline_code_block" near line 3
-- out/inout in a stored procedure cannot be set to a constant.
openGauss=# DECLARE
 result integer;
 a integer := 2;
 b integer := NULL;
BEGIN
 result := func1(a => a, b => 10);
 raise info 'b is: %', b;
 raise info 'result is: %', result;
END;
/
ERROR: when invoking function func1, no destination for arguments "b"
CONTEXT: compilation of PL/SQL function "inline_code_block" near line 3
```

## Helpful Links

[ALTER FUNCTION](#) and [DROP FUNCTION](#)

## 7.13.53 CREATE GROUP

### Description

Creates a user group.

### Precautions

CREATE GROUP is an alias for CREATE ROLE, and it is not a standard SQL syntax and not recommended. Users can use CREATE ROLE directly.

### Syntax

```
CREATE GROUP group_name [[WITH] option [...]] [ENCRYPTED | UNENCRYPTED] { PASSWORD | IDENTIFIED BY } { 'password' [EXPIRED] | DISABLE };
```

The syntax of the option clause is as follows:

```
{SYSADMIN | NOSYSADMIN}
| {MONADMIN | NOMONADMIN}
```

```
| {OPRADMIN | NOOPRADMIN}
| {POLADMIN | NOPOLADMIN}
| {AUDITADMIN | NOAUDITADMIN}
| {CREATEDB | NOCREATEDB}
| {USEFT | NOUSEFT}
| {CREATEROLE | NOCREATEROLE}
| {INHERIT | NOINHERIT}
| {LOGIN | NOLOGIN}
| {REPLICATION | NOREPLICATION}
| {VCADMIN | NOVCADMIN}
| {PERSISTENCE | NOPERSISTENCE}
| CONNECTION LIMIT connlimit
| VALID BEGIN 'timestamp'
| VALID UNTIL 'timestamp'
| PERM SPACE 'spacelimit'
| TEMP SPACE 'tmpspacelimit'
| SPILL SPACE 'spillspacelimit'
| IN ROLE role_name [, ...]
| IN GROUP role_name [, ...]
| ROLE role_name [, ...]
| ADMIN role_name [, ...]
| USER role_name [, ...]
| SYSID uid
| DEFAULT TABLESPACE tablespace_name
| PROFILE DEFAULT
| PROFILE profile_name
| PGUSER
```

## Parameters

See [Parameters](#) in "CREATE ROLE."

## Examples

```
-- Create group my_group.
openGauss=# CREATE GROUP my_group PASSWORD '*****';

-- Drop the group.
openGauss=# DROP GROUP my_group;
```

## Helpful Links

[ALTER GROUP](#), [DROP GROUP](#), and [CREATE ROLE](#)

## 7.13.54 CREATE INCREMENTAL MATERIALIZED VIEW

### Description

CREATE INCREMENTAL MATERIALIZED VIEW creates a fast-refresh materialized view, and you can refresh the data of the materialized view by using REFRESH MATERIALIZED VIEW (complete-refresh) and REFRESH INCREMENTAL MATERIALIZED VIEW (fast-refresh).

CREATE INCREMENTAL MATERIALIZED VIEW is similar to CREATE TABLE AS, but it remembers the query used to initialize the view, so it can refresh data later. A materialized view has many attributes that are the same as those of a table, but does not support temporary materialized views.

### Precautions

- Fast-refresh materialized views cannot be created on temporary tables or global temporary tables.

- Fast-refresh materialized views support only simple filter queries and UNION ALL queries of base tables.
- Distribution keys cannot be specified when an incremental MV is created.
- After a fast-refresh materialized view is created, most DDL operations in the base table are no longer supported.
- IUD operations cannot be performed on fast-refresh materialized views.
- After a fast-refresh materialized view is created, you need to run the **REFRESH** command to synchronize the materialized view with the base table when the base table data changes.

## Syntax

```
CREATE INCREMENTAL MATERIALIZED VIEW mv_name
 [(column_name [, ...])]
 [TABLESPACE tablespace_name]
 AS query;
```

## Parameters

- **mv\_name**  
Name (optionally schema-qualified) of the materialized view to be created.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Column name in the new materialized view. The materialized view supports specified columns. The number of specified columns must be the same as the number of columns in the result of the subsequent query statement. If no column name is provided, the column name is obtained from the output column name of the query.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **TABLESPACE tablespace\_name**  
Specifies a tablespace to which the new materialized view belongs. If it is not specified, the default tablespace is used.
- **AS query**  
**SELECT** or **TABLE** command. This query will be run in a security-constrained operation.

## Examples

```
-- Create an ordinary table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int);

-- Create a fast-refresh materialized view.
openGauss=# CREATE INCREMENTAL MATERIALIZED VIEW my_imv AS SELECT * FROM my_table;

-- Write data to the base table.
openGauss=# INSERT INTO my_table VALUES(1,1),(2,2);

-- Fast refresh the fast-refresh materialized view my_imv.
openGauss=# REFRESH INCREMENTAL MATERIALIZED VIEW my_imv;

-- Drop a fast-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW my_imv;

-- Drop the ordinary table my_table.
openGauss=# DROP TABLE my_table;
```

## Helpful Links

[ALTER MATERIALIZED VIEW](#), [CREATE MATERIALIZED VIEW](#), [CREATE TABLE](#), [DROP MATERIALIZED VIEW](#), [REFRESH INCREMENTAL MATERIALIZED VIEW](#), and [REFRESH MATERIALIZED VIEW](#)

## 7.13.55 CREATE INDEX

### Description

Creates an index in a specified table.

Indexes are primarily used to enhance database performance (though inappropriate use can result in database performance deterioration). You are advised to create indexes on:

- Columns that are often queried.
- Join conditions. For a query on joined columns, you are advised to create a composite index on the columns. For example, for **select \* from t1 join t2 on t1.a=t2.a and t1.b=t2.b**, you can create a composite index on columns **a** and **b** in table **t1**.
- Columns having filter criteria (especially scope criteria) of a WHERE clause.
- Columns that appear after ORDER BY, GROUP BY, and DISTINCT

Partitioned tables do not support partial index creation when indexes contain the GLOBAL or LOCAL keyword or the created index is a GLOBAL index.

### Precautions

- Indexes consume storage and computing resources. Creating too many indexes has negative impact on database performance (especially the performance of data import. Therefore, you are advised to import the data before creating indexes). Therefore, create indexes only when they are necessary.
- All functions and operators used in an index definition must be immutable, that is, their results must depend only on their parameters and never on any outside influence (such as the contents of another table or the current time). This restriction ensures that the behavior of the index is well-defined. To use a user-defined function in an index or WHERE clause, mark it as an immutable function.
- Partitioned table indexes are classified into local indexes and global indexes. A local index binds to a specific partition, and a global index corresponds to the entire partitioned table.
- A user granted with the CREATE ANY INDEX permission can create indexes in both the public and user schemas.
- If a user-defined function is called in the expression index, the expression index function is executed based on the permission of the function creator.

### Syntax

- Create an index on a table.  

```
CREATE [UNIQUE] INDEX [CONCURRENTLY] [[schema_name.]index_name] ON table_name
[USING method]
```

```
(({ column_name | (expression) } [COLLATE collation] [opclass] [ASC | DESC] [NULLS
{ FIRST | LAST }] }, ...)
[INCLUDE (column_name [, ...])]
[WITH ({ storage_parameter = value } [, ...])]
[TABLESPACE tablespace_name]
[WHERE predicate];
```

- **Create an index on a partitioned table.**

```
CREATE [UNIQUE] INDEX [[schema_name.]index_name] ON table_name [USING method]
(({ column_name | (expression) } [COLLATE collation] [opclass] [ASC | DESC] [NULLS
LAST] }, ...)
[LOCAL [({ PARTITION index_partition_name | SUBPARTITION index_subpartition_name
[TABLESPACE index_partition_tablespace] } [, ...])] | GLOBAL]
[INCLUDE (column_name [, ...])]
[WITH ({ storage_parameter = value } [, ...])]
[TABLESPACE tablespace_name]
[WHERE predicate];
```

## Parameters

- **UNIQUE**

Creates a unique index. In this way, the system checks whether new values are unique in the index column. Attempts to insert or update data which would result in duplicate entries will generate an error.

Currently, only B-tree indexes and UB-tree indexes support unique indexes.

- **CONCURRENTLY**

Creates an index (with ShareUpdateExclusiveLock) in a mode that does not block DML statements. When an index is created, other statements cannot access the table on which the index depends. If this keyword is specified, DML is not blocked during the creation.

- This option can only specify a name of one index.
- The CREATE INDEX statement can be run within a transaction, but CREATE INDEX CONCURRENTLY cannot.
- Partitioned tables do not support index creation in CONCURRENTLY mode. For temporary tables, you can use CONCURRENTLY to create indexes. However, indexes are created in blocking mode because no other sessions concurrently access the temporary tables and the blocking mode is more cost-effective.

 **NOTE**

- This keyword is specified when an index is created. The entire table needs to be scanned twice and built. When the table is scanned for the first time, an index is created and the read and write operations are not blocked. During the second scan, changes that have occurred since the first scan are merged and updated.
  - The table needs to be scanned and built twice, and all existing transactions that may modify the table must be completed. This means that the creation of the index takes a longer time than normal. In addition, the CPU and I/O consumption also affects other services.
  - If an index build fails, it leaves an "unusable" index. This index is ignored by the query, but it still consumes the update overhead. In this case, you are advised to delete the index and try CONCURRENTLY again to create the index.
  - After the second scan, index creation must wait for any transaction that holds a snapshot earlier than the snapshot taken by the second scan to terminate. In addition, ShareUpdateExclusiveLock (level 4) added during index creation conflicts with a lock whose level is greater than or equal to 4. Therefore, when such an index is created, the system is prone to hang or deadlock. For example:
    - If two sessions create an index by using CONCURRENTLY for the same table, a deadlock occurs.
    - If a session creates an index by using CONCURRENTLY for a table and another session drops a table, a deadlock occurs.
    - There are three sessions. Session 1 locks table **a** and does not commit it. Session 2 creates an index by using CONCURRENTLY for table **b**. Session 3 writes data to table **a**. Before the transaction of session 1 is committed, session 2 is blocked.
    - The transaction isolation level is set to repeatable read (read committed by default). Two sessions are started. Session 1 writes data to table **a** and does not commit it. Session 2 creates an index by using CONCURRENTLY for table **b**. Before the transaction of session 1 is committed, session 2 is blocked.
    - When the I/O and CPU resources are not limited, the service performance deterioration caused by online index creation can be controlled within 10%. However, in special scenarios, the service performance deterioration may exceed 10%. This is because online index creation is a long transaction that consumes a large number of I/O and CPU resources. It consumes more resources than offline index creation. The longer the online index creation transaction lasts, the greater the impact on service performance. The time for creating indexes online is positively correlated with the data volume of base tables and the data volume generated by concurrent DML statements. When the I/O and CPU resources are not limited, the time for creating indexes online is about two to six times that for creating indexes offline. However, when the number of concurrent transactions is large (> 10000 TPS) or resource contention occurs, the time may be even longer. You can create indexes in parallel to shorten the index creation time. The performance of online parallel index creation increases to a certain value and becomes stable as the number of parallel worker threads increases. Compared with the performance of creating indexes in serial mode, the performance of creating indexes in parallel online is improved by about 30%. You are advised to create indexes online during off-peak hours to avoid great impact on services. Although online index creation provides the capability of uninterrupted services to some extent, it still needs to be implemented with caution.
- **schema\_name**  
Specifies the schema name.  
Value range: an existing schema name.
  - **index\_name**  
Specifies the name of the index to be created. The schema of the index is the same as that of the table.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **table\_name**

Specifies the name of the table to be indexed (optionally schema-qualified).

Value range: an existing table name.

- **USING method**

Specifies the name of the index method to be used.

Value range:

- **btree**: B-tree indexes store key values of data in a B+ tree structure. This structure helps users to quickly search for indexes. B-tree indexes support comparison queries with ranges specified. When an index is created in a Ustore table, the index is automatically changed to UB-tree.
- **ubtree**: Multi-version B-tree index used only for Ustore tables. The index page contains transaction information and can be recycled. By default, the INSERTPT function is enabled for UB-tree indexes.

Row-store tables (Astore) support the index type: **btree** (default). Row-store tables (Ustore) support the index type: **ubtree**.

- **column\_name**

Specifies the name of the column on which an index is to be created.

Multiple columns can be specified if the index method supports composite indexes. A global index supports a maximum of 31 columns, and other indexes support a maximum of 32 columns.

- **expression**

Specifies an expression based on one or more columns of the table. The expression usually must be written with surrounding parentheses, as shown in the syntax. However, the parentheses can be omitted if the expression has the form of a function call.

Expression can be used to obtain fast access to data based on some transformation of the basic data. For example, an index computed on upper(col) would allow the clause WHERE upper(col) = 'JIM' to use an index.

If an expression contains IS NULL, the index for this expression is invalid. In this case, you are advised to create a partial index.

- **COLLATE collation**

Assigns a collation to the column (which must be of a collatable data type). If no collation is specified, the default collation is used. You can run the **select \* from pg\_collation** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.

- **opclass**

Specifies the name of an operator class. An operator class can be specified for each column of an index. The operator class identifies the operators to be used by the index for that column. For example, a B-tree index on the type int4 would use the **int4\_ops** class; this operator class includes comparison functions for values of type int4. In practice, the default operator class for the column's data type is sufficient. The operator class applies to data with multiple sorts. For example, if you want to sort complex numbers by absolute value or real number, you can define two operator classes and select a proper class when creating an index.



- **ASC**  
Specifies an ascending (default) sort order.
- **DESC**  
Specifies a descending sort order.
- **NULLS FIRST**  
Specifies that null values appear before non-null values in the sort ordering. This is the default when **DESC** is specified.
- **NULLS LAST**  
Specifies that null values appear after non-null values in the sort ordering. This is the default when **DESC** is not specified.
- **LOCAL**  
Specifies that the partitioned index to be created is a local index.
- **GLOBAL**  
Specifies the partitioned index to be created as a global index. If no keyword is specified, a global index is created by default.
- **INCLUDE ( column\_name [, ...] )**  
The optional INCLUDE clause specifies that some non-key columns are included in indexes. Non-key columns cannot be used as search criteria for accelerating index scans, and they are omitted when the unique constraints of the indexes are checked.  
  
An index-only scan can directly return content in the non-key columns without accessing the heap table corresponding to the indexes.  
  
Exercise caution when adding non-key columns as INCLUDE columns, especially for wide columns. If the size of an index tuple exceeds the maximum size allowed by the index type, data insertion fails. Note that in any case, adding non-key columns to an index increases the space occupied by the index, which may slow down the search speed.  
  
Currently, only UB-tree indexes access mode supports this feature. Non-key columns are stored in the index leaf tuple corresponding to the heap tuple and are not included in the tuple on the upper-layer index page.
- **WITH ( {storage\_parameter = value} [, ... ] )**  
Specifies the storage parameter used for an index.  
Value range:  
Indexes except for psort support the **FILLFACTOR** parameter. Only UB-tree indexes support **INDEXSPLIT**.
  - **FILLFACTOR**  
The fill factor of an index is a percentage from 10 to 100.  
Value range: 10–100.
  - **INDEXSPLIT**  
Specifies the splitting policy of UB-tree indexes. The **DEFAULT** policy is the same as the splitting policy of B-tree indexes. The **INSERTPT** policy can significantly reduce the index space usage in some scenarios.  
Value range: **INSERTPT** and **DEFAULT**.  
Default value: **INSERTPT**.

- **TABLESPACE tablespace\_name**  
Specifies the tablespace for an index. If no tablespace is specified, the default tablespace is used.  
Value range: an existing tablespace name.
- **WHERE predicate**  
Creates a partial index. A partial index is an index that contains entries for only a portion of a table, usually a portion that is more useful for indexing than the rest of the table. For example, if you have a table that contains both billed and unbilled orders where the unbilled orders take up a small fraction of the total table and yet that is an often used portion, you can improve performance by creating an index on just that portion. In addition, the WHERE clause with a UNIQUE constraint can be used to enforce uniqueness over a subset of a table.  
Value range: The predicate expression can only refer to columns of the underlying table, but it can use all columns, not just the ones being indexed. Currently, subqueries and aggregate expressions are forbidden in WHERE. You are advised not to use numeric types such as int for predicate, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and **0** is implicitly converted to **false**), which may cause unexpected results.  
For a partitioned table index, if the created index contains the GLOBAL or LOCAL keyword or the created index is a GLOBAL index, the WHERE clause cannot be used to create an index.
- **PARTITION index\_partition\_name**  
Specifies the name of an index partition.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **SUBPARTITION index\_subpartition\_name**  
Specifies the name of a level-2 index partition.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **TABLESPACE index\_partition\_tablespace**  
Specifies the tablespace of an index partition.  
Value range: If this parameter is not specified, the value of **index\_tablespace** is used.

## Examples

```
-- Create the tpcds.ship_mode_t1 table.
openGauss=# CREATE SCHEMA tpcds;
openGauss=# CREATE TABLE tpcds.ship_mode_t1
(
 SM_SHIP_MODE_SK INTEGER NOT NULL,
 SM_SHIP_MODE_ID CHAR(16) NOT NULL,
 SM_TYPE CHAR(30) ,
 SM_CODE CHAR(10) ,
 SM_CARRIER CHAR(20) ,
 SM_CONTRACT CHAR(20)
)
;

-- Create a common unique index on the SM_SHIP_MODE_SK column in the tpcds.ship_mode_t1 table.
openGauss=# CREATE UNIQUE INDEX ds_ship_mode_t1_index1 ON
tpcds.ship_mode_t1(SM_SHIP_MODE_SK);
```

```
-- Create a B-tree index on the SM_SHIP_MODE_SK column in the tpcds.ship_mode_t1 table.
openGauss=# CREATE INDEX ds_ship_mode_t1_index4 ON tpcds.ship_mode_t1 USING
btree(SM_SHIP_MODE_SK);

-- Create an expression index on the SM_CODE column in the tpcds.ship_mode_t1 table:
openGauss=# CREATE INDEX ds_ship_mode_t1_index2 ON tpcds.ship_mode_t1(SUBSTR(SM_CODE,1,4));

-- Create a partial index on the SM_SHIP_MODE_SK column where SM_SHIP_MODE_SK is greater than 10
in the tpcds.ship_mode_t1 table.
openGauss=# CREATE UNIQUE INDEX ds_ship_mode_t1_index3 ON
tpcds.ship_mode_t1(SM_SHIP_MODE_SK) WHERE SM_SHIP_MODE_SK>10;

-- Rename an existing index.
openGauss=# ALTER INDEX tpcds.ds_ship_mode_t1_index1 RENAME TO ds_ship_mode_t1_index5;

-- Set the index as unusable.
openGauss=# ALTER INDEX tpcds.ds_ship_mode_t1_index2 UNUSABLE;

-- Rebuild an index.
openGauss=# ALTER INDEX tpcds.ds_ship_mode_t1_index2 REBUILD;

-- Drop an existing index.
openGauss=# DROP INDEX tpcds.ds_ship_mode_t1_index2;

-- Drop the table.
openGauss=# DROP TABLE tpcds.ship_mode_t1;

-- Create a tablespace.
openGauss=# CREATE TABLESPACE example1 RELATIVE LOCATION 'tablespace1/tablespace_1';
openGauss=# CREATE TABLESPACE example2 RELATIVE LOCATION 'tablespace2/tablespace_2';
openGauss=# CREATE TABLESPACE example3 RELATIVE LOCATION 'tablespace3/tablespace_3';
openGauss=# CREATE TABLESPACE example4 RELATIVE LOCATION 'tablespace4/tablespace_4';
-- Create the tpcds.customer_address_p1 table.
openGauss=# CREATE TABLE tpcds.customer_address_p1
(
 CA_ADDRESS_SK INTEGER NOT NULL,
 CA_ADDRESS_ID CHAR(16) NOT NULL,
 CA_STREET_NUMBER CHAR(10) ,
 CA_STREET_NAME VARCHAR(60) ,
 CA_STREET_TYPE CHAR(15) ,
 CA_SUITE_NUMBER CHAR(10) ,
 CA_CITY VARCHAR(60) ,
 CA_COUNTY VARCHAR(30) ,
 CA_STATE CHAR(2) ,
 CA_ZIP CHAR(10) ,
 CA_COUNTRY VARCHAR(20) ,
 CA_GMT_OFFSET DECIMAL(5,2) ,
 CA_LOCATION_TYPE CHAR(20)
)
TABLESPACE example1
PARTITION BY RANGE(CA_ADDRESS_SK)
(
 PARTITION p1 VALUES LESS THAN (3000),
 PARTITION p2 VALUES LESS THAN (5000) TABLESPACE example1,
 PARTITION p3 VALUES LESS THAN (MAXVALUE) TABLESPACE example2
)
ENABLE ROW MOVEMENT;
-- Create the partitioned table index ds_customer_address_p1_index1 without specifying the index
partition name.
openGauss=# CREATE INDEX ds_customer_address_p1_index1 ON
tpcds.customer_address_p1(CA_ADDRESS_SK) LOCAL;
-- Create the partitioned table index ds_customer_address_p1_index2 with the name of the index partition
specified.
openGauss=# CREATE INDEX ds_customer_address_p1_index2 ON
tpcds.customer_address_p1(CA_ADDRESS_SK) LOCAL
(
 PARTITION CA_ADDRESS_SK_index1,
 PARTITION CA_ADDRESS_SK_index2 TABLESPACE example3,
 PARTITION CA_ADDRESS_SK_index3 TABLESPACE example4
)
```

```
)
TABLESPACE example2;

-- Create a global partitioned index.
openGauss=# CREATE INDEX ds_customer_address_p1_index3 ON
tpcds.customer_address_p1(CA_ADDRESS_ID) GLOBAL;

-- If no keyword is specified, a global partitioned index is created by default.
openGauss=# CREATE INDEX ds_customer_address_p1_index4 ON
tpcds.customer_address_p1(CA_ADDRESS_ID);

-- Change the tablespace of the partitioned table index CA_ADDRESS_SK_index2 to example1.
openGauss=# ALTER INDEX tpcds.ds_customer_address_p1_index2 MOVE PARTITION
CA_ADDRESS_SK_index2 TABLESPACE example1;

-- Change the tablespace of the partitioned table index CA_ADDRESS_SK_index3 to example2.
openGauss=# ALTER INDEX tpcds.ds_customer_address_p1_index2 MOVE PARTITION
CA_ADDRESS_SK_index3 TABLESPACE example2;

-- Rename a partitioned table index.
openGauss=# ALTER INDEX tpcds.ds_customer_address_p1_index2 RENAME PARTITION
CA_ADDRESS_SK_index1 TO CA_ADDRESS_SK_index4;

-- Drop the indexes and the partitioned table.
openGauss=# DROP INDEX tpcds.ds_customer_address_p1_index1;
openGauss=# DROP INDEX tpcds.ds_customer_address_p1_index2;
openGauss=# DROP TABLE tpcds.customer_address_p1;
-- Drop tablespaces.
openGauss=# DROP TABLESPACE example1;
openGauss=# DROP TABLESPACE example2;
openGauss=# DROP TABLESPACE example3;
openGauss=# DROP TABLESPACE example4;
```

## Helpful Links

[ALTER INDEX](#) and [DROP INDEX](#)

## Suggestions

- create index

You are advised to create indexes on:

- Columns that are often queried.
- Join conditions. For a query on joined columns, you are advised to create a composite index on the columns. For example, for **select \* from t1 join t2 on t1.a=t2.a and t1.b=t2.b**, you can create a composite index on columns **a** and **b** in table **t1**.
- Columns having filter criteria (especially scope criteria) of a WHERE clause.
- Columns that appear after ORDER BY, GROUP BY, and DISTINCT

Constraints:

- An index of an ordinary table supports a maximum of 32 columns. A GLOBAL index of a partitioned table supports a maximum of 31 columns.
- The size of a single index cannot exceed the size of the index page (8 KB). The size of a B-tree or UB-tree index cannot exceed one-third of the page size.
- Partial indexes cannot be created in a partitioned table.
- When a GLOBAL index is created on a partitioned table, the following constraints apply:

- Expression indexes and partial indexes are not supported.
- Only B-tree indexes are supported.
- In the same attribute column, the local index and global index of a partition cannot coexist.
- If the ALTER statement does not contain UPDATE GLOBAL INDEX, the original global index is invalid. In this case, other indexes are used for query. If the ALTER statement contains UPDATE GLOBAL INDEX, the original global index is still valid and the index function is correct.

## 7.13.56 CREATE LANGUAGE

This version does not support this syntax.

## 7.13.57 CREATE MASKING POLICY

### Description

Creates a masking policy.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

The masking policy takes effect only after the security policy is enabled, that is, **enable\_security\_policy** is set to **on**.

### Syntax

```
CREATE MASKING POLICY policy_name masking_clause[, ...]* policy_filter [
policy_filter_clause
][ENABLE | DISABLE];
```

- **masking\_clause**  
masking\_function ON LABEL(label\_name[, ...])

- **masking\_function**

The maskall function is not preset. It is hard-coded and cannot be displayed by running **\df**.

The masking methods during presetting are as follows:

```
{ maskall | randommasking | creditcardmasking | basicemailmasking | fullemailmasking |
shufflemasking | alldigitsmasking | regexprmasking }
```

- **policy\_filter\_clause:**  
FILTER ON { FILTER\_TYPE ( filter\_value [, ...] ) } [, ...]
- **FILTER\_TYPE:**  
IP | APP | ROLES

### Parameters

- **policy\_name**  
Specifies the audit policy name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).

- **label\_name**  
Specifies the resource label name.
- **masking\_clause**  
Specifies the masking function to be used to anonymize database resources labeled by **label\_name**. **schema.function** can be used to specify the masking function.
- **policy\_filter**  
Specifies the users for which the masking policy takes effect. If this parameter is left empty, the masking policy takes effect for all users.
- **FILTER\_TYPE**  
Specifies the types of information to be filtered by the policy, including **IP**, **APP**, and **ROLES**.
- **filter\_value**  
Specifies the detailed information to be filtered, such as the IP address, app name, and username.
- **ENABLE|DISABLE**  
Enables or disables the masking policy. If **ENABLE|DISABLE** is not specified, **ENABLE** is used by default.

## Examples

```
-- Create users dev_mask and bob_mask.
openGauss=# CREATE USER dev_mask PASSWORD '*****';
openGauss=# CREATE USER bob_mask PASSWORD '*****';

-- Create table tb_for_masking.
openGauss=# CREATE TABLE tb_for_masking(col1 text, col2 text, col3 text);

-- Create a resource label for sensitive column col1.
openGauss=# CREATE RESOURCE LABEL mask_lb1 ADD COLUMN(tb_for_masking.col1);

-- Create a resource label for sensitive column col2.
openGauss=# CREATE RESOURCE LABEL mask_lb2 ADD COLUMN(tb_for_masking.col2);

-- Create a masking policy for the operation of accessing sensitive column col1.
openGauss=# CREATE MASKING POLICY maskpol1 maskall ON LABEL(mask_lb1);

-- Create a masking policy that takes effect only for scenarios where users are dev_mask and bob_mask,
the client tool is gsql, and IP addresses are 10.20.30.40, and 127.0.0.0/24.
openGauss=# CREATE MASKING POLICY maskpol2 randommasking ON LABEL(mask_lb2) FILTER ON
ROLES(dev_mask, bob_mask), APP(gsql), IP('10.20.30.40', '127.0.0.0/24');

-- Drop masking policies.
openGauss=# DROP MASKING POLICY maskpol1, maskpol2;

-- Drop resource labels.
openGauss=# DROP RESOURCE LABEL mask_lb1, mask_lb2;

-- Drop the tb_for_masking table.
openGauss=# DROP TABLE tb_for_masking;

-- Drop the dev_mask and bob_mask users.
openGauss=# DROP USER dev_mask, bob_mask;
```

## Helpful Links

[ALTER MASKING POLICY](#) and [DROP MASKING POLICY](#)

## 7.13.58 CREATE MATERIALIZED VIEW

CREATE MATERIALIZED VIEW creates a complete-refresh materialized view, and you can use REFRESH MATERIALIZED VIEW (full refresh) to refresh the data in the materialized view.

CREATE MATERIALIZED VIEW is similar to CREATE TABLE AS, but it remembers the query used to initialize the view, so it can refresh data later. A materialized view has many attributes that are the same as those of a table, but does not support temporary materialized views.

### Precautions

- Complete-refresh materialized views cannot be created in temporary tables or global temporary tables.
- Complete-refresh materialized views do not support NodeGroups.
- After a complete-refresh materialized view is created, most DDL operations in the base table are no longer supported.
- IUD operations cannot be performed on complete-refresh materialized views.
- After a complete-refresh materialized view is created, if the base table data changes, you need to run the **REFRESH** command to synchronize the materialized view with the base table.
- Ustore does not support the creation and use of materialized views.

### Syntax

```
CREATE MATERIALIZED VIEW mv_name
[(column_name [, ...])]
[WITH ({storage_parameter = value} [, ...])]
[TABLESPACE tablespace_name]
AS query
[WITH [NO] DATA];
```

### Parameters

- **mv\_name**  
Name (optionally schema-qualified) of the materialized view to be created.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Column name in the new materialized view. The materialized view supports specified columns. The number of specified columns must be the same as the number of columns in the result of the subsequent query statement. If no column name is provided, the column name is obtained from the output column name of the query.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **WITH ( storage\_parameter [= value] [, ...] )**  
Specifies an optional storage parameter for a table or an index. For details, see [CREATE TABLE](#).
- **TABLESPACE tablespace\_name**  
Specifies a tablespace to which the new materialized view belongs. If it is not specified, the default tablespace is used.

- **AS query**

Specifies the **SELECT**, **TABLE**, or **VALUES** command. This query will be run in a security-constrained operation.

## Examples

```
-- Create an ordinary table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int);

-- Create a complete-refresh materialized view.
openGauss=# CREATE MATERIALIZED VIEW my_mv AS SELECT * FROM my_table;

-- Write data to the base table.
openGauss=# INSERT INTO my_table VALUES(1,1),(2,2);

-- Completely refresh the complete-refresh materialized view my_mv.
openGauss=# REFRESH MATERIALIZED VIEW my_mv;

-- Drop a complete-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW my_mv;

-- Drop the ordinary table my_table.
openGauss=# DROP TABLE my_table;
```

## Helpful Links

[ALTER MATERIALIZED VIEW](#), [CREATE INCREMENTAL MATERIALIZED VIEW](#), [CREATE TABLE](#), [DROP MATERIALIZED VIEW](#), [REFRESH INCREMENTAL MATERIALIZED VIEW](#), and [REFRESH MATERIALIZED VIEW](#)

## 7.13.59 CREATE OPERATOR

### Description

Defines a new operator.

### Precautions

CREATE OPERATOR defines a new name operator. The user who defines the operator becomes the owner of the operator. If a schema name is given, the operator is created in the specified schema. Otherwise, it will be created in the current schema.

The operator name is a character string consisting of the following characters:

+ - \* / < > = ~ ! @ # % ^ & | ` ?

When selecting a name, note the following restrictions:

- -- and /\* cannot appear anywhere in the operator name, because they are regarded as the beginning of a comment.
- A multi-character operator cannot end with + or - unless the name contains at least one of the following characters:  
~ ! @ # % ^ & | ` ?
- => is no longer used in an operator name.

Operator! = is mapped to <> when being entered. Therefore, the two names are always equivalent.



At least one LEFTARG and one RIGHTARG must be defined. For binocular operators, both need to be defined. For the right operator, only LEFTARG needs to be defined. For the left operator, only RIGHTARG needs to be defined.

Also, the function\_name procedure must have been defined with CREATE FUNCTION, and must be defined to accept the correct number of specified type parameters (one or two).

Other clauses declare optional operator optimization clauses, as defined in [35.13. Operator Optimization Information](#).

To create an operator, you must have the USAGE permission on the parameter type and return type, and the EXECUTE permission on the underlying function. If exchange or negative operators are specified, you must have them.

## Syntax

```
CREATE OPERATOR name (
 PROCEDURE = function_name
 [, LEFTARG = left_type] [, RIGHTARG = right_type]
 [, COMMUTATOR = com_op] [, NEGATOR = neg_op]
 [, RESTRICT = res_proc] [, JOIN = join_proc]
 [, HASHES] [, MERGES]
)
```

## Parameters

- **name**  
Operator to be defined. The available characters are listed above. The name can be schema-qualified, for example, CREATE OPERATOR myschema.+ (...). If there is no schema, the operator is created in the current schema. Two operators in the same schema can have the same name as long as they perform operations on different data types. This is a reloading process.
- **function\_name**  
Function used to implement the operator.
- **left\_type**  
Parameter data type on the left of the operator, if any. This parameter can be omitted if the left operator is used.
- **right\_type**  
Parameter data type on the right of the operator, if any. This parameter can be omitted if the right-view operator is used.
- **com\_op**  
Exchange operator corresponding to the operator.
- **neg\_op**  
Negative operator corresponding to the operator.
- **res\_proc**  
This operator constrains the selectivity evaluation function.
- **join\_proc**  
This operator joins the selectivity evaluation function.
- **HASHES**  
Indicates that the operator supports hash joins.

- **MERGES**

Indicates that this operator supports a merge join.

Use the OPERATOR() syntax to provide a schema-qualified operator name in com\_op or other optional parameters. For example:

```
COMMUTATOR = OPERATOR(myschema.===) ,
```

## Examples

The following command defines a new operator: equal area for the box data type.

```
CREATE OPERATOR === (
 LEFTARG = box,
 RIGHTARG = box,
 PROCEDURE = area_equal_procedure,
 COMMUTATOR = ===,
 NEGATOR = !==,
 RESTRICT = area_restriction_procedure,
 JOIN = area_join_procedure,
 HASHES, MERGES
);
```

## Helpful Links

[ALTER OPERATOR](#) and [DROP OPERATOR](#)

## 7.13.60 CREATE OPERATOR CLASS

### Description

Defines a new operator class.

### Precautions

CREATE OPERATOR CLASS defines a new operator class. An operator class defines how to use a specified data type together with an index. The operator class declares that certain operators will provide particular roles or "strategies" for this data type and this index method. The operator class also declares the supported programs used by the index method when the operator class is selected for an index column. All the operators and functions used by an operator class must be defined before the operator class is created.

If a schema name is given, then the operator class is created in the specified schema. Otherwise, it is created in the current schema. Two operator classes in the same schema can have the same name only if they are for different index methods.

The user who defines an operator class becomes the owner. Currently, the creator must be a super user.

CREATE OPERATOR CLASS does not check whether the class definition includes all the operators and functions required by the index method, nor whether the operators and functions form a self-contained set.

Related operator classes can be grouped into operator families. To add a new operator class to an existing family, specify a FAMILY option in CREATE OPERATOR CLASS. Without this option, the new class is placed into a family with the same name (a family is created if it does not exist).

## Syntax

```
CREATE OPERATOR CLASS
name [DEFAULT] FOR TYPE data_type
USING index_method [
FAMILY family_name] AS
{ OPERATOR strategy_number operator_name [(
op_type, op_type)] [FOR SEARCH | FOR ORDER BY sort_family_name]
| FUNCTION
support_number [(op_type [, op_type])] function_name (argument_type [,
...])
| STORAGE storage_type
} [, ...]
```

## Parameters

- **name**  
Specifies the name of the operator class to be created (which can be schema-qualified).
- **default**  
Specifies that the operator class will become the default operator class for its data type. At most one operator class can be the default for a specific data type and index method.
- **data\_type**  
Specifies the column data type processed by the operator class.
- **index\_method**  
Specifies the name of the index method processed by the operator class.
- **family\_name**  
Specifies the name of an existing operator family where an operator class is added. If it is not specified, the class is placed into a family with the same name (a family is created if it does not exist).
- **strategy\_number**  
Specifies the strategy number of the index method associated with the operator class.
- **operator\_name**  
Specifies the name (which can be schema-qualified) of the operator associated with the operator class.
- **op\_type**  
Specifies the data type of the operator's operand in an OPERATOR clause. The value **NONE** represents a left unary operator or right unary operator. The data type of the operand can be omitted if it is the same as that of the operator class.  
  
In a FUNCTION clause, if the data type of the function's operand is different from the input data type of the function (such as B-tree comparison functions and hash functions) or the class's data type, the data type of the operand supported by this function must be included in this clause. These default values are correct. Therefore, **op\_type** need not be specified in FUNCTION clauses, except in cases where the B-tree sort support function supports cross-type comparisons.
- **sort\_family\_name**

Specifies the name of an existing B-tree operator family that describes the sort ordering associated with a sort operator.

The default value is **FOR SEARCH**.

- **support\_number**

Specifies the number of an index method for a function associated with the operator class.

- **function\_name**

Specifies a function name of an index method for an operator class.

- **argument\_type**

Specifies a data type of a function parameter.

- **storage\_type**

Specifies the data type stored in the index. Normally, this is the same as the column data type, but some index methods allow it to be different. The **STORAGE** clause must be omitted unless the index method allows a different type to be used.

## Examples

```
-- Define a function.
CREATE FUNCTION func_add_sql(num1 integer, num2 integer) RETURN integer AS BEGIN
RETURN num1 + num2;
END;
/
-- Create an operator class and associate it with the preceding function.
CREATE OPERATOR CLASS oc1 DEFAULT FOR TYPE _int4 USING btree AS FUNCTION 1 func_add_sql
(integer, integer);
```

## 7.13.61 CREATE PACKAGE

### Description

Creates a package.

### Precautions

- Only the centralized mode support packages.
- The functions or stored procedures declared in the package specification must be defined in the package body.
- During instantiation, the stored procedure with **commit** or **rollback** cannot be called.
- Package functions cannot be called in triggers.
- Variables in a package cannot be directly used in external SQL statements.
- Private variables and stored procedures in a package cannot be called outside the package.
- Usage that other stored procedures do not support are not supported. For example, if **commit** or **rollback** cannot be called in a function, they cannot be called in the function of a package, either.
- The name of a schema cannot be the same as that of a package.
- Only A-version stored procedures and function definitions are supported.

- Variables with the same name in a package, including parameters with the same name in a package, are not supported.
- The global variables in a package are at the session level. The variables in packages cannot be shared in different sessions.
- When a function of an autonomous transaction is called in a package, the cursor variables in the package and recursive functions that use the cursor variables in the package are not allowed.
- The package does not declare the ref cursor variables.
- The default permission on a package is SECURITY INVOKER. To change the default permission to SECURITY DEFINER, set the GUC parameter **behavior\_compat\_options** to '**plsql\_security\_definer**'.
- A user granted with the CREATE ANY PACKAGE permission can create packages in the public and user schemas.
- If the name of a package to be created contains special characters, spaces cannot be contained between special characters. You are advised to set the GUC parameter **behavior\_compat\_options** to "**skip\_insert\_gs\_source**". Otherwise, an error may occur.
- When you create a package function, the default parameter value cannot contain variables.

## Syntax

- CREATE PACKAGE SPECIFICATION

```
CREATE [OR REPLACE] PACKAGE [schema] package_name
[invoker_rights_clause] { IS | AS } item_list_1 END package_name;
```

**invoker\_rights\_clause** can be declared as **AUTHID DEFINER** or **AUTHID INVOKER**, which indicate the definer permission and invoker permission, respectively.

**item\_list\_1** can be a declared variable, stored procedure, or function.

The package specification (header) declares public variables, functions, and exceptions in a package, which can be called by external functions or stored procedures. It can only declare stored procedures and functions but cannot define them.

- CREATE PACKAGE BODY

```
CREATE [OR REPLACE] PACKAGE BODY [schema] package_name
{ IS | AS } declare_section [initialize_section] END package_name;
```

The package body defines private variables and functions in a package. If a variable or function is not declared by the package specification, it is a private variable or function.

The package body also has an initialization part to initialize the package. For details, see the example.

## Examples

- Example of CREATE PACKAGE SPECIFICATION

```
CREATE OR REPLACE PACKAGE emp_bonus IS
var1 int:=1;-- Public variable
var2 int:=2;
PROCEDURE testpro1(var3 int);-- Public stored procedure, which can be called by external systems.
END emp_bonus;
/
```

- Example of CREATE PACKAGE BODY

```
drop table if exists test1;
create or replace package body emp_bonus is
```

```
var3 int:=3;
var4 int:=4;
procedure testpro1(var3 int)
is
begin
create table if not exists test1(col1 int);
insert into test1 values(var1);
insert into test1 values(var4);
end;
begin: --The instantiation starts.
var4:=9;
testpro1(var4);
end emp_bonus;
/
```

- Example of ALTER PACKAGE OWNER

```
ALTER PACKAGE emp_bonus OWNER TO omm;
-- Change the owner of PACKAGE emp_bonus to omm.
```

- Example of calling a package

```
call emp_bonus.testpro1(1); -- Use call to call the stored procedure of a package.
select emp_bonus.testpro1(1); -- Use select to call the stored procedure of a package.
-- Call the stored procedure of a package in an anonymous block.
begin
emp_bonus.testpro1(1);
end;
/
```

## 7.13.62 CREATE PROCEDURE

### Description

Creates a stored procedure.

### Precautions

- If the parameters or return values of a stored procedure have precision, the precision is not checked.
- When creating a stored procedure, you are advised to explicitly specify the schemas of all operations on table objects in the stored procedure definition. Otherwise, the stored procedure may fail to be executed.
- **current\_schema** and **search\_path** specified by **SET** during stored procedure creation are invalid. **search\_path** and **current\_schema** before and after function execution should be the same.
- If a SELECT statement calls a stored procedure with output parameters, output parameters cannot be specified. When a CALL statement calls a stored procedure involving a non-overloaded function, output parameters must be specified. When a CALL statement calls an overloaded PACKAGE function, the output parameters may not be specified. For details, see examples in [CALL](#).
- A stored procedure with the PACKAGE attribute can use overloaded functions.
- When you create a procedure, you cannot nest average functions within aggregate or other functions.
- When stored procedures without parameters are called in another stored procedure, you can omit brackets and call stored procedures using their names directly.
- When functions with output parameters are called in a stored procedure which is an assignment expression, you can omit the output parameters of the called functions.

- The stored procedure supports viewing, exporting, and importing parameter comments.
- The stored procedure supports viewing, exporting, and importing parameter comments between IS/AS and plsql\_body.
- The default permission on a stored procedure is SECURITY INVOKER. To change the default permission to SECURITY DEFINER, set the GUC parameter **behavior\_compat\_options** to '**plsql\_security\_definer**'.
- Users granted the CREATE ANY FUNCTION permission can create or replace stored procedures in the user schemas.
- An **out** or **inout** parameter must be set to a variable but not a constant.
- In a centralized environment, if you want to call a stored procedure with the same in parameters but different out parameters, you need to set the GUC parameter **behavior\_compat\_options** to '**proc\_outparam\_override**'. After the parameter is enabled, you must add **out** parameters no matter whether you use the SELECT or CALL statement to call the stored procedure. After the parameter is enabled, you cannot use **perform** to call a stored procedure or function.
- When an overloaded stored procedure is called, the variable type must be specified.
- If a stored procedure with the definer specified is created in a schema of another user, the stored procedure will be executed by another user, which may cause unauthorized operations. Therefore, exercise caution when performing this operation.

## Syntax

```
CREATE [OR REPLACE] PROCEDURE procedure_name
[(([argname] [argmode] argtype [{ DEFAULT | := | = } expression]) [, ...])]
[
 { IMMUTABLE | STABLE | VOLATILE }
 | { SHIPPABLE | NOT SHIPPABLE }
 | { PACKAGE }
 | [NOT] LEAKPROOF
 | { CALLED ON NULL INPUT | RETURNS NULL ON NULL INPUT | STRICT }
 | [{ EXTERNAL } SECURITY INVOKER | [EXTERNAL] SECURITY DEFINER | AUTHID DEFINER | AUTHID
CURRENT_USER }
 | COST execution_cost
 | SET configuration_parameter { TO value | = value | FROM CURRENT }
][...]
{ IS | AS }
plsql_body
/
```

## Parameters

- **OR REPLACE**  
Replaces the original definition when two stored procedures are with the same name.
- **procedure\_name**  
Specifies the name of the stored procedure that is created (optionally with schema names).  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **argmode**  
Specifies the mode of an argument.

---

**NOTICE**

**VARIADIC** specifies parameters of the array type.

---

Value range: **IN**, **OUT**, **INOUT**, and **VARIADIC**. The default value is **IN**. Only the parameters in **OUT** mode can follow the **VARIADIC** parameter.

- **argname**

Specifies the argument name.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **argtype**

Specifies the type of an argument. **%TYPE** or **%ROWTYPE** can be used to indirectly reference a variable or table type. For details, see [Variable Definition Statements](#).

Value range: a valid data type.

 **NOTE**

In **PROCEDURE argtype** outside *PACKAGE*, **%TYPE** cannot reference the type of the *PACKAGE* variable.

- **expression**

Specifies the default expression of a parameter.

 **NOTE**

- It is recommended that you define all default parameters after all non-default parameters.

- **configuration\_parameter**

- **value**

Sets the specified configuration parameter to a specified value. If the **value** is **DEFAULT**, the default setting is used in the new session. **OFF** disables the setting.

Value range: a string.

- **DEFAULT**
    - **OFF**
    - Specified default value

- **from current**

Uses the value of **configuration\_parameter** of the current session.

- **IMMUTABLE, STABLE,...**

Specifies a constraint. The function of each parameter is similar to that of **CREATE FUNCTION**. For details, see [CREATE FUNCTION](#).

- **plsql\_body**

Specifies the PL/SQL stored procedure body.



**NOTICE**

When you perform password-related operations, such as user creation, password change, and encryption/decryption, in a stored procedure, the password will be recorded in the system catalogs and logs in plaintext. To prevent sensitive information leakage, you are advised not to perform password-related operations in a stored procedure.

 **NOTE**

No specific order is applied to **argname** and **argname**. The following order is advised: **argname**, **argmode**, and **argtype**.

## Helpful Links

[DROP PROCEDURE](#)

## Suggestions

- analyse | analyze
  - Do not run **ANALYZE** in a transaction or anonymous block.
  - Do not run **ANALYZE** in a function or stored procedure.

## 7.13.63 CREATE RESOURCE LABEL

### Description

Creates a resource label.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

### Syntax

```
CREATE RESOURCE LABEL [IF NOT EXISTS] label_name ADD label_item_list[, ...]*;
```

- label\_item\_list  
resource\_type(resource\_path[, ...]\*)
- resource\_type  
TABLE | COLUMN | SCHEMA | VIEW | FUNCTION

### Parameters

- **IF NOT EXISTS**  
If a resource label with the same name already exists, no error is generated. Instead, a message is displayed, indicating that the resource label already exists.
- **label\_name**  
Specifies the resource label name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).

- **resource\_type**  
Specifies the type of database resources to be labeled.
- **resource\_path**  
Specifies the path of database resources.

## Examples

```
-- Create table tb_for_label.
openGauss=# CREATE TABLE tb_for_label(col1 text, col2 text, col3 text);

-- Create schema schema_for_label.
openGauss=# CREATE SCHEMA schema_for_label;

-- Create view view_for_label.
openGauss=# CREATE VIEW view_for_label AS SELECT 1;

-- Create function func_for_label.
openGauss=# CREATE FUNCTION func_for_label RETURNS TEXT AS $$ SELECT col1 FROM tb_for_label; $$
LANGUAGE SQL;

-- Create a resource label based on a table.
openGauss=# CREATE RESOURCE LABEL IF NOT EXISTS table_label add TABLE(public.tb_for_label);

-- Create a resource label based on columns.
openGauss=# CREATE RESOURCE LABEL IF NOT EXISTS column_label add
COLUMN(public.tb_for_label.col1);

-- Create a resource label based on a schema.
openGauss=# CREATE RESOURCE LABEL IF NOT EXISTS schema_label add SCHEMA(schema_for_label);

-- Create a resource label based on a view.
openGauss=# CREATE RESOURCE LABEL IF NOT EXISTS view_label add VIEW(view_for_label);

-- Create a resource label based on a function.
openGauss=# CREATE RESOURCE LABEL IF NOT EXISTS func_label add FUNCTION(func_for_label);

-- Drop resource labels.
openGauss=# DROP RESOURCE LABEL func_label, view_label, schema_label, column_label, table_label;

-- Drop the func_for_label function.
openGauss=# DROP FUNCTION func_for_label;

-- Drop the view_for_label view.
openGauss=# DROP VIEW view_for_label;

-- Drop the schema_for_label schema.
openGauss=# DROP SCHEMA schema_for_label;

-- Drop the tb_for_label table.
openGauss=# DROP TABLE tb_for_label;
```

## Helpful Links

[ALTER RESOURCE LABEL](#) and [DROP RESOURCE LABEL](#)

## 7.13.64 CREATE ROLE

### Description

Creates a role.

A role is an entity that owns database objects and permissions. In different environments, a role can be considered a user, a group, or both.

## Precautions

- If a role is added to the database, the role does not have the login permission.
- Only the user who has the CREATE ROLE permission or a system administrator is allowed to create roles.

## Syntax

```
CREATE ROLE role_name [[WITH] option [...]] [ENCRYPTED | UNENCRYPTED] { PASSWORD | IDENTIFIED BY } { 'password' [EXPIRED] | DISABLE };
```

The syntax of option clauses for configuring role information is as follows:

```
{SYSADMIN | NOSYSADMIN}
| {MONADMIN | NOMONADMIN}
| {OPRADMIN | NOOPRADMIN}
| {POLADMIN | NOPOLADMIN}
| {AUDITADMIN | NOAUDITADMIN}
| {CREATEDB | NOCREATEDB}
| {USEFT | NOUSEFT}
| {CREATEROLE | NOCREATEROLE}
| {INHERIT | NOINHERIT}
| {LOGIN | NOLOGIN}
| {REPLICATION | NOREPLICATION}
| {VCADMIN | NOVCADMIN}
| {PERSISTENCE | NOPERSISTENCE}
| CONNECTION LIMIT connlimit
| VALID BEGIN 'timestamp'
| VALID UNTIL 'timestamp'
| USER GROUP 'groupuser'
| PERM SPACE 'spacelimit'
| TEMP SPACE 'tmpspacelimit'
| SPILL SPACE 'spillspacelimit'
| NODE GROUP logic_cluster_name
| IN ROLE role_name [, ...]
| IN GROUP role_name [, ...]
| ROLE role_name [, ...]
| ADMIN rol e_name [, ...]
| USER role_name [, ...]
| SYSID uid
| DEFAULT TABLESPACE tablespace_name
| PROFILE DEFAULT
| PROFILE profile_name
| PGUSER
```

## Parameters

- **role\_name**

Specifies the name of a role.

Value range: a string that follows the [Identifier Naming Conventions](#) with a maximum of 63 characters. If a value contains more than 63 characters, the database truncates it and retains the first 63 characters as the role name.

When a role is created, the database will display a message.

 **NOTE**

The identifier must be letters, underscores (\_), digits (0-9), or dollar signs (\$) and must start with a letter (a-z) or underscore (\_).

- **password**

Specifies the login password.

A new password must:

- Contain at least eight characters. This is the default length.

- Differ from the username or the username spelled backward.
- Contain at least three of the following character types: uppercase characters, lowercase characters, digits, and special characters (limited to ~!@#\$%^&\*()-\_+=\|[{]}<.>/?).
- The password can also be a ciphertext character string that meets the format requirements. This mode is mainly used to import user data. You are advised not to use it directly. If a ciphertext password is used, you need to know the plaintext corresponding to the ciphertext password and ensure a complex plaintext password. The database does not verify the complexity of the ciphertext password, so you should ensure the password security.
- When creating a role, enclose the user password in single quotation marks.

Value range: a character string that cannot be empty.

- **EXPIRED**

When creating a user, you can specify the **EXPIRED** parameter to create a user whose password is invalid. The user cannot perform simple or extended queries. The statement can be executed only after the password is changed.

- **DISABLE**

By default, you can change your password unless it is disabled. To disable the password of a user, use this parameter. After the password of a user is disabled, the password will be deleted from the system. The user can connect to the database only through external authentication, for example, Kerberos authentication. Only administrators can enable or disable a password. Common users cannot disable the password of an initial user. To enable a password, run **ALTER USER** and specify the password.

- **ENCRYPTED | UNENCRYPTED**

Controls whether the password is stored encrypted in the system catalogs. According to product security requirement, the password must be stored encrypted. Therefore, **UNENCRYPTED** is forbidden in GaussDB. If the password string has already been encrypted in the SHA256 format, it is stored encrypted as it was, regardless of whether **ENCRYPTED** or **UNENCRYPTED** is specified (since the system cannot decrypt the specified encrypted password string). This allows reloading of encrypted passwords during dump/restore.

- **SYSADMIN | NOSYSADMIN**

Specifies whether a new role is SYSADMIN, which has the highest permission. The default value is **NOSYSADMIN**.

When separation of duties is disabled, users with the SYSADMIN permission can create users with the SYSADMIN, REPLICATION, CREATEROLE, AUDITADMIN, MONADMIN, POLADMIN, or CREATEDB permission and common users.

When separation of duties is enabled, users with the SYSADMIN permission do not have the permission to create users.

- **MONADMIN | NOMONADMIN**

Specifies whether a role is a MONADMIN.

The default value is **NOMONADMIN**.

- **OPRADMIN | NOOPRADMIN**

Specifies whether a role is an OPRADMIN.

The default value is **NOOPRADMIN**.

- **POLADMIN | NOPOLADMIN**

Specifies whether a role is a POLADMIN.

The default value is **NOPOLADMIN**.

- **AUDITADMIN | NOAUDITADMIN**

Specifies whether a role has the audit and management attributes.

The default value is **NOAUDITADMIN**.

- **CREATEDB | NOCREATEDB**

Specifies a role's permission to create databases.

A new role does not have the permission to create databases.

The default value is **NOCREATEDB**.

- **USEFT | NOUSEFT**

This parameter is reserved and not used in this version.

- **CREATEROLE | NOCREATEROLE**

Specifies whether a role will be permitted to create new roles (that is, execute **CREATE ROLE** and **CREATE USER**). A role with the **CREATEROLE** permission can also modify and delete other roles.

The default value is **NOCREATEROLE**.

When separation of duties is disabled, users with the **CREATEROLE** permission can create users with the **CREATEROLE**, **AUDITADMIN**, **MONADMIN**, **POLADMIN**, or **CREATEDB** permission and common users.

When separation of duties is enabled, users with the **CREATEROLE** permission can create users with the **CREATEROLE**, **MONADMIN**, **POLADMIN**, or **CREATEDB** permission and common users.

- **INHERIT | NOINHERIT**

Specifies whether a role "inherits" the permissions of roles in the same group. It is not recommended.

- **LOGIN | NOLOGIN**

Specifies whether a role is allowed to log in to a database. A role with the **LOGIN** attribute can be considered as a user.

The default value is **NOLOGIN**.

- **REPLICATION | NOREPLICATION**

Specifies whether a role is allowed to initiate streaming replication or put the system in and out of backup mode. A role with **REPLICATION** attribute is specific to replication.

The default value is **NOREPLICATION**.

- **VCADMIN | NOVADMIN**

This parameter has no actual meaning.

- **PERSISTENCE | NOPERSISTENCE**

Defines a permanent user. Only the initial user is allowed to create, modify, and delete permanent users with the **PERSISTENCE** attribute.

- **CONNECTION LIMIT**

Specifies how many concurrent connections the role can make.

---

**NOTICE**

- System administrators are not restricted by this parameter.
- The number of concurrent connections of each primary database node is calculated separately (which is the value of **connlimit**). Number of all connections of the database = Value of **connlimit** x Number of normal primary database nodes.

---

Value range: an integer in the range  $[-1, 2^{31}-1]$ . The default value is **-1**, indicating that there is no limit.

- **VALID BEGIN**  
Sets a date and time when the role's password takes effect. If this clause is omitted, the password takes effect immediately.
- **VALID UNTIL**  
Sets a date and time after which the role's password is no longer valid. If this clause is omitted, the password will be valid for all time.
- **USER GROUP**  
Creates a sub-user. This function is not supported in the current version.
- **PERM SPACE**  
Sets the space available for a user.
- **TEMP SPACE**  
Sets the space allocated to the temporary table of a user.
- **SPILL SPACE**  
Sets the operator disk flushing space of a user.
- **NODE GROUP**  
Name of the node group associated with a user. This function is not supported in the current version.
- **IN ROLE**  
Lists one or more existing roles to which the new role will be immediately added as a new member. It is not recommended.
- **IN GROUP**  
Specifies an obsolete spelling of IN ROLE. It is not recommended.
- **ROLE**  
Lists one or more existing roles which are automatically added as members of the new role.
- **ADMIN**  
Similar to ROLE. However, ADMIN grants permissions of new roles to other roles.
- **USER**  
Specifies an obsolete spelling of the ROLE clause.
- **SYSID**

The SYSID clause is ignored.

- **DEFAULT TABLESPACE**

The DEFAULT TABLESPACE clause is ignored.

- **PROFILE**

The PROFILE clause is ignored.

- **PGUSER**

In the current version, this attribute is reserved only for forward compatibility.

## Examples

```
-- Create a role manager whose password is *****.
openGauss=# CREATE ROLE manager IDENTIFIED BY '*****';

-- Create a role with its validity from January 1, 2015 to January 1, 2026.
openGauss=# CREATE ROLE miriam WITH LOGIN PASSWORD '*****' VALID BEGIN '2015-01-01' VALID
UNTIL '2026-01-01';

-- Change the password of role manager to *****.
openGauss=# ALTER ROLE manager IDENTIFIED BY '*****' REPLACE '*****';

-- Change role manager to SYSADMIN.
openGauss=# ALTER ROLE manager SYSADMIN;

-- Drop role manager.
openGauss=# DROP ROLE manager;

-- Drop role miriam.
openGauss=# DROP GROUP miriam;
```

## Helpful Links

[SET ROLE](#), [ALTER ROLE](#), [DROP ROLE](#), and [GRANT](#)

## 7.13.65 CREATE ROW LEVEL SECURITY POLICY

### Description

Creates a row-level security policy for a table.

The policy takes effect only after row-level security is enabled (by running **ALTER TABLE... ENABLE ROW LEVEL SECURITY**). Otherwise, this statement does not take effect.

Currently, row-level security affects the read (SELECT, UPDATE, and DELETE) of data tables and does not affect the write (INSERT and MERGE INTO) of data tables. The table owner or system administrators can create an expression in the USING clause. When the client reads the data table, the database server combines the expressions that meet the condition and applies it to the execution plan in the statement rewriting phase of a query. For each tuple in a data table, if the expression returns **TRUE**, the tuple is visible to the current user; if the expression returns **FALSE** or **NULL**, the tuple is invisible to the current user.

A row-level security policy name is specific to a table. A data table cannot have row-level security policies with the same name. Different data tables can have the same row-level security policy.

Row-level security policies can be applied to specified operations (**SELECT**, **UPDATE**, **DELETE**, and **ALL**). **ALL** indicates that **SELECT**, **UPDATE**, and **DELETE**

will be affected. For a new row-level security policy, the default value **ALL** will be used if you do not specify the operations that will be affected.

Row-level security policies can be applied to a specified user (role) or to all users (**PUBLIC**). For a new row-level security policy, the default value **PUBLIC** will be used if you do not specify the user who will be affected.

## Precautions

- Row-level security policies can be defined for row-store tables, row-store partitioned tables, unlogged tables, and hash tables.
- Row-level security policies cannot be defined for foreign tables and local temporary tables.
- Row-level security policies cannot be defined for views.
- A maximum of 100 row-level security policies can be defined for a table.
- System administrators are not affected by row-level security policies and can view all data in a table.
- Tables queried by using SQL statements, views, functions, and stored procedures are affected by row-level security policies.
- The data type of a table column to which a row-level security policy has been added cannot be changed.

## Syntax

```
CREATE [ROW LEVEL SECURITY] POLICY policy_name ON table_name
[AS { PERMISSIVE | RESTRICTIVE }]
[FOR { ALL | SELECT | UPDATE | DELETE }]
[TO { role_name | PUBLIC | CURRENT_USER | SESSION_USER } [, ...]]
USING (using_expression)
```

## Parameters

- **policy\_name**  
Specifies the name of a row-level security policy to be created. The names of row-level security policies for a table must be unique.
- **table\_name**  
Specifies the name of a table to which a row-level security policy is applied.
- **PERMISSIVE | RESTRICTIVE**  
**PERMISSIVE** enables the permissive policy for row-level security. The conditions of the permissive policy are joined through the OR expression.  
**RESTRICTIVE** enables the restrictive policy for row-level security. The conditions of the restrictive policy are joined through the AND expression. The join methods are as follows:  

```
(using_expression_permissive_1 OR using_expression_permissive_2 ...) AND
(using_expression_restrictive_1 AND using_expression_restrictive_2 ...)
```

  
The default value is **PERMISSIVE**.
- **command**  
Specifies the SQL operations limited by a row-level security policy, including **ALL**, **SELECT**, **UPDATE**, and **DELETE**. If this parameter is not specified, the default value **ALL** will be used, covering **SELECT**, **UPDATE**, and **DELETE**.



If *command* is set to **SELECT**, only tuple data that meets the condition (the return value of *using\_expression* is **TRUE**) can be queried. The operations that are affected include **SELECT**, **UPDATE.... RETURNING**, and **DELETE... RETURNING**.

If *command* is set to **UPDATE**, only tuple data that meets the condition (the return value of *using\_expression* is **TRUE**) can be updated. The operations that are affected include **UPDATE**, **UPDATE ... RETURNING**, and **SELECT ... FOR UPDATE/SHARE**.

If *command* is set to **DELETE**, only tuple data that meets the condition (the return value of *using\_expression* is **TRUE**) can be deleted. The operations that are affected include **DELETE** and **DELETE ... RETURNING**.

The following table describes the relationship between row-level security policies and SQL statements.

**Table 7-105** Relationship between row-level security policies and SQL statements.

| Command                        | SELECT/ALL policy | UPDATE/ALL policy | DELETE/ALL policy |
|--------------------------------|-------------------|-------------------|-------------------|
| <b>SELECT</b>                  | Existing row      | No                | No                |
| <b>SELECT FOR UPDATE/SHARE</b> | Existing row      | Existing row      | No                |
| <b>UPDATE</b>                  | No                | Existing row      | No                |
| <b>UPDATE RETURNING</b>        | Existing row      | Existing row      | No                |
| <b>DELETE</b>                  | No                | No                | Existing row      |
| <b>DELETE RETURNING</b>        | Existing row      | No                | Existing row      |

- **role\_name**

Specifies database users limited by a row-level security policy.

- **CURRENT\_USER**: username of the current operating environment.
- **SESSION\_USER**: session username.
- If this parameter is not specified, the default value **PUBLIC** will be used, indicating that all database users will be limited. You can specify multiple database users.

---

**NOTICE**

System administrators are not limited by the row-level security policy.

---

- **using\_expression**

Specifies an expression defined for a row-level security policy (return type: Boolean).

The expression cannot contain aggregate functions or window functions. In the statement rewriting phase of a query, if row-level security for a data table is enabled, the expressions that meet the specified conditions will be added to the plan tree. The expression is calculated for each tuple in the data table. For SELECT, UPDATE, and DELETE, row data is visible to the current user only when the return value of the expression is **TRUE**. If the expression returns **FALSE**, the tuple is invisible to the current user. In this case, the user cannot view the tuple through the SELECT statement, update the tuple through the UPDATE statement, or delete the tuple through the DELETE statement.

## Examples

```
-- Create user alice.
openGauss=# CREATE USER alice PASSWORD '*****';

-- Create user bob.
openGauss=# CREATE USER bob PASSWORD '*****';

-- Create data table all_data.
openGauss=# CREATE TABLE public.all_data(id int, role varchar(100), data varchar(100));

-- Insert data into the data table.
openGauss=# INSERT INTO all_data VALUES(1, 'alice', 'alice data');
openGauss=# INSERT INTO all_data VALUES(2, 'bob', 'bob data');
openGauss=# INSERT INTO all_data VALUES(3, 'peter', 'peter data');

-- Grant the read permission on the all_data table to users alice and bob.
openGauss=# GRANT SELECT ON all_data TO alice, bob;

-- Enable the row-level security policy.
openGauss=# ALTER TABLE all_data ENABLE ROW LEVEL SECURITY;

-- Create a row-level security policy to specify that the current user can view only their own data.
openGauss=# CREATE ROW LEVEL SECURITY POLICY all_data_rls ON all_data USING(role =
CURRENT_USER);

-- View information about the all_data table.
openGauss=# \d+ all_data
 Table "public.all_data"
Column | Type | Modifiers | Storage | Stats target | Description
-----+-----+-----+-----+-----+-----
id | integer | | plain | |
role | character varying(100) | | extended | |
data | character varying(100) | | extended | |
Row Level Security Policies:
 POLICY "all_data_rls"
 USING (((role)::name = "current_user"()))
Has OIDs: no
Options: orientation=row, enable_rowsecurity=true

-- Run SELECT.
openGauss=# SELECT * FROM all_data;
 id | role | data
-----+-----+-----
 1 | alice | alice data
 2 | bob | bob data
 3 | peter | peter data
(3 rows)

openGauss=# EXPLAIN(COSTS OFF) SELECT * FROM all_data;
 QUERY PLAN

Seq Scan on all_data
(1 row)

-- Switch to user alice and run SELECT.
```

```
openGauss=# SELECT * FROM all_data;
 id | role | data
-----+-----+-----
 1 | alice | alice data
(1 row)

openGauss=# EXPLAIN(COSTS OFF) SELECT * FROM all_data;
QUERY PLAN

Seq Scan on all_data
 Filter: ((role)::name = 'alice'::name)
Notice: This query is influenced by row level security feature
(3 rows)

-- Drop a row-level security policy.
openGauss=# DROP ROW LEVEL SECURITY POLICY all_data_rls ON all_data;

-- Drop the all_data table.
openGauss=# DROP TABLE public.all_data;

-- Drop users alice and bob.
openGauss=# DROP USER alice, bob;
```

## Helpful Links

[DROP ROW LEVEL SECURITY POLICY](#) and [ALTER ROW LEVEL SECURITY POLICY](#)

## 7.13.66 CREATE RULE

### Description

Defines a new rewriting rule.

### Precautions

- To define or modify rules for a table, you must be the owner of the table.
- If multiple rules of the same type are defined for the same table, the rules are triggered one by one by name in alphabetical order.
- In the view, the RETURNING clause can be added to the INSERT, UPDATE, and DELETE rules to return columns by view. If a rule is triggered by the **INSERT RETURNING**, **UPDATE RETURNING**, or **DELETE RETURNING** command, these clauses are used to calculate the output result. If a rule is triggered by a command without RETURNING, the RETURNING clause of the rule is ignored. Currently, only unconditional INSTEAD rules can contain the RETURNING clause, and only one RETURNING clause can exist in all rules of one event. This ensures that only one RETURNING clause can be used for result calculation. If the RETURNING clause does not exist in any valid rule, the RETURNING query in this view will be rejected.
- Currently, ON SELECT rules must be unconditional INSTEAD rules and must have actions consisting of a single SELECT query. Therefore, an ON SELECT rule actually turns a table into a view whose visible content is the content returned by the **SELECT** command of the rule, rather than the content in the table (if any).

### Syntax

```
CREATE [OR REPLACE] RULE name AS ON event
 TO table_name [WHERE condition]
 DO [ALSO | INSTEAD] { NOTHING | command | (command ; command ...) }
```

Events include:

```
SELECT
INSERT
DELETE
UPDATE
```

## Parameters

- **name**  
Name of the created rule. It must be unique among all the rules for the same table.  
Value range: a string, which complies with the [Identifier Naming Conventions](#). A value can contain a maximum of 63 characters.
- **event**  
One of the SELECT, INSERT, UPDATE, and DELETE events.
- **table\_name**  
Name (optionally schema-qualified) of the table or view to which the rule applies.
- **condition**  
SQL condition expression that returns a Boolean value, which determines whether to execute the rule. Expressions cannot reference any table except **NEW** and **OLD**, and cannot have aggregate functions. You are advised not to use numeric types such as int for **condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and **0** is implicitly converted to **false**), which may cause unexpected results.
- **INSTEAD**  
**INSTEAD** indicates that the initial event is replaced with this command.
- **ALSO**  
**ALSO** indicates that the command should be executed after the initial event. If neither **ALSO** nor **INSTEAD** is specified, **ALSO** is the default value.
- **command**  
Command that composes the rule action. A valid command is one of the SELECT, INSERT, UPDATE, and DELETE statements.

## Examples

```
CREATE RULE "_RETURN" AS
ON SELECT TO t1
DO INSTEAD
SELECT * FROM t2;
```

### NOTICE

- The name specified after ON SELECT rules must be "\_RETURN".
- Currently, the ON SELECT rule must be INSTEAD SELECT, and the table specified by TO is converted to a view. The prerequisite is that the table is empty and does not have restrictions such as triggers, indexes, and child tables. That is, the table must be an initial empty table. Therefore, you are advised to create a view directly instead of using this method.

## 7.13.67 CREATE SCHEMA

### Description

Creates a schema.

Named objects are accessed either by "qualifying" their names with the schema name as a prefix, or by setting a search path that includes the desired schema. When creating named objects, you can also use the schema name as a prefix.

Optionally, CREATE SCHEMA can include sub-commands to create objects within the new schema. The sub-commands are treated essentially the same as separate commands issued after creating the schema. If the AUTHORIZATION clause is used, all the created objects are owned by this user.

### Precautions

- Only a user with the CREATE permission on the current database can perform this operation.
- The owner of an object created by a system administrator in a schema with the same name as a common user is the common user, not a system administrator.

### Syntax

- Create a schema based on a specified name.  

```
CREATE SCHEMA schema_name
[AUTHORIZATION user_name] [schema_element [...]];
```
- Create a schema based on a username.  

```
CREATE SCHEMA AUTHORIZATION user_name [schema_element [...]];
```

### Parameters

- **schema\_name**  
Specifies the schema name.

---

#### NOTICE

The name must be unique.

The schema name cannot start with **pg\_**.

---

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **AUTHORIZATION user\_name**  
Specifies the owner of a schema. If **schema\_name** is not specified, **user\_name** will be used as the schema name. In this case, **user\_name** can only be a role name.  
Value range: an existing username or role name.
- **schema\_element**  
Specifies an SQL statement defining an object to be created within the schema. Currently, only the CREATE TABLE, CREATE VIEW, CREATE INDEX,

CREATE PARTITION, CREATE SEQUENCE, CREATE TRIGGER, and GRANT clauses are supported.

Objects created by sub-commands are owned by the user specified by **AUTHORIZATION**.

#### NOTE

If objects in the schema on the current search path are with the same name, specify the schemas for different objects. You can run **SHOW SEARCH\_PATH** to check the schemas on the current search path.

## Examples

```
-- Create the role1 role.
openGauss=# CREATE ROLE role1 IDENTIFIED BY '*****1';

-- Create a schema named role1 for the role1 role. The owner of the films and winners tables created by
the clause is role1.
openGauss=# CREATE SCHEMA AUTHORIZATION role1
CREATE TABLE films (title text, release date, awards text[])
CREATE VIEW winners AS
SELECT title, release FROM films WHERE awards IS NOT NULL;

-- Drop the schema.
openGauss=# DROP SCHEMA role1 CASCADE;
-- Drop the user.
openGauss=# DROP USER role1 CASCADE;
```

## Helpful Links

[ALTER SCHEMA](#) and [DROP SCHEMA](#)

## 7.13.68 CREATE SEQUENCE

### Function

**CREATE SEQUENCE** adds a sequence to the current database. The owner of a sequence is the user who creates the sequence.

### Precautions

- A sequence is a special table that stores arithmetic progressions. It has no actual meaning and is usually used to generate unique identifiers for rows or tables.
- If a schema name is given, the sequence is created in the specified schema; otherwise, it is created in the current schema. The sequence name must be different from the names of other sequences, tables, indexes, views in the same schema.
- After the sequence is created, functions **nextval()** and **generate\_series(1,N)** insert data to the table. Make sure that the number of times for invoking **nextval** is greater than or equal to N+1. Otherwise, errors will be reported because the number of times for invoking function **generate\_series()** is N+1.
- By default, the maximum value of **Sequence** is  $2^{63} - 1$ . If a large identifier is used, the maximum value can be  $2^{127} - 1$ .
- A user granted with the **CREATE ANY SEQUENCE** permission can create sequences in the public and user schemas.

## Syntax

```
CREATE [LARGE] SEQUENCE name [INCREMENT [BY] increment]
 [MINVALUE minvalue | NO MINVALUE | NOMINVALUE] [MAXVALUE maxvalue | NO MAXVALUE |
 NOMAXVALUE]
 [START [WITH] start] [CACHE cache] [[NO] CYCLE | NOCYCLE]
 [OWNED BY { table_name.column_name | NONE }];
```

## Parameter Description

- **name**  
Specifies the name of a sequence to be created.  
Value range: a sting containing only lowercase letters, uppercase letters, special characters #\_\$, and digits.
- **increment**  
Specifies the step for a sequence. A positive number generates an ascending sequence, and a negative number generates a decreasing sequence.  
The default value is **1**.
- **MINVALUE minvalue | NO MINVALUE| NOMINVALUE**  
Specifies the minimum value of the sequence. If **MINVALUE** is not declared, or **NO MINVALUE** is declared, the default value of the ascending sequence is **1**, and that of the descending sequence is **-2<sup>63</sup>-1**. **NOMINVALUE** is equivalent to **NO MINVALUE**.
- **MAXVALUE maxvalue | NO MAXVALUE| NOMAXVALUE**  
Specifies the maximum value of the sequence. If **MAXVALUE** is not declared, or **NO MAXVALUE** is declared, the default value of the ascending sequence is **2<sup>63</sup>-1**, and that of the descending sequence is **-1**. **NOMAXVALUE** is equivalent to **NO MAXVALUE**.
- **start**  
Specifies the start value of the sequence. The default value for an ascending sequence is **minvalue** and that for a descending sequence is **maxvalue**.
- **cache**  
Specifies the number of sequences stored in the memory for quick access purposes.  
Default value **1** indicates that one sequence can be generated each time.  
  
 **NOTE**  
It is not recommended that you define **cache** and **maxvalue** or **minvalue** at the same time. The continuity of sequences cannot be ensured after **cache** is defined because unacknowledged sequences may be generated, causing waste of sequences.
- **CYCLE**  
Recycles sequences after the number of sequences reaches **maxvalue** or **minvalue**.  
If **NO CYCLE** is specified, any invocation of **nextval** would return an error after the number of sequences reaches **maxvalue** or **minvalue**.  
**NOCYCLE** is equivalent to **NO CYCLE**.  
The default value is **NO CYCLE**.  
If **CYCLE** is specified, the sequence uniqueness cannot be ensured.

- **OWNED BY**

Associates a sequence with a specified column included in a table. In this way, the sequence will be deleted when you delete its associated column or the table where the column belongs to. The associated table and sequence must be owned by the same user and in the same schema. **OWNED BY** only establishes the association between a table column and the sequence. Sequences on the column do not increase automatically when data is inserted.

The default value **OWNED BY NONE** indicates that such association does not exist.

---

**NOTICE**

You are not advised to use the sequence created using **OWNED BY** in other tables. If multiple tables need to share a sequence, the sequence must not belong to a specific table.

---

## Examples

Create an ascending sequence named **serial**, which starts from 101.

```
openGauss=# CREATE SEQUENCE serial
START 101
CACHE 20;
```

Select the next number from the sequence.

```
openGauss=# SELECT nextval('serial');
nextval

101
```

Select the next number from the sequence.

```
openGauss=# SELECT nextval('serial');
nextval

102
```

Create a sequence associated with the table.

```
openGauss=# CREATE TABLE customer_address
(
 ca_address_sk integer not null,
 ca_address_id char(16) not null,
 ca_street_number char(10) ,
 ca_street_name varchar(60) ,
 ca_street_type char(15) ,
 ca_suite_number char(10) ,
 ca_city varchar(60) ,
 ca_county varchar(30) ,
 ca_state char(2) ,
 ca_zip char(10) ,
 ca_country varchar(20) ,
 ca_gmt_offset decimal(5,2) ,
 ca_location_type char(20)
);

openGauss=# CREATE SEQUENCE serial1
START 101
CACHE 20
```



```
OWNED BY customer_address.ca_address_sk;
-- Delete a table and sequences.
openGauss=# DROP TABLE customer_address;
openGauss=# DROP SEQUENCE serial cascade;
openGauss=# DROP SEQUENCE serial1 cascade;
```

## Helpful Links

[DROP SEQUENCE](#) and [ALTER SEQUENCE](#)

## 7.13.69 CREATE SERVER

### Function

Defines a new foreign server.

### Precautions

When multi-layer quotation marks are used for sensitive columns (such as **password** and **secret\_access\_key**) in **OPTIONS**, the semantics is different from that in the scenario where quotation marks are not used. Therefore, sensitive columns are not identified for anonymization.

### Syntax

```
CREATE SERVER server_name
 FOREIGN DATA WRAPPER fdw_name
 OPTIONS ({ option_name ' value ' } [, ...]) ;
```

### Parameter Description

- **server\_name**  
Specifies the server name.  
Value range: a string containing no more than 63 characters
- **fdw\_name**  
Specifies the name of the foreign data wrapper.  
Value range: **dist\_fdw**, **log\_fdw**, and **file\_fdw**
- **OPTIONS ( { option\_name ' value ' } [, ...] )**  
Specifies options for the server. These options typically define the connection details of the server, but the actual names and values depend on the foreign data wrapper of the server.
  - Specifies the parameters for the foreign server. The detailed parameter description is as follows:
    - **encrypt**  
Specifies whether data is encrypted. This parameter is available only when **type** is **OBS**. The default value is **on**.  
Value range:
      - **on** indicates that data is encrypted and HTTPS is used for communication.
      - **off** indicates that data is not encrypted and HTTP is used for communication.

- **access\_key**  
Specifies the access key (AK) (obtained by users from the OBS console) used for the OBS access protocol. When you create a foreign table, the AK value is encrypted and saved to the metadata table of the database. This parameter is available only when **type** is set to **OBS**.
- **secret\_access\_key**  
Specifies the secret key (SK) value (obtained by users from the OBS console) used for the OBS access protocol. When you create a foreign table, the SK value is encrypted and saved to the metadata table of the database. This parameter is available only when **type** is set to **OBS**.

In addition to the connection parameters supported by libpq, the following parameters are provided:

- **fdw\_startup\_cost**  
Estimates the startup time required for a foreign table scan, including the time to establish a connection, analyze the request at the remote server, and generate a plan. The default value is **100**.
- **fdw\_tuple\_cost**  
Specifies the additional consumption when each tuple is scanned on a remote server. The value specifies the extra consumption of data transmission between servers. The default value is **0.01**.

## Examples

```
-- Create a server.
openGauss=# create server my_server foreign data wrapper log_fdw;
CREATE SERVER

-- Delete my_server.
openGauss=# DROP SERVER my_server;
DROP SERVER
```

## Helpful Links

[ALTER SERVER](#) and [DROP SERVER](#)

## 7.13.70 CREATE SYNONYM

### Description

Creates a synonym object. A synonym is an alias of a database object and is used to record the mapping between database object names. You can use synonyms to access associated database objects.

### Precautions

- The user of a synonym should be its owner.
- If the schema name is specified, create a synonym in the specified schema. Otherwise create a synonym in the current schema.

- Database objects that can be accessed using synonyms include tables, views, functions, and stored procedures.
- To use synonyms, you must have the required permissions on associated objects.
- The following DML statements support synonyms: SELECT, INSERT, UPDATE, DELETE, EXPLAIN, and CALL.
- You are advised not to create synonyms for temporary tables. To create a synonym, you need to specify the schema name of the target temporary table. Otherwise, the synonym cannot be used. In addition, you need to run the **DROP SYNONYM** command before the current session ends.
- After an object is deleted, the synonym associated with the object will not be deleted in cascading mode. If you continue to access the synonym, an error message is displayed, indicating that the synonym has expired.

## Syntax

```
CREATE [OR REPLACE] SYNONYM synonym_name
FOR object_name;
```

## Parameters

- **synonym**  
Specifies the name of the synonym to be created, which can contain the schema name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **object\_name**  
Specifies the name of an object that is associated (optionally with schema names).  
Value range: a string that complies with the [Identifier Naming Conventions](#).

### NOTE

**object\_name** can be the name of an object that does not exist.

### CAUTION

Do not create aliases for functions that contain passwords and other sensitive information, such as the encryption function `gs_encrypt` and the decryption function `gs_decrypt` or use aliases to call the functions to prevent sensitive information leakage.

## Examples

```
-- Create schema ot.
openGauss=# CREATE SCHEMA ot;

-- Create table ot.t1 and its synonym t1.
openGauss=# CREATE TABLE ot.t1(id int, name varchar2(10));
openGauss=# CREATE OR REPLACE SYNONYM t1 FOR ot.t1;

-- Use synonym t1.
openGauss=# SELECT * FROM t1;
openGauss=# INSERT INTO t1 VALUES (1, 'ada'), (2, 'bob');
```

```
openGauss=# UPDATE t1 SET t1.name = 'cici' WHERE t1.id = 2;

-- Create synonym v1 and its associated view ot.v_t1.
openGauss=# CREATE SYNONYM v1 FOR ot.v_t1;
openGauss=# CREATE VIEW ot.v_t1 AS SELECT * FROM ot.t1;

-- Use synonym v1.
openGauss=# SELECT * FROM v1;

-- Create overloaded function ot.add and its synonym add.
openGauss=# CREATE OR REPLACE FUNCTION ot.add(a integer, b integer) RETURNS integer AS
$$
SELECT $1 + $2
$$
LANGUAGE sql;

openGauss=# CREATE OR REPLACE FUNCTION ot.add(a decimal(5,2), b decimal(5,2)) RETURNS
decimal(5,2) AS
$$
SELECT $1 + $2
$$
LANGUAGE sql;

openGauss=# CREATE OR REPLACE SYNONYM add FOR ot.add;

-- Use synonym add.
openGauss=# SELECT add(1,2);
openGauss=# SELECT add(1.2,2.3);

-- Create stored procedure ot.register and its synonym register.
openGauss=# CREATE PROCEDURE ot.register(n_id integer, n_name varchar2(10))
SECURITY INVOKER
AS
BEGIN
 INSERT INTO ot.t1 VALUES(n_id, n_name);
END;
/

openGauss=# CREATE OR REPLACE SYNONYM register FOR ot.register;

-- Use synonym register to call the stored procedure.
openGauss=# CALL register(3,'mia');

-- Drop synonyms.
openGauss=# DROP SYNONYM t1;
openGauss=# DROP SYNONYM IF EXISTS v1;
openGauss=# DROP SYNONYM IF EXISTS add;
openGauss=# DROP SYNONYM register;
openGauss=# DROP SCHEMA ot CASCADE;
```

## Helpful Links

[ALTER SYNONYM](#) and [DROP SYNONYM](#)

## 7.13.71 CREATE TABLE

### Description

Creates an initially empty table in the current database. The table will be owned by the creator.

## Precautions

- If an error occurs during table creation, after it is fixed, the system may fail to delete the empty disk files created before the last automatic clearance. This problem seldom occurs and does not affect system running of the database.
- When JDBC is used, the **DEFAULT** value can be set through **PreparedStatement**.
- A user granted with the CREATE ANY TABLE permission can create tables in the public and user schemas. To create a table that contains serial columns, you must also obtain the CREATE ANY SEQUENCE permission to create sequences.

## Syntax

Create a table.

```
CREATE [[GLOBAL | LOCAL] [TEMPORARY | TEMP] | UNLOGGED] TABLE [IF NOT EXISTS] table_name
 ({ column_name data_type [COLLATE collation] [column_constraint [...]]
 | table_constraint
 | LIKE source_table [like_option [...]] }
 [, ...])
 [WITH ({storage_parameter = value} [, ...])]
 [ON COMMIT { PRESERVE ROWS | DELETE ROWS | DROP }]
 [TABLESPACE tablespace_name];
```

- **column\_constraint** is as follows:

```
[CONSTRAINT constraint_name]
{ NOT NULL |
 NULL |
 CHECK (expression) |
 DEFAULT default_expr |
 GENERATED ALWAYS AS (generation_expr) STORED |
 UNIQUE index_parameters |
 PRIMARY KEY index_parameters |
 REFERENCES reftable [(refcolumn)] [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
 [ON DELETE action] [ON UPDATE action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- **table\_constraint** is as follows:

```
[CONSTRAINT constraint_name]
{ CHECK (expression) |
 UNIQUE (column_name [, ...]) index_parameters |
 PRIMARY KEY (column_name [, ...]) index_parameters |
 FOREIGN KEY (column_name [, ...]) REFERENCES reftable [(refcolumn [, ...])]
 [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE] [ON DELETE action] [ON UPDATE action]
 |
 PARTIAL CLUSTER KEY (column_name [, ...]) }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- **like\_option** is as follows:

```
{ INCLUDING | EXCLUDING } { DEFAULTS | GENERATED | CONSTRAINTS | INDEXES | STORAGE |
 COMMENTS | PARTITION | REOPTIONS | ALL }
```

- **index\_parameters** is as follows:

```
[WITH ({storage_parameter = value} [, ...])]
[USING INDEX TABLESPACE tablespace_name]
```

## Parameters

- **UNLOGGED**

If this keyword is specified, the created table is an unlogged table. Data written to unlogged tables is not written to the WALs, which makes them considerably faster than ordinary tables. However, the data in the unlogged table is cleared after a conflict occurs, the OS is restarted, the database is

restarted, a switchover is performed, the power supply is cut off, or the database is restarted abnormally, which may cause data loss. Contents of an unlogged table are also not replicated to standby nodes. Any indexes created on an unlogged table are not automatically logged as well.

Usage scenario: Unlogged tables do not ensure data security. Users can back up data before using unlogged tables; for example, users should back up the data before a system upgrade.

Troubleshooting: If data is missing in the indexes of unlogged tables due to an abnormal shutdown or other unexpected operations, users should rebuild indexes with errors.

- **GLOBAL | LOCAL**

When creating a temporary table, you can specify the **GLOBAL** or **LOCAL** keyword before **TEMP** or **TEMPORARY**. If the keyword **GLOBAL** is specified, GaussDB creates a global temporary table. Otherwise, GaussDB creates a local temporary table.

- **TEMPORARY | TEMP**

If **TEMP** or **TEMPORARY** is specified, the created table is a temporary table. Temporary tables are classified into global temporary tables and local temporary tables. If the keyword **GLOBAL** is specified when a temporary table is created, the table is a global temporary table. Otherwise, the table is a local temporary table.

The metadata of the global temporary table is visible to all sessions. After the sessions end, the metadata still exists. The user data, indexes, and statistics of a session are isolated from those of another session. Each session can only view and modify the data committed by itself. Global temporary tables can be created using **ON COMMIT PRESERVE ROWS** or **ON COMMIT DELETE ROWS**. The former creates session-level tables, where user data is automatically cleared when a session ends; the latter creates transaction-level tables, where user data is automatically cleared when a commit or rollback operation is performed. If the **ON COMMIT** option is not specified during table creation, the session level is used by default. Different from local temporary tables, you can specify a schema that does not start with **pg\_temp\_** when creating a global temporary table.

A local temporary table is automatically dropped at the end of the current session. Therefore, you can create and use temporary tables in the current session as long as the connected database node in the session is normal. Temporary tables are created only in the current session. If a DDL statement involves operations on temporary tables, a DDL error will be generated. Therefore, you are advised not to perform operations on temporary tables in DDL statements. **TEMP** is equivalent to **TEMPORARY**.

**NOTICE**

- Local temporary tables are visible to the current session through the schema starting with **pg\_temp**. Users should not delete schemas starting with **pg\_temp** or **pg\_toast\_temp**.
  - If **TEMPORARY** or **TEMP** is not specified when you create a table but its schema is set to that starting with **pg\_temp\_** in the current session, the table will be created as a temporary table.
  - If global temporary tables and indexes are being used by other sessions, do not perform **ALTER** or **DROP** (except the **ALTER INDEX index\_name REBUILD** command).
  - The DDL of a global temporary table affects only the user data and indexes of the current session. For example, TRUNCATE, REINDEX, and ANALYZE are valid only for the current session.
  - You can set the GUC parameter **max\_active\_global\_temporary\_table** to determine whether to enable the global temporary table function. If **max\_active\_global\_temporary\_table** is set to **0**, the global temporary table function is disabled.
  - A temporary table is visible only to the current session. Therefore, it cannot be used together with **\parallel on**.
  - Temporary tables do not support switchover.
  - The global temporary table does not respond to automatic clearance. In persistent connection scenarios, you are advised to use the global temporary table in the ON COMMIT DELETE ROWS clause or periodically and manually execute the VACUUM statement. Otherwise, Clogs may not be recycled.
- 
- **IF NOT EXISTS**  
Sends a notice, but does not throw an error, if a table with the same name exists.
  - **table\_name**  
Specifies the name of the table to be created.

**NOTICE**

Some processing logic of materialized views determines whether a table is the log table of a materialized view or a table associated with a materialized view based on the table name prefix. Therefore, do not create a table whose name prefix is **mlog\_** or **matviewmap\_**. Otherwise, some functions of the table are affected.

- 
- **column\_name**  
Specifies the name of a column to be created in the new table.
  - **data\_type**  
Specifies the data type of the column.
  - **COLLATE collation**  
Assigns a collation to the column (which must be of a collatable data type). If no collation is specified, the default collation is used. You can run the **select \***

**from pg\_collation;** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.

- **LIKE source\_table [ like\_option ... ]**

Specifies a table from which the new table automatically copies all column names, their data types, and their NOT NULL constraints.

The new table and the source table are decoupled after creation is complete. Changes to the source table will not be applied to the new table, and it is not possible to include data of the new table in scans of the source table.

The copied columns and constraints are not merged with similarly named columns and constraints. If the same name is specified explicitly or in another LIKE clause, an error is reported.

- The default expressions are copied from the source table to the new table only if **INCLUDING DEFAULTS** is specified. The default behavior is to exclude default expressions, resulting in the copied columns in the new table having default values **NULL**.
- If **INCLUDING GENERATED** is specified, the generated expression of the source table column is copied to the new table. By default, the generated expression is not copied.
- The CHECK constraints are copied from the source table to the new table only when **INCLUDING CONSTRAINTS** is specified. Other types of constraints are never copied to the new table. NOT NULL constraints are always copied to the new table. These rules also apply to column constraints and table constraints.
- Any indexes on the source table will not be created on the new table, unless the **INCLUDING INDEXES** clause is specified.
- **STORAGE** settings for the copied column definitions are copied only if **INCLUDING STORAGE** is specified. The default behavior is to exclude **STORAGE** settings.
- If **INCLUDING COMMENTS** is specified, comments for the columns, constraints, and indexes of the source table are copied. The default behavior is to exclude comments.
- If **INCLUDING PARTITION** is specified, the partition definitions of the source table are copied to the new table, and the new table no longer uses the PARTITION BY clause. The default behavior is to exclude partition definition of the source table. If the source table has an index, you can use the INCLUDING PARTITION INCLUDING INDEXES syntax. If only INCLUDING INDEXES is used for a partitioned table, the target table will be defined as an ordinary table, but the index is a partitioned index. In this case, an error will be reported because ordinary tables do not support partitioned indexes.
- If **INCLUDING REOPTIONS** is specified, the new table will copy the storage parameter (that is, **WITH** clause) of the source table. The default behavior is to exclude partition definition of the storage parameter of the source table.
- **INCLUDING ALL** contains the meaning of **INCLUDING DEFAULTS**, **INCLUDING GENERATED**, **INCLUDING CONSTRAINTS**, **INCLUDING INDEXES**, **INCLUDING STORAGE**, **INCLUDING COMMENTS**, **INCLUDING PARTITION**, and **INCLUDING REOPTIONS**.



## NOTICE

- If the source table contains a sequence with the serial, bigserial, smallserial, or largeserial data type, or a column in the source table is a sequence by default and the sequence is created for this table by using **CREATE SEQUENCE ... OWNED BY**, these sequences will not be copied to the new table, and another sequence specific to the new table will be created. This is different from earlier versions. To share a sequence between the source table and new table, create a shared sequence (do not use **OWNED BY**) and set a column in the source table to this sequence.
- You are advised not to set a column in the source table to the sequence specific to another table especially when the table is distributed in specific node groups, because doing so may result in **CREATE TABLE ... LIKE** execution failures. In addition, doing so may cause the sequence to become invalid in the source sequence because the sequence will also be deleted from the source table when it is deleted from the table that the sequence is specific to. To share a sequence among multiple tables, you are advised to create a shared sequence for them.
- **EXCLUDING** of a partitioned table must be used together with **INCLUDING ALL**, for example, **INCLUDING ALL EXCLUDING DEFAULTS**, except for **DEFAULTS** of the source partitioned table.

## • **WITH ( { storage\_parameter = value } [, ... ] )**

Specifies an optional storage parameter for a table or an index. The **WITH** clause for a table can contain **OIDS=TRUE** or **OIDS** to specify that each row in the new table is assigned an OID. If **OIDS=FALSE** is specified, no OID is assigned.

### NOTE

When using the numeric type of arbitrary precision to define a column, specify precision **p** and scale **s**. When precision and scale are not specified, the input will be displayed.

The description of parameters is as follows:

#### – **FILLFACTOR**

The fill factor of a table is a percentage from 10 to 100. **100** (complete filling) is the default value. When a smaller fill factor is specified, **INSERT** operations fill table pages only to the indicated percentage. The remaining space on each page is reserved for updating rows on that page. This gives **UPDATE** a chance to place the updated copy of a row on the same page, which is more efficient than placing it on a different page. For a table whose entries are never updated, setting the fill factor to **100** (complete filling) is the best choice, but in heavily updated tables a smaller fill factor would be appropriate.

Value range: 10–100.

#### – **ORIENTATION**

Specifies the storage mode (row-store) of table data. This parameter cannot be modified once it is set.

Value range:

- **ROW** indicates that table data is stored in rows.

Row store applies to OLTP service and scenarios with a large number of point queries or addition/deletion operations.

Default value:

If an ordinary tablespace is specified, the default is **ROW**.

- **STORAGE\_TYPE**

Specifies the storage engine type. This parameter cannot be modified once it is set.

Value range:

- **USTORE** indicates that tables support the inplace-update storage engine. Note that the **track\_counts** and **track\_activities** parameters must be enabled when the Ustore table is used. Otherwise, space bloating may occur.

---

 **CAUTION**

Currently, Ustore does not support the ultimate RTO replay mode. For a primary node, if **recovery\_parse\_workers** is set to a value greater than 1, an error is reported when Ustore tables are created. For a standby node, if the database contains Ustore tables, replay may fail and an error may be reported if the ultimate RTO function is enabled. In serious cases, data on the standby node may be damaged. In this case, you need to rebuild the standby node.

---

- **ASTORE** indicates that tables support the append-only storage engine.

Default value:

If no table is specified, data is stored in append-only mode by default.

- **INIT\_TD**

Specifies the number of TDs to be initialized when a Ustore table is created. This parameter takes effect only when a Ustore table is created. You can run the **ALTER TABLE** command to change the parameter value.

Value ranges: 2–128. The default value is **4**.

- **segment**

The data is stored in segment-page mode. This parameter supports only row-store tables. Temporary tables and unlogged tables are not supported. Ustore is not supported.

Value range: **on** and **off**

Default value: **off**

- **parallel\_workers**

Number of bgworker threads started when an index is created. For example, value **2** indicates that two bgworker threads are started to create indexes concurrently.

Value range: [0,32], int type. The value **0** indicates that this function is disabled.

Default value: If this parameter is not set, the concurrent index creation function is disabled.

- hasuids

If this parameter is set to **on**, a unique table-level ID is allocated to a tuple when the tuple is updated.

Value range: **on** and **off**

Default value: **off**

- min\_tuples

The optimizer selects the larger value of the estimated statistics and the parameter to calculate the data volume based on the estimation table of statistics.

Value range:  $[0, DBL\_MAX)$

**Default value: 0**

- **WITHOUT OIDS**

It is equivalent to **WITH(OIDS=FALSE)**.

- **ON COMMIT { PRESERVE ROWS | DELETE ROWS | DROP }**

**ON COMMIT** determines what to do when you commit a temporary table creation operation. The three options are as follows (but only **PRESERVE ROWS** and **DELETE ROWS** can be used):

- **PRESERVE ROWS** (default): No special action is taken at the ends of transactions. The temporary table and its table data are unchanged.
- **DELETE ROWS**: All rows in the temporary table will be deleted at the end of each transaction block.
- **DROP**: The temporary table will be dropped at the end of the current transaction block. Only local temporary tables can be dropped. Global temporary tables cannot be dropped.

- **TABLESPACE tablespace\_name**

Specifies the tablespace where the new table is created. If it is not specified, the default tablespace is used.

- **CONSTRAINT constraint\_name**

Specifies the name of a column or table constraint. The optional constraint clauses specify constraints that new or updated rows must satisfy for an INSERT or UPDATE operation to succeed.

There are two ways to define constraints:

- A column constraint is defined as part of a column definition, and it is bound to a particular column.
- A table constraint is not bound to a particular column but can apply to more than one column.

- **NOT NULL**

Specifies that the column value cannot be **NULL**.

- **NULL**

Specifies that the column value can be **NULL**, which is the default value.

This clause is only provided for compatibility with non-standard SQL databases. It is not recommended.

- **CHECK ( expression )**

Specifies an expression producing a Boolean result where the INSERT or UPDATE operation of new or updated rows can succeed only when the expression result is true or unknown; otherwise, an error is thrown and the database is not altered.

A CHECK constraint specified as a column constraint should reference only the column's value, while an expression in a table constraint can reference multiple columns.

 **NOTE**

<>NULL and !=NULL are invalid in an expression. Change them to **is NOT NULL**.

- **DEFAULT default\_expr**

Assigns a default data value to a column. The value can be any variable-free expressions. (Subqueries and cross-references to other columns in the current table are not allowed.) The data type of the default expression must match that of the column.

The default expression will be used in any INSERT operation that does not specify a value for the column. If there is no default value for a column, then the default value is **NULL**.

- **GENERATED ALWAYS AS ( generation\_expr ) STORED**

This clause creates a column as a generated column. The value of the generated column is calculated by **generation\_expr** when data is written (inserted or updated). **STORED** indicates that the value of the generated column is stored as an ordinary column.

 **NOTE**

- The generation expression cannot refer to data other than the current row in any way. The generation expression cannot reference other generation columns or system columns. The generation expression cannot return a result set. No subquery, aggregate function, or window function can be used. The function called by the generation expression can only be an immutable function.
  - Default values cannot be specified for generated columns.
  - The generated column cannot be used as a part of the partition key.
  - Do not specify the generated column and the CASCADE, SET NULL, and SET DEFAULT actions of the ON UPDATE constraint at the same time. Do not specify the generated column and the SET NULL, and SET DEFAULT actions of the ON DELETE constraint at the same time.
  - The method of modifying and deleting generated columns is the same as that of ordinary columns. Delete the ordinary column that the generated column depends on. The generated column is automatically deleted. The type of the column on which the generated column depends cannot be changed.
  - The generated column cannot be directly written. In the INSERT or UPDATE statement, values cannot be specified for generated columns, but the keyword DEFAULT can be specified.
  - The permission control for generated columns is the same as that for ordinary columns.
- **UNIQUE index\_parameters**  
**UNIQUE ( column\_name [, ... ] ) index\_parameters**  
Specifies that a group of one or more columns of a table can contain only unique values.

For the purpose of a unique constraint, null is not considered equal.

- **PRIMARY KEY** *index\_parameters*

**PRIMARY KEY** ( *column\_name* [, ... ] ) *index\_parameters*

Specifies that a column or columns of a table can contain only unique (non-duplicate) and non-NULL values.

Only one primary key can be specified for a table.

- **REFERENCES** *reftable* [ ( *refcolumn* ) ] [ **MATCH** *matchtype* ] [ **ON DELETE** *action* ] [ **ON UPDATE** *action* ] ( *column constraint* )

**FOREIGN KEY** ( *column\_name* [, ... ] ) **REFERENCES** *reftable* [ ( *refcolumn* [, ... ] ) ] [ **MATCH** *matchtype* ] [ **ON DELETE** *action* ] [ **ON UPDATE** *action* ] ( *table constraint* )

The foreign key constraint requires that the group consisting of one or more columns in the new table should contain and match only the referenced column values in the referenced table. If **refcolumn** is omitted, the primary key of **reftable** is used. The referenced column should be the only column or primary key in the referenced table. A foreign key constraint cannot be defined between a temporary table and a permanent table.

There are three types of matching between a reference column and a referenced column:

- **MATCH FULL**: A column with multiple foreign keys cannot be NULL unless all foreign key columns are NULL.
- **MATCH SIMPLE** (default): Any unexpected foreign key column can be NULL.
- **MATCH PARTIAL**: This option is not supported currently.

In addition, when certain operations are performed on the data in the referenced table, the operations are performed on the corresponding columns in the new table. The ON DELETE clause specifies the operations to be executed after a referenced row in the referenced table is deleted. The ON UPDATE clause specifies the operation to be performed when the referenced column data in the referenced table is updated. Possible responses to the ON DELETE and ON UPDATE clauses are as follows:

- **NO ACTION** (default): When a foreign key is deleted or updated, an error indicating that the foreign key constraint is violated is reported. If the constraint is deferrable and there are still any referenced rows, this error will occur when the constraint is checked.
- **RESTRICT**: When a foreign key is deleted or updated, an error indicating that the foreign key constraint is violated is created. It is the same as **NO ACTION** except that the constraint is not deferrable.
- **CASCADE**: deletes any row that references the deleted row from the new table, or update the column value of the referenced row in the new table to the new value of the referenced column.
- **SET NULL**: sets the referenced columns to NULL.
- **SET DEFAULT**: sets the referenced columns to their default values.

- **DEFERRABLE | NOT DEFERRABLE**

Determines whether the constraint can be deferred. A constraint that is not deferrable will be checked immediately after every command. Checking of constraints that are deferrable can be postponed until the end of the

transaction using the **SET CONSTRAINTS** command. **NOT DEFERRABLE** is the default value. Currently, only UNIQUE constraints, primary key constraints, and foreign key constraints accept this clause. All the other constraints are not deferrable.

#### NOTE

Ustore tables do not support the keywords **DEFERRABLE** and **INITIALLY DEFERRED**.

- **INITIALLY IMMEDIATE | INITIALLY DEFERRED**

If a constraint is deferrable, this clause specifies the default time to check the constraint.

- If the constraint is **INITIALLY IMMEDIATE** (default value), it is checked after each statement.
- If the constraint is **INITIALLY DEFERRED**, it is checked only at the end of the transaction.

The constraint check time can be altered by executing the **SET CONSTRAINTS** command.

- **USING INDEX TABLESPACE tablespace\_name**

Specifies a tablespace in which an index associated with a UNIQUE or PRIMARY KEY constraint will be created. If this clause is not used, such index will be created in **default\_tablespace**. If **default\_tablespace** is empty, the default tablespace of the database is used.

## Examples

```
-- Create a simple table.
openGauss=# CREATE TABLE tpcds.warehouse_t1
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);

openGauss=# CREATE TABLE tpcds.warehouse_t2
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);
```

```
-- Create a table and set the default value of the W_STATE column to GA.
openGauss=# CREATE TABLE tpcds.warehouse_t3
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) DEFAULT 'GA',
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create a table and check whether the W_WAREHOUSE_NAME column is unique at the end of its
creation.
openGauss=# CREATE TABLE tpcds.warehouse_t4
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) UNIQUE DEFERRABLE,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create a table with its fill factor set to 70%.
openGauss=# CREATE TABLE tpcds.warehouse_t5
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2),
 UNIQUE(W_WAREHOUSE_NAME) WITH(fillfactor=70)
);

-- Alternatively, user the following syntax:
openGauss=# CREATE TABLE tpcds.warehouse_t6
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) UNIQUE,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
```

```
W_CITY VARCHAR(60)
W_COUNTY VARCHAR(30)
W_STATE CHAR(2)
W_ZIP CHAR(10)
W_COUNTRY VARCHAR(20)
W_GMT_OFFSET DECIMAL(5,2)
) WITH(fillfactor=70);

-- Create a table and specify that its data is not written to WALs.
openGauss=# CREATE UNLOGGED TABLE tpcds.warehouse_t7
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create a temporary table.
openGauss=# CREATE TEMPORARY TABLE warehouse_t24
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create a local temporary table and specify that this table is dropped when the transaction is committed.
openGauss=# CREATE TEMPORARY TABLE warehouse_t25
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
) ON COMMIT DELETE ROWS;

-- Create a global temporary table and specify that the temporary table data is deleted when the session
ends. The current Ustore does not support global temporary tables.
openGauss=# CREATE GLOBAL TEMPORARY TABLE gtt1
(
```



```
ID INTEGER NOT NULL,
NAME CHAR(16) NOT NULL,
ADDRESS VARCHAR(50)
POSTCODE CHAR(6)
) ON COMMIT PRESERVE ROWS;

-- Create a table and specify that no error is reported for duplicate tables (if any).
openGauss=# CREATE TABLE IF NOT EXISTS tpcds.warehouse_t8
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create an ordinary tablespace.
openGauss=# CREATE TABLESPACE DS_TABLESPACE1 RELATIVE LOCATION 'tablespace/tablespace_1';
-- Specify a tablespace when creating a table.
openGauss=# CREATE TABLE tpcds.warehouse_t9
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
) TABLESPACE DS_TABLESPACE1;

-- Separately specify the index tablespace for W_WAREHOUSE_NAME when creating the table.
openGauss=# CREATE TABLE tpcds.warehouse_t10
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) UNIQUE USING INDEX TABLESPACE
DS_TABLESPACE1,
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
);
-- Create a table with a primary key constraint.
openGauss=# CREATE TABLE tpcds.warehouse_t11
(
 W_WAREHOUSE_SK INTEGER PRIMARY KEY,
```

```
W_WAREHOUSE_ID CHAR(16) NOT NULL,
W_WAREHOUSE_NAME VARCHAR(20)
W_WAREHOUSE_SQ_FT INTEGER
W_STREET_NUMBER CHAR(10)
W_STREET_NAME VARCHAR(60)
W_STREET_TYPE CHAR(15)
W_SUITE_NUMBER CHAR(10)
W_CITY VARCHAR(60)
W_COUNTY VARCHAR(30)
W_STATE CHAR(2)
W_ZIP CHAR(10)
W_COUNTRY VARCHAR(20)
W_GMT_OFFSET DECIMAL(5,2)
);

-- An alternative for the preceding syntax is as follows:
openGauss=# CREATE TABLE tpcds.warehouse_t12
(
W_WAREHOUSE_SK INTEGER NOT NULL,
W_WAREHOUSE_ID CHAR(16) NOT NULL,
W_WAREHOUSE_NAME VARCHAR(20)
W_WAREHOUSE_SQ_FT INTEGER
W_STREET_NUMBER CHAR(10)
W_STREET_NAME VARCHAR(60)
W_STREET_TYPE CHAR(15)
W_SUITE_NUMBER CHAR(10)
W_CITY VARCHAR(60)
W_COUNTY VARCHAR(30)
W_STATE CHAR(2)
W_ZIP CHAR(10)
W_COUNTRY VARCHAR(20)
W_GMT_OFFSET DECIMAL(5,2),
PRIMARY KEY(W_WAREHOUSE_SK)
);

-- Or use the following statement to specify the name of the constraint:
openGauss=# CREATE TABLE tpcds.warehouse_t13
(
W_WAREHOUSE_SK INTEGER NOT NULL,
W_WAREHOUSE_ID CHAR(16) NOT NULL,
W_WAREHOUSE_NAME VARCHAR(20)
W_WAREHOUSE_SQ_FT INTEGER
W_STREET_NUMBER CHAR(10)
W_STREET_NAME VARCHAR(60)
W_STREET_TYPE CHAR(15)
W_SUITE_NUMBER CHAR(10)
W_CITY VARCHAR(60)
W_COUNTY VARCHAR(30)
W_STATE CHAR(2)
W_ZIP CHAR(10)
W_COUNTRY VARCHAR(20)
W_GMT_OFFSET DECIMAL(5,2),
CONSTRAINT W_CSTR_KEY1 PRIMARY KEY(W_WAREHOUSE_SK)
);

-- Create a table with a compound primary key constraint.
openGauss=# CREATE TABLE tpcds.warehouse_t14
(
W_WAREHOUSE_SK INTEGER NOT NULL,
W_WAREHOUSE_ID CHAR(16) NOT NULL,
W_WAREHOUSE_NAME VARCHAR(20)
W_WAREHOUSE_SQ_FT INTEGER
W_STREET_NUMBER CHAR(10)
W_STREET_NAME VARCHAR(60)
W_STREET_TYPE CHAR(15)
W_SUITE_NUMBER CHAR(10)
W_CITY VARCHAR(60)
W_COUNTY VARCHAR(30)
W_STATE CHAR(2)
```

```
W_ZIP CHAR(10) ,
W_COUNTRY VARCHAR(20) ,
W_GMT_OFFSET DECIMAL(5,2),
CONSTRAINT W_CSTR_KEY2 PRIMARY KEY(W_WAREHOUSE_SK, W_WAREHOUSE_ID)
);

-- Define a column check constraint.
openGauss=# CREATE TABLE tpcds.warehouse_t19
(
 W_WAREHOUSE_SK INTEGER PRIMARY KEY CHECK (W_WAREHOUSE_SK > 0),
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) CHECK (W_WAREHOUSE_NAME IS NOT NULL),
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);

openGauss=# CREATE TABLE tpcds.warehouse_t20
(
 W_WAREHOUSE_SK INTEGER PRIMARY KEY,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) CHECK (W_WAREHOUSE_NAME IS NOT NULL),
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) ,
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2),
 CONSTRAINT W_CONSTR_KEY2 CHECK(W_WAREHOUSE_SK > 0 AND W_WAREHOUSE_NAME IS NOT
NULL)
);

-- Create a table with a foreign key constraint.
openGauss=# CREATE TABLE tpcds.city_t23
(
 W_CITY VARCHAR(60) PRIMARY KEY,
 W_ADDRESS TEXT
);

openGauss=# CREATE TABLE tpcds.warehouse_t23
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER ,
 W_STREET_NUMBER CHAR(10) ,
 W_STREET_NAME VARCHAR(60) ,
 W_STREET_TYPE CHAR(15) ,
 W_SUITE_NUMBER CHAR(10) ,
 W_CITY VARCHAR(60) REFERENCES tpcds.city_t23(W_CITY),
 W_COUNTY VARCHAR(30) ,
 W_STATE CHAR(2) ,
 W_ZIP CHAR(10) ,
 W_COUNTRY VARCHAR(20) ,
 W_GMT_OFFSET DECIMAL(5,2)
);
```

```
-- An alternative for the preceding syntax is as follows:
openGauss=# CREATE TABLE tpcds.warehouse_t23
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 ,
 W_WAREHOUSE_SQ_FT INTEGER
 ,
 W_STREET_NUMBER CHAR(10)
 ,
 W_STREET_NAME VARCHAR(60)
 ,
 W_STREET_TYPE CHAR(15)
 ,
 W_SUITE_NUMBER CHAR(10)
 ,
 W_CITY VARCHAR(60)
 ,
 W_COUNTY VARCHAR(30)
 ,
 W_STATE CHAR(2)
 ,
 W_ZIP CHAR(10)
 ,
 W_COUNTRY VARCHAR(20)
 ,
 W_GMT_OFFSET DECIMAL(5,2)
 ,
 FOREIGN KEY(W_CITY) REFERENCES tpcds.city_t23(W_CITY)
);

-- Or use the following statement to specify the name of the constraint:
openGauss=# CREATE TABLE tpcds.warehouse_t23
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 ,
 W_WAREHOUSE_SQ_FT INTEGER
 ,
 W_STREET_NUMBER CHAR(10)
 ,
 W_STREET_NAME VARCHAR(60)
 ,
 W_STREET_TYPE CHAR(15)
 ,
 W_SUITE_NUMBER CHAR(10)
 ,
 W_CITY VARCHAR(60)
 ,
 W_COUNTY VARCHAR(30)
 ,
 W_STATE CHAR(2)
 ,
 W_ZIP CHAR(10)
 ,
 W_COUNTRY VARCHAR(20)
 ,
 W_GMT_OFFSET DECIMAL(5,2)
 ,
 CONSTRAINT W_FORE_KEY1 FOREIGN KEY(W_CITY) REFERENCES tpcds.city_t23(W_CITY)
);

-- Add a varchar column to the tpcds.warehouse_t19 table.
openGauss=# ALTER TABLE tpcds.warehouse_t19 ADD W_GOODS_CATEGORY varchar(30);

-- Add a CHECK constraint to the tpcds.warehouse_t19 table.
openGauss=# ALTER TABLE tpcds.warehouse_t19 ADD CONSTRAINT W_CONSTR_KEY4 CHECK (W_STATE IS NOT NULL);

-- Use one statement to alter the types of two existing columns.
openGauss=# ALTER TABLE tpcds.warehouse_t19
 ALTER COLUMN W_GOODS_CATEGORY TYPE varchar(80),
 ALTER COLUMN W_STREET_NAME TYPE varchar(100);

-- This statement is equivalent to the preceding statement.
openGauss=# ALTER TABLE tpcds.warehouse_t19 MODIFY (W_GOODS_CATEGORY varchar(30),
W_STREET_NAME varchar(60));

-- Add a NOT NULL constraint to an existing column.
openGauss=# ALTER TABLE tpcds.warehouse_t19 ALTER COLUMN W_GOODS_CATEGORY SET NOT NULL;

-- Remove NOT NULL constraints from an existing column.
openGauss=# ALTER TABLE tpcds.warehouse_t19 ALTER COLUMN W_GOODS_CATEGORY DROP NOT NULL;

-- Move a table to another tablespace.
openGauss=# ALTER TABLE tpcds.warehouse_t19 SET TABLESPACE PG_DEFAULT;
-- Create the joe schema.
openGauss=# CREATE SCHEMA joe;

-- Move a table to another schema.
openGauss=# ALTER TABLE tpcds.warehouse_t19 SET SCHEMA joe;
```

```
-- Rename an existing table.
openGauss=# ALTER TABLE joe.warehouse_t19 RENAME TO warehouse_t23;

-- Drop a column from the warehouse_t23 table.
openGauss=# ALTER TABLE joe.warehouse_t23 DROP COLUMN W_STREET_NAME;

-- Drop the tablespace, schema joe, and schema tables warehouse.
openGauss=# DROP TABLE tpceds.warehouse_t1;
openGauss=# DROP TABLE tpceds.warehouse_t2;
openGauss=# DROP TABLE tpceds.warehouse_t3;
openGauss=# DROP TABLE tpceds.warehouse_t4;
openGauss=# DROP TABLE tpceds.warehouse_t5;
openGauss=# DROP TABLE tpceds.warehouse_t6;
openGauss=# DROP TABLE tpceds.warehouse_t7;
openGauss=# DROP TABLE tpceds.warehouse_t8;
openGauss=# DROP TABLE tpceds.warehouse_t9;
openGauss=# DROP TABLE tpceds.warehouse_t10;
openGauss=# DROP TABLE tpceds.warehouse_t11;
openGauss=# DROP TABLE tpceds.warehouse_t12;
openGauss=# DROP TABLE tpceds.warehouse_t13;
openGauss=# DROP TABLE tpceds.warehouse_t14;
openGauss=# DROP TABLE tpceds.warehouse_t15;
openGauss=# DROP TABLE tpceds.warehouse_t16;
openGauss=# DROP TABLE tpceds.warehouse_t17;
openGauss=# DROP TABLE tpceds.warehouse_t18;
openGauss=# DROP TABLE tpceds.warehouse_t20;
openGauss=# DROP TABLE tpceds.warehouse_t21;
openGauss=# DROP TABLE tpceds.warehouse_t22;
openGauss=# DROP TABLE joe.warehouse_t23;
openGauss=# DROP TABLE tpceds.warehouse_t24;
openGauss=# DROP TABLE tpceds.warehouse_t25;
openGauss=# DROP TABLESPACE DS_TABLESPACE1;
openGauss=# DROP SCHEMA IF EXISTS joe CASCADE;
```

## Helpful Links

[ALTER TABLE](#), [DROP TABLE](#), and [CREATE TABLESPACE](#)

## Suggestions

- UNLOGGED
  - The unlogged table and its indexes do not use the WAL mechanism during data writing. Their write speed is much higher than that of ordinary tables. Therefore, they can be used for storing intermediate result sets of complex queries to improve query performance.
  - The unlogged table does not have the primary/standby mechanism. If the database is restarted due to a system fault or abnormal breakpoint, data in the unlogged table will be deleted, which may cause data loss. Therefore, the unlogged table cannot be used to store basic data.
- TEMPORARY | TEMP
  - A temporary table is automatically dropped at the end of the current session.
- LIKE
  - The new table automatically inherits all column names, data types, and NOT NULL constraints from this table. The new table is irrelevant to the source table after the creation.
- LIKE INCLUDING DEFAULTS
  - The default expressions are copied from the source table to the new table only if **INCLUDING DEFAULTS** is specified. The default behavior is to

exclude default expressions, resulting in the copied columns in the new table having default values **NULL**.

- **LIKE INCLUDING CONSTRAINTS**
  - The CHECK constraints are copied from the source table to the new table only when **INCLUDING CONSTRAINTS** is specified. Other types of constraints are never copied to the new table. NOT NULL constraints are always copied to the new table. These rules also apply to column constraints and table constraints.
- **LIKE INCLUDING INDEXES**
  - Any indexes on the source table will not be created on the new table, unless the **INCLUDING INDEXES** clause is specified.
- **LIKE INCLUDING STORAGE**
  - **STORAGE** settings for the copied column definitions are copied only if **INCLUDING STORAGE** is specified. The default behavior is to exclude **STORAGE** settings.
- **LIKE INCLUDING COMMENTS**
  - If **INCLUDING COMMENTS** is specified, comments for the columns, constraints, and indexes of the source table are copied. The default behavior is to exclude comments.
- **LIKE INCLUDING PARTITION**
  - If **INCLUDING PARTITION** is specified, the partition definitions of the source table are copied to the new table, and the new table no longer uses the PARTITION BY clause. The default behavior is to exclude partition definition of the source table.

---

#### NOTICE

List and hash partitioned tables do not support LIKE INCLUDING PARTITION.

---

- **LIKE INCLUDING REOPTIONS**
  - If **INCLUDING REOPTIONS** is specified, the new table will copy the storage parameter (that is, WITH clause) of the source table. The default behavior is to exclude partition definition of the storage parameter of the source table.
- **LIKE INCLUDING ALL**
  - **INCLUDING ALL** contains the meaning of **INCLUDING DEFAULTS**, **INCLUDING CONSTRAINTS**, **INCLUDING INDEXES**, **INCLUDING STORAGE**, **INCLUDING COMMENTS**, **INCLUDING PARTITION**, and **INCLUDING REOPTIONS**.
- **ORIENTATION ROW**
  - Creates a row-store table. Row-store applies to the OLTP service, which has many interactive transactions. An interaction involves many columns in the table. Using row-store can improve the efficiency.

## 7.13.72 CREATE TABLE AS

### Description

Creates a table from the results of a query.

CREATE TABLE AS creates a table and fills it with data obtained using SELECT. The table columns have the names and data types associated with the output columns of SELECT (except that you can override the SELECT output column names by giving an explicit list of new column names).

CREATE TABLE AS queries a source table once and writes the data in a new table. The result in the query view changes with the source table. In contrast, the view re-computes and defines its SELECT statement at each query.

### Precautions

- This statement cannot be used to create a partitioned table.
- If an error occurs during table creation, after it is fixed, the system may fail to delete the disk files that are created before the last automatic clearance and whose size is not 0. This problem seldom occurs and does not affect system running of the database.

### Syntax

```
CREATE [UNLOGGED] TABLE table_name
[(column_name [, ...])]
[WITH ({storage_parameter = value} [, ...])]
[TABLESPACE tablespace_name]
AS query
[WITH [NO] DATA];
```

### Parameters

- **UNLOGGED**

Specifies that the table is created as an unlogged table. Data written to unlogged tables is not written to the WALs, which makes them considerably faster than ordinary tables. However, it is also insecure. After the database is restarted due to a conflict or abnormal shutdown, the data in the unlogged table is cleared. Contents of an unlogged table are also not replicated to standby nodes. Any indexes created on an unlogged table are automatically unlogged as well.

- Usage scenario: Unlogged tables do not ensure data security. Users can back up data before using unlogged tables; for example, users should back up the data before a system upgrade.
- Troubleshooting: If data is missing in the indexes of unlogged tables due to an abnormal shutdown or other unexpected operations, users should rebuild indexes with errors.

- **GLOBAL | LOCAL**

When creating a temporary table, you can specify the **GLOBAL** or **LOCAL** keyword before **TEMP** or **TEMPORARY**. If the keyword **GLOBAL** is specified, GaussDB creates a global temporary table. Otherwise, GaussDB creates a local temporary table.

- **TEMPORARY | TEMP**

If **TEMP** or **TEMPORARY** is specified, the created table is a temporary table. Temporary tables are classified into global temporary tables and local temporary tables. If the keyword **GLOBAL** is specified when a temporary table is created, the table is a global temporary table. Otherwise, the table is a local temporary table.

The metadata of the global temporary table is visible to all sessions. After the sessions end, the metadata still exists. The user data, indexes, and statistics of a session are isolated from those of another session. Each session can only view and modify the data committed by itself. Global temporary tables can be created using **ON COMMIT PRESERVE ROWS** or **ON COMMIT DELETE ROWS**. The former creates session-level tables, where user data is automatically cleared when a session ends; the latter creates transaction-level tables, where user data is automatically cleared when a commit or rollback operation is performed. If the **ON COMMIT** option is not specified during table creation, the session level is used by default. Different from local temporary tables, you can specify a schema that does not start with **pg\_temp\_** when creating a global temporary table.

A local temporary table is automatically dropped at the end of the current session. Therefore, you can create and use temporary tables in the current session as long as the connected database node in the session is normal. Temporary tables are created only in the current session. If a DDL statement involves operations on temporary tables, a DDL error will be generated. Therefore, you are advised not to perform operations on temporary tables in DDL statements. **TEMP** is equivalent to **TEMPORARY**.

---

**NOTICE**

- Local temporary tables are visible to the current session through the schema starting with **pg\_temp**. Users should not delete schemas starting with **pg\_temp** or **pg\_toast\_temp**.
- If **TEMPORARY** or **TEMP** is not specified when you create a table but its schema is set to that starting with **pg\_temp\_** in the current session, the table will be created as a temporary table.
- If global temporary tables or indexes are being used by other sessions, do not perform **ALTER** or **DROP** on the tables or indexes.
- The DDL of a global temporary table affects only the user data and indexes of the current session. For example, **TRUNCATE**, **REINDEX**, and **ANALYZE** are valid only for the current session.

---

- **table\_name**

Specifies the name of the table to be created.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **column\_name**

Specifies the name of a column to be created in the new table.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **WITH ( storage\_parameter [= value] [, ... ] )**

Specifies an optional storage parameter for a table or an index. See details of parameters below.



- **FILLFACTOR**  
The fill factor of a table is a percentage from 10 to 100. **100** (complete filling) is the default value. When a smaller fill factor is specified, INSERT operations fill table pages only to the indicated percentage. The remaining space on each page is reserved for updating rows on that page. This gives UPDATE a chance to place the updated copy of a row on the same page, which is more efficient than placing it on a different page. For a table whose entries are never updated, setting the fill factor to **100** (complete filling) is the best choice, but in heavily updated tables a smaller fill factor would be appropriate. The parameter is only valid for row-store tables.  
Value range: 10–100.
- **ORIENTATION**  
Value range:  
**COLUMN**: The data will be stored in columns.  
**ROW** (default value): The data will be stored in rows.
- **ON COMMIT { PRESERVE ROWS | DELETE ROWS | DROP }**  
**ON COMMIT** determines what to do when you commit a temporary table creation operation. The three options are as follows (but only **PRESERVE ROWS** and **DELETE ROWS** can be used):
  - **PRESERVE ROWS** (default): No special action is taken at the ends of transactions. The temporary table and its table data are unchanged.
  - **DELETE ROWS**: All rows in the temporary table will be deleted at the end of each transaction block.
  - **DROP**: The temporary table will be dropped at the end of the current transaction block. Only local temporary tables can be dropped. Global temporary tables cannot be dropped.
- **TABLESPACE tablespace\_name**  
Specifies that the new table will be created in the **tablespace\_name** tablespace. If it is not specified, the default tablespace is used.
- **AS query**  
Specifies a **SELECT** or **VALUES** command, or an **EXECUTE** command that runs a prepared **SELECT** or **VALUES** query.
- **[ WITH [ NO ] DATA ]**  
Specifies whether the data produced by the query should be copied to the new table. By default, the data will be copied. If the value **NO** is used, only the table structure will be copied.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.store_returns table.
openGauss=# CREATE TABLE tpcds.store_returns
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 sr_item_sk VARCHAR(20) ,
 W_WAREHOUSE_SQ_FT INTEGER
);
```

```
-- Insert a record into a table.
openGauss=# INSERT INTO tpcds.store_returns(W_WAREHOUSE_SK, W_WAREHOUSE_ID, sr_item_sk,
W_WAREHOUSE_SQ_FT) VALUES (1,
'AAAAAAAAABAAAAAA', '4800', '20');

-- Create the tpcds.store_returns_t1 table and insert numbers that are greater than 4795 into the
sr_item_sk column of the tpcds.store_returns table.
openGauss=# CREATE TABLE tpcds.store_returns_t1 AS SELECT * FROM tpcds.store_returns WHERE
sr_item_sk > '4795';

-- Copy tpcds.store_returns to create the tpcds.store_returns_t2 table.
openGauss=# CREATE TABLE tpcds.store_returns_t2 AS table tpcds.store_returns;

-- Drop the table.
openGauss=# DROP TABLE tpcds.store_returns_t1 ;
openGauss=# DROP TABLE tpcds.store_returns_t2 ;
openGauss=# DROP TABLE tpcds.store_returns;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[CREATE TABLE](#) and [SELECT](#)

### 7.13.73 CREATE TABLE PARTITION

#### Description

Creates a partitioned table. Partitioning refers to splitting what is logically one large table into smaller physical pieces based on specific schemes. The table based on the logic is called a partitioned table, and each physical piece is called a partition. Data is stored on these physical partitions, instead of the logical partitioned table.

The common forms of partitioning include range partitioning, interval partitioning, hash partitioning, list partitioning, and value partitioning. Currently, row-store tables support range partitioning, interval partitioning, hash partitioning, and list partitioning.

In range partitioning, a table is partitioned based on ranges defined by one or more columns, with no overlap between the ranges of values assigned to different partitions. Each range has a dedicated partition for data storage.

The partitioning policy for range partitioning refers to how data is inserted into partitions. Currently, range partitioning only allows the use of the range partitioning policy.

In range partitioning, a table is partitioned based on partition key values. If a record can be mapped to a partition, it is inserted into the partition; if it cannot, an error message is returned. Range partitioning is the most commonly used partitioning policy.

Interval partitioning is a special type of range partitioning. Compared with range partitioning, interval value definition is added. When no matching partition can be found for an inserted record, a partition can be automatically created based on the interval value.

Interval partitioning supports only table-based partitioning of a list where the data type can be `TIMESTAMP[(p)] [WITHOUT TIME ZONE]`, `TIMESTAMP[(p)] [WITH TIME ZONE]` and `DATE`.

Interval partitioning policy: A record is mapped to a created partition based on the partition key value. If the record can be mapped to a created partition, the record is inserted into the corresponding partition. Otherwise, a partition is automatically created based on the partition key value and table definition information, and then the record is inserted into the new partition. The data range of the new partition is equal to the interval value.

In hash partitioning, a modulus and a remainder are specified for each partition based on a column in the table, and records to be inserted into the table are allocated to the corresponding partition, the rows in each partition must meet the following condition: The value of the partition key divided by the specified modulus generates the remainder specified for the partition key.

In hash partitioning, table is partitioned based on partition key values. If a record can be mapped to a partition, it is inserted into the partition; if it cannot, an error message is returned.

List partitioning is to allocate the records to be inserted into a table to the corresponding partition based on the key values in each partition. The key values do not overlap in different partitions. Create a partition for each group of key values to store corresponding data.

In list partitioning, table is partitioned based on partition key values. If a record can be mapped to a partition, it is inserted into the partition; if it cannot, an error message is returned.

Partitioning can provide several benefits:

- Query performance can be improved drastically in certain situations, particularly when most of the heavily accessed rows of the table are in a single partition or a small number of partitions. Partitioning narrows the range of data search and improves data access efficiency.
- When querying or updating most of the records in a partition, continuously scanning that partition instead of accessing the entire table can achieve a significant performance improvement.
- Frequent loading or deletion operations on records in a separate partition can be accomplished by reading or deleting that partition. This not only improves performance but also avoids the `VACUUM` overload caused by bulk `DELETE` operations (hash partitions cannot be deleted).

## Precautions

- If the constraint key of the unique constraint and primary key constraint contains all partition keys, a local index is created for the constraints. Otherwise, a global index is created.
- Currently, hash partitioning and list partitioning support only single-column partitioning, and do not support multi-column partitioning.
- When you have the `INSERT` permission on an interval partitioned table, partitions can be automatically created when you run `INSERT` to write data to the table.

- In the PARTITION FOR (values) syntax for partitioned tables, values can only be constants.
- In the PARTITION FOR (values) syntax for partitioned tables, if data type conversion is required for values, you are advised to use forcible type conversion to prevent the implicit type conversion result from being inconsistent with the expected result.
- The maximum number of partitions is 1048575. Generally, it is impossible to create so many partitions, because too many partitions may cause insufficient memory. Create partitions based on the value of **local\_syscache\_threshold**. The memory used by the partitioned tables is about (number of partitions x 3/1024) MB. Theoretically, the memory occupied by the partitions cannot be greater than the value of **local\_syscache\_threshold**. In addition, some space must be reserved for other functions.
- Currently, the statement specifying a partition cannot perform global index scan.

## Syntax

```
CREATE TABLE [IF NOT EXISTS] partition_table_name
([
 { column_name data_type [COLLATE collation] [column_constraint [...]]
 | table_constraint
 | LIKE source_table [like_option [...]] }, ...]
)
[WITH ({storage_parameter = value} [, ...])]

[TABLESPACE tablespace_name]
PARTITION BY {
 {RANGE (partition_key) [INTERVAL ('interval_expr') [STORE IN (tablespace_name [, ...])]]
 (partition_less_than_item [, ...]) } |
 {RANGE (partition_key) [INTERVAL ('interval_expr') [STORE IN (tablespace_name [, ...])]]
 (partition_start_end_item [, ...]) } |
 {LIST (partition_key) (PARTITION partition_name VALUES (list_values) [TABLESPACE
tablespace_name][, ...]) } |
 {HASH (partition_key) (PARTITION partition_name [TABLESPACE tablespace_name][, ...]) }
} [{ ENABLE | DISABLE } ROW MOVEMENT];
```

- **column\_constraint:**

```
[CONSTRAINT constraint_name]
{ NOT NULL |
 NULL |
 CHECK (expression) |
 DEFAULT default_expr |
 GENERATED ALWAYS AS (generation_expr) STORED |
 UNIQUE index_parameters |
 PRIMARY KEY index_parameters |
 REFERENCES reftable [(refcolumn)] [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
 [ON DELETE action] [ON UPDATE action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```
- **table\_constraint:**

```
[CONSTRAINT constraint_name]
{ CHECK (expression) |
 UNIQUE (column_name [, ...]) index_parameters |
 PRIMARY KEY (column_name [, ...]) index_parameters |
 FOREIGN KEY (column_name [, ...]) REFERENCES reftable [(refcolumn [, ...])]
 [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE] [ON DELETE action] [ON UPDATE
action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```
- **like\_option:**

```
{ INCLUDING | EXCLUDING } { DEFAULTS | GENERATED | CONSTRAINTS | INDEXES | STORAGE |
COMMENTS | REOPTIONS | ALL }
```

- **index\_parameters:**  
[ WITH ( {storage\_parameter = value} [, ... ] ) ]  
[ USING INDEX TABLESPACE tablespace\_name ]
- **partition\_less\_than\_item:**  
PARTITION partition\_name VALUES LESS THAN ( { partition\_value | MAXVALUE } ) [TABLESPACE tablespace\_name]
- **partition\_start\_end\_item:**  
PARTITION partition\_name {  
    {START(partition\_value) END (partition\_value) EVERY (interval\_value)} |  
    {START(partition\_value) END ({partition\_value | MAXVALUE})} |  
    {START(partition\_value)} |  
    {END({partition\_value | MAXVALUE})}  
} [TABLESPACE tablespace\_name]

## Parameters

- **IF NOT EXISTS**  
Sends a notice instead of an error if tables have identical names. The notice prompts that the table relationship already exists.
- **partition\_table\_name**  
Specifies the name of a partitioned table.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Specifies the name of a column to be created in the new table.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **data\_type**  
Specifies the data type of the column.
- **COLLATE collation**  
Assigns a collation to the column (which must be of a collatable data type). If no collation is specified, the default collation is used. You can run the **select \* from pg\_collation;** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.
- **CONSTRAINT constraint\_name**  
Specifies the name of a column or table constraint. The optional constraint clauses specify constraints that new or updated rows must satisfy for an INSERT or UPDATE operation to succeed.  
There are two ways to define constraints:
  - A column constraint is defined as part of a column definition, and it is bound to a particular column.
  - A table constraint is not bound to a particular column but can apply to more than one column.
- **LIKE source\_table [ like\_option ... ]**  
Specifies a table from which the new table automatically copies all column names, their data types, and their NOT NULL constraints.  
Unlike INHERITS, the new table and source table are decoupled after creation is complete. Changes to the source table will not be applied to the new table, and it is not possible to include data of the new table in scans of the source table.

- Default expressions for the copied column definitions will be copied only if **INCLUDING DEFAULTS** is specified. The default behavior is to exclude default expressions, resulting in the copied columns in the new table having default values **NULL**.
  - If **INCLUDING GENERATED** is specified, the generated expression of the source table column is copied to the new table. By default, the generated expression is not copied.
  - NOT NULL constraints are always copied to the new table. CHECK constraints will only be copied if **INCLUDING CONSTRAINTS** is specified; other types of constraints will never be copied. These rules also apply to column constraints and table constraints.
  - Unlike those of **INHERITS**, the copied columns and constraints are not merged with similarly named columns and constraints. If the same name is specified explicitly or in another LIKE clause, an error is reported.
  - Any indexes on the source table will not be created on the new table, unless the **INCLUDING INDEXES** clause is specified.
  - **STORAGE** settings for the copied column definitions are copied only if **INCLUDING STORAGE** is specified. The default behavior is to exclude **STORAGE** settings.
  - If **INCLUDING COMMENTS** is specified, comments for the copied columns, constraints, and indexes are copied. The default behavior is to exclude comments.
  - If **INCLUDING REOPTIONS** is specified, the new table will copy the storage parameter (that is, WITH clause) of the source table. The default behavior is to exclude partition definition of the storage parameter of the source table.
  - **INCLUDING ALL** contains the meaning of **INCLUDING DEFAULTS**, **INCLUDING CONSTRAINTS**, **INCLUDING INDEXES**, **INCLUDING STORAGE**, **INCLUDING COMMENTS**, **INCLUDING PARTITION**, and **INCLUDING REOPTIONS**.
- **WITH ( storage\_parameter [= value] [, ... ] )**  
Specifies an optional storage parameter for a table or an index. Optional parameters are as follows:
    - **FILLFACTOR**  
The fill factor of a table is a percentage from 10 to 100. **100** (complete filling) is the default value. When a smaller fill factor is specified, INSERT operations fill table pages only to the indicated percentage. The remaining space on each page is reserved for updating rows on that page. This gives UPDATE a chance to place the updated copy of a row on the same page, which is more efficient than placing it on a different page. For a table whose entries are never updated, setting the fill factor to **100** (complete filling) is the best choice, but in heavily updated tables a smaller fill factor would be appropriate.  
Value range: 10–100.
    - **ORIENTATION**  
Determines the data storage mode of the table.  
Value range:

- **COLUMN**: The data will be stored in columns.
- **ROW** (default value): The data will be stored in rows.

---

**NOTICE**

**ORIENTATION** cannot be modified.

---

– **STORAGE\_TYPE**

Specifies the storage engine type. This parameter cannot be modified once it is set.

Value range:

- **USTORE** indicates that tables support the inplace-update storage engine. Note that the **track\_counts** and **track\_activities** parameters must be enabled when the Ustore table is used. Otherwise, space bloating may occur.
- **ASTORE** indicates that tables support the append-only storage engine.

Default value:

If no table is specified, data is stored in append-only mode by default.

– **segment**

The data is stored in segment-page mode. This parameter supports only row-store tables. Temporary tables and unlogged tables are not supported. Ustore is not supported.

Value range: **on** and **off**

Default value: **off**

• **TABLESPACE tablespace\_name**

Specifies that the new table will be created in the **tablespace\_name** tablespace. If it is not specified, the default tablespace is used.

• **PARTITION BY RANGE(partition\_key)**

Creates a range partition. **partition\_key** is the name of the partition key.

(1) Assume that the VALUES LESS THAN syntax is used.

---

**NOTICE**

In this case, a maximum of four partition keys are supported.

---

Data types supported by the partition keys are as follows: SMALLINT, INTEGER, BIGINT, DECIMAL, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER VARYING(*n*), VARCHAR(*n*), CHARACTER(*n*), CHAR(*n*), CHARACTER, CHAR, TEXT, NVARCHAR, NVARCHAR2, NAME, TIMESTAMP[(*p*)] [WITHOUT TIME ZONE], TIMESTAMP[(*p*)] [WITH TIME ZONE], and DATE.

(2) Assume that the START END syntax is used.

#### NOTICE

If the START END clause is used, a range partitioning policy supports only a one-column partition key.

Data types supported by the partition key are as follows: SMALLINT, INTEGER, BIGINT, DECIMAL, NUMERIC, REAL, DOUBLE PRECISION, TIMESTAMP[(p)] [WITHOUT TIME ZONE], TIMESTAMP[(p)] [WITH TIME ZONE], and DATE.

(3) Assume that the INTERVAL syntax is used.

#### NOTICE

In this case, only one partition key is supported.

In this case, the data types supported by the partition key are TIMESTAMP[(p)] [WITHOUT TIME ZONE], TIMESTAMP[(p)] [WITH TIME ZONE] and DATE.

- **PARTITION partition\_name VALUES LESS THAN ( { partition\_value | MAXVALUE } )**

Specifies the information of partitions. **partition\_name** is the name of a range partition. **partition\_value** is the upper limit of a range partition, and the value depends on the type of **partition\_key**. **MAXVALUE** usually specifies the upper limit of the last range partition.

#### NOTICE

- Each partition requires an upper limit.
- The data type of the upper limit must be the same as that of the partition key.
- In a partition list, partitions are arranged in ascending order of upper limits. A partition with a smaller upper limit value is placed before another partition with a larger one.

- **PARTITION partition\_name {START (partition\_value) END (partition\_value) EVERY (interval\_value)} | {START (partition\_value) END (partition\_value|MAXVALUE)} | {START(partition\_value)} | {END (partition\_value | MAXVALUE)}**

Specifies the information of partitions.

- **partition\_name**: name or name prefix of a range partition. It is the name prefix only in the following cases (assuming that **partition\_name** is **p1**):
  - If **START+END+EVERY** is used, the names of partitions will be defined as **p1\_1**, **p1\_2**, and the like. For example, if **PARTITION p1 START(1) END(4) EVERY(1)** is defined, the generated partitions are [1, 2), [2, 3), and [3, 4), and their names are **p1\_1**, **p1\_2**, and **p1\_3**. In this case, **p1** is a name prefix.
  - If the defined statement is in the first place and has **START** specified, the range (**MINVALUE**, **START**) will be automatically used as the first



actual partition, and its name will be **p1\_0**. The other partitions are then named **p1\_1**, **p1\_2**, and the like. For example, if **PARTITION p1 START(1), PARTITION p2 START(2)** is defined, generated partitions are (*MINVALUE*, 1), [1, 2), and [2, *MAXVALUE*), and their names will be **p1\_0**, **p1\_1**, and **p2**. In this case, **p1** is a name prefix and **p2** is a partition name. **MINVALUE** means the minimum value.

- **partition\_value**: start value or end value of a range partition. The value depends on **partition\_key** and cannot be *MAXVALUE*.
- **interval\_value**: width of each partition for dividing the [**START**, **END**) range. It cannot be *MAXVALUE*. If the value of (**END** - **START**) divided by **EVERY** has a remainder, the width of only the last partition is less than the value of **EVERY**.
- *MAXVALUE* usually specifies the upper limit of the last range partition.

---

**NOTICE**

1. If the defined statement is in the first place and has **START** specified, the range (*MINVALUE*, **START**) will be automatically used as the first actual partition.
  2. The START END syntax must comply with the following rules:
    - The value of **START** (if any, same for the following situations) in each **partition\_start\_end\_item** must be smaller than that of **END**.
    - In two adjacent **partition\_start\_end\_item** statements, the value of the first **END** must be equal to that of the second **START**.
    - The value of **EVERY** in each **partition\_start\_end\_item** must be a positive number (in ascending order) and must be smaller than **END** minus **START**.
    - Each partition includes the start value (unless it is *MINVALUE*) and excludes the end value. The format is as follows: [**START**, **END**).
    - Partitions created by the same **partition\_start\_end\_item** belong to the same tablespace.
    - If **partition\_name** is a name prefix of a partition, the length must not exceed 57 bytes. If there are more than 57 bytes, the prefix will be automatically truncated.
    - When creating or modifying a partitioned table, ensure that the total number of partitions in the table does not exceed the maximum value **1048575**.
  3. In statements for creating partitioned tables, START END and LESS THAN cannot be used together.
  4. The START END syntax in a partitioned table creation SQL statement will be replaced by the VALUES LESS THAN syntax when *gs\_dump* is executed.
- 
- **INTERVAL ('interval\_expr') [ STORE IN (tablespace\_name [, ... ] ) ]**  
Defines interval partitioning.
    - **interval\_expr**: interval for automatically creating partitions, for example, 1 day or 1 month.
    - **STORE IN (tablespace\_name [, ... ] )**: Specifies the list of tablespaces for storing automatically created partitions. If this parameter is specified, the

automatically created partitions are cyclically selected from the tablespace list. Otherwise, the default tablespace of the partitioned table is used.

- **PARTITION BY LIST(partition\_key)**

Create a list partition. **partition\_key** is the name of the partition key.

- For **partition\_key**, the list partitioning policy supports only one column of partition keys.
- For the clause syntax VALUES (list\_values), if **list\_values** contains the key values of the corresponding partition, it is recommended that the number of key values of each partition be less than or equal to 64.

Partition keys support the following data types: INT1, INT2, INT4, INT8, NUMERIC, VARCHAR(*n*), CHAR, BPCHAR, NVARCHAR, NVARCHAR2, TIMESTAMP[(*p*)] [WITHOUT TIME ZONE], TIMESTAMP[(*p*)] [WITH TIME ZONE], and DATE. The number of partitions cannot exceed 1048575.

- **PARTITION BY HASH(partition\_key)**

Create a hash partition. **partition\_key** is the name of the partition key.

For **partition\_key**, the hash partitioning policy supports only one column of partition keys.

Partition keys support the following data types: INT1, INT2, INT4, INT8, NUMERIC, VARCHAR(*n*), CHAR, BPCHAR, TEXT, NVARCHAR, NVARCHAR2, TIMESTAMP[(*p*)] [WITHOUT TIME ZONE], TIMESTAMP[(*p*)] [WITH TIME ZONE], and DATE. The number of partitions cannot exceed 1048575.

- **{ ENABLE | DISABLE } ROW MOVEMENT**

Specifies whether to enable row movement.

If the tuple value is updated on the partition key during the UPDATE operation, the partition where the tuple is located is altered. Setting this parameter enables error messages to be reported or movement of the tuple between partitions.

Value range:

- **ENABLE** (default value): Row movement is enabled.
- **DISABLE**: Row movement is disabled.

If the row movement is enabled, an error may be reported when UPDATE and DELETE operations are performed concurrently. The causes are as follows:

The old data is marked as deleted under the UPDATE and DELETE operations. If the row movement is enabled, the cross-partition update occurs when the partition key is updated, the kernel marks the old data in the old partition as deleted and adds a data to the new partition. As a result, the new data cannot be found by querying the old data.

If data in the same row is concurrently operated, the cross-partition and non-cross-partition data results have different behaviors in the following three concurrency scenarios: UPDATE and UPDATE concurrency, DELETE and DELETE concurrency, as well as UPDATE and DELETE concurrency.

- a. For non-cross-partition data, no error is reported for the second operation after the first operation is performed.

If the first operation is UPDATE, the latest data can be found and operated after the second operation is performed.

If the first operation is DELETE, the second operation is terminated if the current data is deleted and the latest data cannot be found.

- b. For the cross-partition data result, an error is reported for the second operation after the first operation is performed.

If the first operation is UPDATE, the second operation cannot find the latest data because the new data is in the new partition. Therefore, the second operation fails and an error is reported.

If the first operation is DELETE, performing the second operation can find that the current data is deleted and the latest data cannot be found, but cannot determine whether the operation of deleting the old data is UPDATE or DELETE. If the operation is UPDATE, an error is reported. If the operation is DELETE, the second operation is terminated. To ensure the data correctness, an error is reported.

If the UPDATE and UPDATE concurrency, and UPDATE and DELETE concurrency are performed, the error can be solved only when the operations are performed serially. If the DELETE and DELETE concurrency are performed, the error can be solved by disabling the row movement.

- **NOT NULL**

Specifies that the column value cannot be **NULL**. **ENABLE** can be omitted.

- **NULL**

Specifies that the column value can be **NULL**, which is the default value.

This clause is only provided for compatibility with non-standard SQL databases. It is not recommended.

- **CHECK (condition) [ NO INHERIT ]**

Specifies an expression producing a Boolean result where the INSERT or UPDATE operation of new or updated rows can succeed only when the expression result is true or unknown; otherwise, an error is thrown and the database is not altered.

A CHECK constraint specified as a column constraint should reference only the column's value, while an expression in a table constraint can reference multiple columns.

A constraint marked with **NO INHERIT** will not propagate to child tables.

**ENABLE** can be omitted.

- **DEFAULT default\_expr**

Assigns a default data value to a column. The value can be any variable-free expressions. (Subqueries and cross-references to other columns in the current table are not allowed.) The data type of the default expression must match that of the column.

The default expression will be used in any INSERT operation that does not specify a value for the column. If there is no default value for a column, then the default value is **NULL**.

- **GENERATED ALWAYS AS ( generation\_expr ) STORED**

This clause creates a column as a generated column. The value of the generated column is calculated by **generation\_expr** when data is written (inserted or updated). **STORED** indicates that the value of the generated column is stored as an ordinary column.

 **NOTE**

- The generation expression cannot refer to data other than the current row in any way. The generation expression cannot reference other generation columns or system columns. The generation expression cannot return a result set. No subquery, aggregate function, or window function can be used. The function called by the generation expression can only be an immutable function.
  - Default values cannot be specified for generated columns.
  - The generated column cannot be used as a part of the partition key.
  - Do not specify the generated column and the CASCADE, SET NULL, and SET DEFAULT actions of the ON UPDATE constraint at the same time. Do not specify the generated column and the SET NULL, and SET DEFAULT actions of the ON DELETE constraint at the same time.
  - The method of modifying and deleting generated columns is the same as that of ordinary columns. Delete the ordinary column that the generated column depends on. The generated column is automatically deleted. The type of the column on which the generated column depends cannot be changed.
  - The generated column cannot be directly written. In the INSERT or UPDATE statement, values cannot be specified for generated columns, but the keyword DEFAULT can be specified.
  - The permission control for generated columns is the same as that for ordinary columns.
- **UNIQUE index\_parameters**  
**UNIQUE ( column\_name [, ... ] ) index\_parameters**  
Specifies that a group of one or more columns of a table can contain only unique values.  
For the purpose of a unique constraint, null is not considered equal.
  - **PRIMARY KEY index\_parameters**  
**PRIMARY KEY ( column\_name [, ... ] ) index\_parameters**  
Specifies that a column or columns of a table can contain only unique (non-duplicate) and non-NULL values.  
Only one primary key can be specified for a table.
  - **DEFERRABLE | NOT DEFERRABLE**  
Determines whether the constraint can be deferred. A constraint that is not deferrable will be checked immediately after every command. Checking of constraints that are deferrable can be postponed until the end of the transaction using the **SET CONSTRAINTS** command. **NOT DEFERRABLE** is the default value. Currently, only UNIQUE constraints, primary key constraints, and foreign key constraints accept this clause. All the other constraints are not deferrable.
  - **INITIALLY IMMEDIATE | INITIALLY DEFERRED**  
If a constraint is deferrable, this clause specifies the default time to check the constraint.
    - If the constraint is **INITIALLY IMMEDIATE** (default value), it is checked after each statement.
    - If the constraint is **INITIALLY DEFERRED**, it is checked only at the end of the transaction.The constraint check time can be altered by executing the **SET CONSTRAINTS** command.

- **USING INDEX TABLESPACE tablespace\_name**

Specifies a tablespace in which an index associated with a UNIQUE or PRIMARY KEY constraint will be created. If this clause is not used, such index will be created in **default\_tablespace**. If **default\_tablespace** is empty, the default tablespace of the database is used.

## Examples

- Example 1: Create a range-partitioned table **tpcds.web\_returns\_p1**. The table has eight partitions and their partition keys are of the integer type. The ranges of the partitions are:  $wr\_returned\_date\_sk < 2450815$ ,  $2450815 \leq wr\_returned\_date\_sk < 2451179$ ,  $2451179 \leq wr\_returned\_date\_sk < 2451544$ ,  $2451544 \leq wr\_returned\_date\_sk < 2451910$ ,  $2451910 \leq wr\_returned\_date\_sk < 2452275$ ,  $2452275 \leq wr\_returned\_date\_sk < 2452640$ ,  $2452640 \leq wr\_returned\_date\_sk < 2453005$ , and  $wr\_returned\_date\_sk \geq 2453005$ .

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.web_returns table.
openGauss=# CREATE TABLE tpcds.web_returns
(
 W_WAREHOUSE_SK INTEGER NOT NULL,
 W_WAREHOUSE_ID CHAR(16) NOT NULL,
 W_WAREHOUSE_NAME VARCHAR(20)
 W_WAREHOUSE_SQ_FT INTEGER
 W_STREET_NUMBER CHAR(10)
 W_STREET_NAME VARCHAR(60)
 W_STREET_TYPE CHAR(15)
 W_SUITE_NUMBER CHAR(10)
 W_CITY VARCHAR(60)
 W_COUNTY VARCHAR(30)
 W_STATE CHAR(2)
 W_ZIP CHAR(10)
 W_COUNTRY VARCHAR(20)
 W_GMT_OFFSET DECIMAL(5,2)
);

-- Create a range-partitioned table tpcds.web_returns_p1.
openGauss=# CREATE TABLE tpcds.web_returns_p1
(
 WR_RETURNED_DATE_SK INTEGER
 WR_RETURNED_TIME_SK INTEGER
 WR_ITEM_SK INTEGER NOT NULL,
 WR_REFUNDED_CUSTOMER_SK INTEGER
 WR_REFUNDED_CDEMO_SK INTEGER
 WR_REFUNDED_HDEMO_SK INTEGER
 WR_REFUNDED_ADDR_SK INTEGER
 WR_RETURNING_CUSTOMER_SK INTEGER
 WR_RETURNING_CDEMO_SK INTEGER
 WR_RETURNING_HDEMO_SK INTEGER
 WR_RETURNING_ADDR_SK INTEGER
 WR_WEB_PAGE_SK INTEGER
 WR_REASON_SK INTEGER
 WR_ORDER_NUMBER BIGINT NOT NULL,
 WR_RETURN_QUANTITY INTEGER
 WR_RETURN_AMT DECIMAL(7,2)
 WR_RETURN_TAX DECIMAL(7,2)
 WR_RETURN_AMT_INC_TAX DECIMAL(7,2)
 WR_FEE DECIMAL(7,2)
 WR_RETURN_SHIP_COST DECIMAL(7,2)
 WR_REFUNDED_CASH DECIMAL(7,2)
 WR_REVERSED_CHARGE DECIMAL(7,2)
 WR_ACCOUNT_CREDIT DECIMAL(7,2)
 WR_NET_LOSS DECIMAL(7,2)
)
PARTITION BY RANGE(WR_RETURNED_DATE_SK)
```

```
(
 PARTITION P1 VALUES LESS THAN(2450815),
 PARTITION P2 VALUES LESS THAN(2451179),
 PARTITION P3 VALUES LESS THAN(2451544),
 PARTITION P4 VALUES LESS THAN(2451910),
 PARTITION P5 VALUES LESS THAN(2452275),
 PARTITION P6 VALUES LESS THAN(2452640),
 PARTITION P7 VALUES LESS THAN(2453005),
 PARTITION P8 VALUES LESS THAN(MAXVALUE)
);

-- Import data from the example data table.
openGauss=# INSERT INTO tpcds.web_returns_p1 SELECT * FROM tpcds.web_returns;

-- Drop the P8 partition.
openGauss=# ALTER TABLE tpcds.web_returns_p1 DROP PARTITION P8;

-- Add a partition WR_RETURNED_DATE_SK with values ranging from 2453005 to 2453105.
openGauss=# ALTER TABLE tpcds.web_returns_p1 ADD PARTITION P8 VALUES LESS THAN (2453105);

-- Add a partition WR_RETURNED_DATE_SK with values ranging from 2453105 to MAXVALUE.
openGauss=# ALTER TABLE tpcds.web_returns_p1 ADD PARTITION P9 VALUES LESS THAN
(MAXVALUE);

-- Drop the P8 partition.
openGauss=# ALTER TABLE tpcds.web_returns_p1 DROP PARTITION FOR (2453005);

-- Rename the P7 partition to P10.
openGauss=# ALTER TABLE tpcds.web_returns_p1 RENAME PARTITION P7 TO P10;

-- Rename the P6 partition to P11.
openGauss=# ALTER TABLE tpcds.web_returns_p1 RENAME PARTITION FOR (2452639) TO P11;

-- Query the number of rows in the P10 partition.
openGauss=# SELECT count(*) FROM tpcds.web_returns_p1 PARTITION (P10);
count

0
(1 row)

-- Query the number of rows in the P1 partition.
openGauss=# SELECT COUNT(*) FROM tpcds.web_returns_p1 PARTITION FOR (2450815);
count

0
(1 row)

-- Drop the tpcds.web_returns_p1 table.
openGauss=# DROP TABLE tpcds.web_returns_p1;

-- Drop the tpcds.web_returns table.
openGauss=# DROP TABLE tpcds.web_returns;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

- Example 2: Create a range-partitioned table **tpcds.web\_returns\_p2**. The table has eight partitions and their partition keys are of the integer type. The upper limit of the eighth partition is *MAXVALUE*.

The ranges of the partitions are:  $wr\_returned\_date\_sk < 2450815$ ,  $2450815 \leq wr\_returned\_date\_sk < 2451179$ ,  $2451179 \leq wr\_returned\_date\_sk < 2451544$ ,  $2451544 \leq wr\_returned\_date\_sk < 2451910$ ,  $2451910 \leq wr\_returned\_date\_sk < 2452275$ ,  $2452275 \leq wr\_returned\_date\_sk < 2452640$ ,  $2452640 \leq wr\_returned\_date\_sk < 2453005$ , and  $wr\_returned\_date\_sk \geq 2453005$ .

The tablespace of the **tpcds.web\_returns\_p2** partitioned table is **example1**. Partitions **P1** to **P7** have no specified tablespaces, and use the **example1**

tablespace of the **tpcds.web\_returns\_p2** partitioned table. The tablespace of the **P8** partitioned table is **example2**.

Assume that the following data directories of the database nodes are empty directories for which user **dwsadmin** has the read and write permissions: **/pg\_location/mount1/path1**, **/pg\_location/mount2/path2**, **/pg\_location/mount3/path3**, and **/pg\_location/mount4/path4**.

```
openGauss=# CREATE TABLESPACE example1 RELATIVE LOCATION 'tablespace1/tablespace_1';
openGauss=# CREATE TABLESPACE example2 RELATIVE LOCATION 'tablespace2/tablespace_2';
openGauss=# CREATE TABLESPACE example3 RELATIVE LOCATION 'tablespace3/tablespace_3';
openGauss=# CREATE TABLESPACE example4 RELATIVE LOCATION 'tablespace4/tablespace_4';
```

-- Create a schema.

```
openGauss=# CREATE SCHEMA tpcds;
```

```
openGauss=# CREATE TABLE tpcds.web_returns_p2
```

```
(
 WR_RETURNED_DATE_SK INTEGER ,
 WR_RETURNED_TIME_SK INTEGER ,
 WR_ITEM_SK INTEGER NOT NULL,
 WR_REFUNDED_CUSTOMER_SK INTEGER ,
 WR_REFUNDED_CDEMO_SK INTEGER ,
 WR_REFUNDED_HDEMO_SK INTEGER ,
 WR_REFUNDED_ADDR_SK INTEGER ,
 WR_RETURNING_CUSTOMER_SK INTEGER ,
 WR_RETURNING_CDEMO_SK INTEGER ,
 WR_RETURNING_HDEMO_SK INTEGER ,
 WR_RETURNING_ADDR_SK INTEGER ,
 WR_WEB_PAGE_SK INTEGER ,
 WR_REASON_SK INTEGER ,
 WR_ORDER_NUMBER BIGINT NOT NULL,
 WR_RETURN_QUANTITY INTEGER ,
 WR_RETURN_AMT DECIMAL(7,2) ,
 WR_RETURN_TAX DECIMAL(7,2) ,
 WR_RETURN_AMT_INC_TAX DECIMAL(7,2) ,
 WR_FEE DECIMAL(7,2) ,
 WR_RETURN_SHIP_COST DECIMAL(7,2) ,
 WR_REFUNDED_CASH DECIMAL(7,2) ,
 WR_REVERSED_CHARGE DECIMAL(7,2) ,
 WR_ACCOUNT_CREDIT DECIMAL(7,2) ,
 WR_NET_LOSS DECIMAL(7,2)
)
```

```
TABLESPACE example1
```

```
PARTITION BY RANGE(WR_RETURNED_DATE_SK)
```

```
(
 PARTITION P1 VALUES LESS THAN(2450815),
 PARTITION P2 VALUES LESS THAN(2451179),
 PARTITION P3 VALUES LESS THAN(2451544),
 PARTITION P4 VALUES LESS THAN(2451910),
 PARTITION P5 VALUES LESS THAN(2452275),
 PARTITION P6 VALUES LESS THAN(2452640),
 PARTITION P7 VALUES LESS THAN(2453005),
 PARTITION P8 VALUES LESS THAN(MAXVALUE) TABLESPACE example2
)
```

```
ENABLE ROW MOVEMENT;
```

-- Create a partitioned table using LIKE.

```
openGauss=# CREATE TABLE tpcds.web_returns_p3 (LIKE tpcds.web_returns_p2 INCLUDING
PARTITION);
```

-- Change the tablespace of the **P1** partition to **example2**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 MOVE PARTITION P1 TABLESPACE example2;
```

-- Change the tablespace of the **P2** partition to **example3**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 MOVE PARTITION P2 TABLESPACE example3;
```

-- Split the **P8** partition at **2453010**.

```
openGauss=# ALTER TABLE tpcds.web_returns_p2 SPLIT PARTITION P8 AT (2453010) INTO
(
```

```
 PARTITION P9,
 PARTITION P10
);

-- Merge the P6 and P7 partitions into one.
openGauss=# ALTER TABLE tpcds.web_returns_p2 MERGE PARTITIONS P6, P7 INTO PARTITION P8;

-- Modify the migration attribute of the partitioned table.
openGauss=# ALTER TABLE tpcds.web_returns_p2 DISABLE ROW MOVEMENT;
-- Drop tables and tablespaces.
openGauss=# DROP TABLE tpcds.web_returns_p1;
openGauss=# DROP TABLE tpcds.web_returns_p2;
openGauss=# DROP TABLE tpcds.web_returns_p3;
openGauss=# DROP TABLESPACE example1;
openGauss=# DROP TABLESPACE example2;
openGauss=# DROP TABLESPACE example3;
openGauss=# DROP TABLESPACE example4;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

- Example 3: Use START END to create and modify a range-partitioned table.

Assume that **/home/omm/startend\_tbs1**, **/home/omm/startend\_tbs2**, **/home/omm/startend\_tbs3**, and **/home/omm/startend\_tbs4** are empty directories for which user **omm** has the read and write permissions.

```
-- Create tablespaces.
openGauss=# CREATE TABLESPACE startend_tbs1 LOCATION '/home/omm/startend_tbs1';
openGauss=# CREATE TABLESPACE startend_tbs2 LOCATION '/home/omm/startend_tbs2';
openGauss=# CREATE TABLESPACE startend_tbs3 LOCATION '/home/omm/startend_tbs3';
openGauss=# CREATE TABLESPACE startend_tbs4 LOCATION '/home/omm/startend_tbs4';

-- Create a temporary schema.
openGauss=# CREATE SCHEMA tpcds;
openGauss=# SET CURRENT_SCHEMA TO tpcds;

-- Create a partitioned table with the partition key of the integer type.
openGauss=# CREATE TABLE tpcds.startend_pt (c1 INT, c2 INT)
TABLESPACE startend_tbs1
PARTITION BY RANGE (c2) (
 PARTITION p1 START(1) END(1000) EVERY(200) TABLESPACE startend_tbs2,
 PARTITION p2 END(2000),
 PARTITION p3 START(2000) END(2500) TABLESPACE startend_tbs3,
 PARTITION p4 START(2500),
 PARTITION p5 START(3000) END(5000) EVERY(1000) TABLESPACE startend_tbs4
)
ENABLE ROW MOVEMENT;

-- View the information of the partitioned table.
openGauss=# SELECT relname, boundaries, spcname FROM pg_partition p JOIN pg_tablespace t ON
p.reltablespace=t.oid and p.parentid='tpcds.startend_pt'::regclass ORDER BY 1;
 relname | boundaries | spcname
-----+-----+-----
p1_0 | {1} | startend_tbs2
p1_1 | {201} | startend_tbs2
p1_2 | {401} | startend_tbs2
p1_3 | {601} | startend_tbs2
p1_4 | {801} | startend_tbs2
p1_5 | {1000} | startend_tbs2
p2 | {2000} | startend_tbs1
p3 | {2500} | startend_tbs3
p4 | {3000} | startend_tbs1
p5_1 | {4000} | startend_tbs4
p5_2 | {5000} | startend_tbs4
startend_pt | | startend_tbs1
(12 rows)

-- Import data and check the data volume in a partition.
openGauss=# INSERT INTO tpcds.startend_pt VALUES (GENERATE_SERIES(0, 4999),
```



```
GENERATE_SERIES(0, 4999));
openGauss=# SELECT COUNT(*) FROM tpcds.startend_pt PARTITION FOR (0);
count

1
(1 row)

openGauss=# SELECT COUNT(*) FROM tpcds.startend_pt PARTITION (p3);
count

500
(1 row)

-- Add partitions [5000, 5300), [5300, 5600), [5600, 5900), and [5900, 6000).
openGauss=# ALTER TABLE tpcds.startend_pt ADD PARTITION p6 START(5000) END(6000)
EVERY(300) TABLESPACE startend_tbs4;

-- Add the partition p7, specified by MAXVALUE.
openGauss=# ALTER TABLE tpcds.startend_pt ADD PARTITION p7 END(MAXVALUE);

-- Rename the partition p7 to p8.
openGauss=# ALTER TABLE tpcds.startend_pt RENAME PARTITION p7 TO p8;

-- Drop the partition p8.
openGauss=# ALTER TABLE tpcds.startend_pt DROP PARTITION p8;

-- Rename the partition where 5950 is located to p71.
openGauss=# ALTER TABLE tpcds.startend_pt RENAME PARTITION FOR(5950) TO p71;

-- Split the partition [4000, 5000) where 4500 is located.
openGauss=# ALTER TABLE tpcds.startend_pt SPLIT PARTITION FOR(4500) INTO(PARTITION q1
START(4000) END(5000) EVERY(250) TABLESPACE startend_tbs3);

-- Change the tablespace of the partition p2 to startend_tbs4.
openGauss=# ALTER TABLE tpcds.startend_pt MOVE PARTITION p2 TABLESPACE startend_tbs4;

-- View the partition status.
openGauss=# SELECT relname, boundaries, spcname FROM pg_partition p JOIN pg_tablespace t ON
p.reltablespace=t.oid and p.parentid='tpcds.startend_pt'::regclass ORDER BY 1;
 relname | boundaries | spcname
-----+-----+-----
p1_0 | {1} | startend_tbs2
p1_1 | {201} | startend_tbs2
p1_2 | {401} | startend_tbs2
p1_3 | {601} | startend_tbs2
p1_4 | {801} | startend_tbs2
p1_5 | {1000} | startend_tbs2
p2 | {2000} | startend_tbs4
p3 | {2500} | startend_tbs3
p4 | {3000} | startend_tbs1
p5_1 | {4000} | startend_tbs4
p6_1 | {5300} | startend_tbs4
p6_2 | {5600} | startend_tbs4
p6_3 | {5900} | startend_tbs4
p71 | {6000} | startend_tbs4
q1_1 | {4250} | startend_tbs3
q1_2 | {4500} | startend_tbs3
q1_3 | {4750} | startend_tbs3
q1_4 | {5000} | startend_tbs3
startend_pt | | startend_tbs1
(19 rows)

-- Drop tables and tablespaces.
openGauss=# DROP SCHEMA tpcds CASCADE;
openGauss=# DROP TABLESPACE startend_tbs1;
openGauss=# DROP TABLESPACE startend_tbs2;
openGauss=# DROP TABLESPACE startend_tbs3;
openGauss=# DROP TABLESPACE startend_tbs4;
```

- Example 4: Create interval partitioned table **sales**. The table initially contains two partitions and the partition key is of the DATE type. Ranges of the two partitions are as follows: **time\_id** < '2019-02-01 00:00:00' and '2019-02-01 00:00:00' ≤ **time\_id** < '2019-02-02 00:00:00', respectively.

```
-- Create table sales.
openGauss=# CREATE TABLE sales
(prod_id NUMBER(6),
 cust_id NUMBER,
 time_id DATE,
 channel_id CHAR(1),
 promo_id NUMBER(6),
 quantity_sold NUMBER(3),
 amount_sold NUMBER(10,2)
)
PARTITION BY RANGE (time_id)
INTERVAL('1 day')
(PARTITION p1 VALUES LESS THAN ('2019-02-01 00:00:00'),
 PARTITION p2 VALUES LESS THAN ('2019-02-02 00:00:00')
);

-- Insert data into partition p1.
openGauss=# INSERT INTO sales VALUES(1, 12, '2019-01-10 00:00:00', 'a', 1, 1, 1);

-- Insert data into partition p2.
openGauss=# INSERT INTO sales VALUES(1, 12, '2019-02-01 00:00:00', 'a', 1, 1, 1);

-- View the partition information.
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'sales' AND t1.parttype = 'p';
 relname | partstrategy | boundaries
-----+-----+-----
p1 | r | {"2019-02-01 00:00:00"}
p2 | r | {"2019-02-02 00:00:00"}
(2 rows)
```

-- If the data to be inserted does not match any partition, create a partition, and insert the data into the new partition.

-- The range of the new partition is '2019-02-05 00:00:00' ≤ time\_id < '2019-02-06 00:00:00'.

```
openGauss=# INSERT INTO sales VALUES(1, 12, '2019-02-05 00:00:00', 'a', 1, 1, 1);
```

-- If the data to be inserted does not match any partition, create a partition, and insert the data into the new partition.

-- The range of the new partition is '2019-02-03 00:00:00' ≤ time\_id < '2019-02-04 00:00:00'.

```
openGauss=# INSERT INTO sales VALUES(1, 12, '2019-02-03 00:00:00', 'a', 1, 1, 1);
```

-- View the partition information.

```
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'sales' AND t1.parttype = 'p';
 relname | partstrategy | boundaries
-----+-----+-----
sys_p1 | i | {"2019-02-06 00:00:00"}
sys_p2 | i | {"2019-02-04 00:00:00"}
p1 | r | {"2019-02-01 00:00:00"}
p2 | r | {"2019-02-02 00:00:00"}
(4 rows)
```

- Example 5: Create list partitioned table **test\_list**. The table initially contains four partitions and the partition key is of the INT type. The ranges of the four partitions are 2000, 3000, 4000, and 5000 respectively.

```
-- Create the test_list table.
openGauss=# create table test_list (col1 int, col2 int)
partition by list(col1)
(
partition p1 values (2000),
partition p2 values (3000),
partition p3 values (4000),
partition p4 values (5000)
);
```

```
-- Insert data.
openGauss=# INSERT INTO test_list VALUES(2000, 2000);
INSERT 0 1
openGauss=# INSERT INTO test_list VALUES(3000, 3000);
INSERT 0 1

-- View the partition information.
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'test_list' AND t1.parttype = 'p';
 relname | partstrategy | boundaries
-----+-----+-----
p1 | l | {2000}
p2 | l | {3000}
p3 | l | {4000}
p4 | l | {5000}
(4 rows)

-- The inserted data does not match the partition, and an error is reported.
openGauss=# INSERT INTO test_list VALUES(6000, 6000);
ERROR: inserted partition key does not map to any table partition

-- Add a partition.
openGauss=# alter table test_list add partition p5 values (6000);
ALTER TABLE
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'test_list' AND t1.parttype = 'p';
 relname | partstrategy | boundaries
-----+-----+-----
p5 | l | {6000}
p4 | l | {5000}
p1 | l | {2000}
p2 | l | {3000}
p3 | l | {4000}
(5 rows)
openGauss=# INSERT INTO test_list VALUES(6000, 6000);
INSERT 0 1

-- Exchange data between the partitioned table and ordinary table.
openGauss=# create table t1 (col1 int, col2 int);
CREATE TABLE
openGauss=# select * from test_list partition (p1);
 col1 | col2
-----+-----
2000 | 2000
(1 row)
openGauss=# alter table test_list exchange partition (p1) with table t1;
ALTER TABLE
openGauss=# select * from test_list partition (p1);
 col1 | col2
-----+-----
(0 rows)
openGauss=# select * from t1;
 col1 | col2
-----+-----
2000 | 2000
(1 row)

-- Truncate the partition.
openGauss=# select * from test_list partition (p2);
 col1 | col2
-----+-----
3000 | 3000
(1 row)
openGauss=# alter table test_list truncate partition p2;
ALTER TABLE
openGauss=# select * from test_list partition (p2);
 col1 | col2
-----+-----
```

(0 rows)

-- Drop the partition.

```
openGauss=# alter table test_list drop partition p5;
```

```
ALTER TABLE
```

```
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'test_list' AND t1.parttype = 'p';
relname | partstrategy | boundaries
```

```
-----+-----
p4 | l | {5000}
p1 | l | {2000}
p2 | l | {3000}
p3 | l | {4000}
```

(4 rows)

```
openGauss=# INSERT INTO test_list VALUES(6000, 6000);
```

```
ERROR: inserted partition key does not map to any table partition
```

-- Drop the partitioned table.

```
openGauss=# drop table test_list;
```

- Example 6: Create a hash partitioned table **test\_hash**. The table initially contains two partitions and the partition key is of the INT type.

-- Create the **test\_hash** table.

```
openGauss=# create table test_hash (col1 int, col2 int)
```

```
partition by hash(col1)
```

```
(
```

```
partition p1,
```

```
partition p2
```

```
);
```

-- Insert data.

```
openGauss=# INSERT INTO test_hash VALUES(1, 1);
```

```
INSERT 0 1
```

```
openGauss=# INSERT INTO test_hash VALUES(2, 2);
```

```
INSERT 0 1
```

```
openGauss=# INSERT INTO test_hash VALUES(3, 3);
```

```
INSERT 0 1
```

```
openGauss=# INSERT INTO test_hash VALUES(4, 4);
```

```
INSERT 0 1
```

-- View the partition information.

```
openGauss=# SELECT t1.relname, partstrategy, boundaries FROM pg_partition t1, pg_class t2 WHERE
t1.parentid = t2.oid AND t2.relname = 'test_hash' AND t1.parttype = 'p';
relname | partstrategy | boundaries
```

```
-----+-----
p1 | h | {0}
p2 | h | {1}
```

(2 rows)

-- View the data.

```
openGauss=# select * from test_hash partition (p1);
```

```
col1 | col2
```

```
-----+-----
3 | 3
4 | 4
```

(2 rows)

```
openGauss=# select * from test_hash partition (p2);
```

```
col1 | col2
```

```
-----+-----
1 | 1
2 | 2
```

(2 rows)

-- Exchange data between the partitioned table and ordinary table.

```
openGauss=# create table t1 (col1 int, col2 int);
```

```
CREATE TABLE
```

```
openGauss=# alter table test_hash exchange partition (p1) with table t1;
```

```
ALTER TABLE
```

```
openGauss=# select * from test_hash partition (p1);
col1 | col2
-----+-----
(0 rows)
openGauss=# select * from t1;
col1 | col2
-----+-----
 3 | 3
 4 | 4
(2 rows)

-- Truncate the partition.
openGauss=# alter table test_hash truncate partition p2;
ALTER TABLE
openGauss=# select * from test_hash partition (p2);
col1 | col2
-----+-----
(0 rows)

-- Drop the partitioned table.
openGauss=# drop table test_hash;
```

## Helpful Links

[ALTER TABLE PARTITION](#) and [DROP TABLE](#)

## 7.13.74 CREATE TABLESPACE

### Description

Creates a tablespace in a database.

### Precautions

- A system administrator or a user who inherits permissions of the built-in role `gs_role_tablespace` can create a tablespace.
- Do not run **CREATE TABLESPACE** in a transaction block.
- If executing **CREATE TABLESPACE** fails but the internal directory (or file) has been created, the directory (or file) will remain. You need to manually clear it before creating the tablespace again. If there are residual files of soft links for the tablespace in the data directory, delete the residual files, and then perform O&M operations.
- **CREATE TABLESPACE** cannot be used for two-phase transactions. If it fails on some nodes, the execution cannot be rolled back.
- For details about how to prepare for creating tablespaces, see the description of parameters below.
- You are advised not to use user-defined tablespaces. This is because user-defined tablespaces are usually used with storage media other than the main storage (storage device where the default tablespace is located, such as a disk) to isolate I/O resources that can be used by different services. Storage devices use standard configurations and do not have other available storage media. If the user-defined tablespace is not properly used, the system cannot run stably for a long time and the overall performance is affected. Therefore, you are advised to use the default tablespace.

## Syntax

```
CREATE TABLESPACE tablespace_name
[OWNER user_name] [RELATIVE] LOCATION 'directory' [MAXSIZE 'space_size']
[with_option_clause];
```

The with\_option\_clause syntax for creating an ordinary tablespace is as follows:

```
WITH ({filesystem= { 'general'| "general" | general} |
 random_page_cost = { 'value ' | value } |
 seq_page_cost = { 'value ' | value }}[,...])
```

## Parameters

- **tablespace\_name**  
Specifies name of a tablespace to be created.  
The tablespace name must be distinct from the name of any existing tablespace in the database and cannot start with "pg", which are reserved for system catalog spaces.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **OWNER user\_name**  
Specifies the name of the user who will own the tablespace. If omitted, the default owner is the current user.  
Only system administrators can create tablespaces, but they can use the OWNER clause to assign ownership of tablespaces to non-system-administrators.  
Value range: a string. It must be the name an existing user.
- **RELATIVE**  
If this parameter is specified, a relative path is used. The location directory is relative to each CN/DN data directory.  
Directory hierarchy: *data directory/pg\_location/relative path* A relative path contains a maximum of two levels.  
If this parameter is not specified, the absolute tablespace path is used. The location directory must be an absolute path.
- **LOCATION directory**  
Specifies a directory for the table space. When creating an absolute tablespace path, ensure that the directory meets the following requirements:
  - The GaussDB system user must have the read and write permissions on the directory, and the directory must be empty. If the directory does not exist, the system automatically creates it.
  - The directory must be an absolute path, and does not contain special characters, such as dollar sign (\$) and greater-than sign (>).
  - The directory cannot be specified under the database data directory.
  - The directory must be a local path.Value range: a string. It must be a valid directory.
- **MAXSIZE 'space\_size'**  
Specifies the maximum value of the tablespace in a single database node.  
Value range: a string consisting of a positive integer and unit. The unit can be KB, MB, GB, TB, or PB currently. The unit of parsed value is KB and cannot

exceed the range that can be expressed in 64 bits, which is 1 KB to 9007199254740991 KB.

- **random\_page\_cost**

Specifies the cost of randomly reading the page overhead.

Value range: 0 to 1.79769e+308.

Default value: value of the GUC parameter **random\_page\_cost**

- **seq\_page\_cost**

Specifies the cost of reading the page overhead in specified order.

Value range: 0 to 1.79769e+308.

Default value: value of GUC parameter **seq\_page\_cost**

## Examples

```
-- Create a tablespace.
openGauss=# CREATE TABLESPACE ds_location1 RELATIVE LOCATION 'tablespace/tablespace_1';

-- Create user joe.
openGauss=# CREATE ROLE joe IDENTIFIED BY '*****!';

-- Create user jay.
openGauss=# CREATE ROLE jay IDENTIFIED BY '*****!';

-- Create a tablespace and set its owner to user joe.
openGauss=# CREATE TABLESPACE ds_location2 OWNER joe RELATIVE LOCATION 'tablespace/
tablespace_1';

-- Rename the ds_location1 tablespace to ds_location3.
openGauss=# ALTER TABLESPACE ds_location1 RENAME TO ds_location3;

-- Change the owner of the ds_location2 tablespace.
openGauss=# ALTER TABLESPACE ds_location2 OWNER TO jay;

-- Drop tablespaces.
openGauss=# DROP TABLESPACE ds_location2;
openGauss=# DROP TABLESPACE ds_location3;

-- Drop users.
openGauss=# DROP ROLE joe;
openGauss=# DROP ROLE jay;
```

## Helpful Links

[CREATE DATABASE](#), [CREATE TABLE](#), [CREATE INDEX](#), [DROP TABLESPACE](#), and [ALTER TABLESPACE](#)

## Suggestions

- create tablespace

You are advised not to create tablespaces in a transaction.

## 7.13.75 CREATE TABLE SUBPARTITION

### Description

Creates a level-2 partitioned table. Partitioning refers to splitting what is logically one large table into smaller physical pieces based on specific schemes. The table based on the logic is called a partitioned table, and each physical piece is called a

partition. Data is stored on these physical partitions, instead of the logical partitioned table. For a level-2 partitioned table, the top-level node table and level-1 partitioned table are logical tables and do not store data. Only the level-2 partitioned (leaf node) stores data.

The partitioning solution of a level-2 partitioned table is a combination of the partitioning solutions of two level-1 partitions. For details about the partitioning solution of a level-1 partitioned table, see [CREATE TABLE PARTITION](#).

Common combination solutions for level-2 partitioned tables include range-range partitioning, range-list partitioning, range-hash partitioning, list-range partitioning, list-list partitioning, list-hash partitioning, hash-range partitioning, hash-list partitioning, and hash-hash partitioning. Currently, level-2 partitioned tables can only be row-store tables.

## Precautions

- A level-2 partitioned table has two partition keys, and each partition key supports only one column. The two partition keys cannot be the same column.
- If the constraint key of the unique constraint and primary key constraint contains all partition keys, a local index is created for the constraints. Otherwise, a global index is created. If a local unique index is created, all partition keys must be included.
- When a level-2 partitioned table is created, if the specified level-2 partition is not displayed under the level-1 partition, a level-2 partition with the same range is automatically created.
- The number of level-2 partitions (leaf nodes) in a level-2 partitioned table cannot exceed 1048575. There is no limit on the number of level-1 partitions, but there must be at least one level-2 partition under a level-1 partition.
- The maximum total number of partitions (including level-1 and level-2 partitions) in a level-2 partitioned table is 1048575. Generally, it is impossible for services to create so many partitions and the memory will be insufficient if so. Create partitions based on the value of **local\_syscache\_threshold**. The memory allocated to the level-2 partitioned tables can be calculated as follows: total number of partitions x 3/1024, in MB. Theoretically, the memory occupied by the partitions cannot be greater than the value of **local\_syscache\_threshold**. In addition, some space must be reserved for other functions.
- Level-2 partitioned tables support only row store and do not support hash bucket.
- Clusters are not supported.
- When specifying a partition for query, for example, running **SELECT \* FROM tablename PARTITION/SUBPARTITION (partitionname)**, ensure that **PARTITION** and **SUBPARTITION** are correct. If they are incorrect, no error is reported during the query. In this case, the query is performed based on the table alias.
- Row-level security is not supported.
- In the **PARTITION FOR (values)** syntax for level-2 partitioned tables, values can only be constants.
- In the **PARTITION/SUBPARTITION FOR (values)** syntax for level-2 partitioned tables, if data type conversion is required for values, you are advised to use



forcible type conversion to prevent the implicit type conversion result from being inconsistent with the expected result.

- Currently, the statement specifying a partition cannot perform global index scan.

## Syntax

```
CREATE TABLE [IF NOT EXISTS] subpartition_table_name
(
{ column_name data_type [COLLATE collation] [column_constraint [...]]
| table_constraint
| LIKE source_table [like_option [...]] }, ...]
)
[WITH ({storage_parameter = value} [, ...])]
[TABLESPACE tablespace_name]
PARTITION BY {RANGE | LIST | HASH} (partition_key) SUBPARTITION BY {RANGE | LIST | HASH}
(subpartition_key)
(
PARTITION partition_name1 [VALUES LESS THAN (val1) | VALUES (val1[, ...])] [TABLESPACE tablespace]
[(
{ SUBPARTITION subpartition_name1 [VALUES LESS THAN (val1_1) | VALUES (val1_1[, ...])]
| TABLESPACE tablespace } [, ...]
)][, ...]
)[{ ENABLE | DISABLE } ROW MOVEMENT];
```

- column\_constraint:**

```
[CONSTRAINT constraint_name]
{ NOT NULL |
NULL |
CHECK (expression) |
DEFAULT default_expr |
GENERATED ALWAYS AS (generation_expr) STORED |
UNIQUE index_parameters |
PRIMARY KEY index_parameters |
REFERENCES reftable [(refcolumn)] [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
[ON DELETE action] [ON UPDATE action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- table\_constraint:**

```
[CONSTRAINT constraint_name]
{ CHECK (expression) |
UNIQUE (column_name [, ...]) index_parameters |
PRIMARY KEY (column_name [, ...]) index_parameters |
FOREIGN KEY (column_name [, ...]) REFERENCES reftable [(refcolumn [, ...])]
[MATCH FULL | MATCH PARTIAL | MATCH SIMPLE] [ON DELETE action] [ON UPDATE
action] }
[DEFERRABLE | NOT DEFERRABLE | INITIALLY DEFERRED | INITIALLY IMMEDIATE]
```

- like\_option:**

```
{ INCLUDING | EXCLUDING } { DEFAULTS | GENERATED | CONSTRAINTS | INDEXES | STORAGE |
COMMENTS | REOPTIONS | ALL }
```

- index\_parameters:**

```
[WITH ({storage_parameter = value} [, ...])]
[USING INDEX TABLESPACE tablespace_name]
```

## Parameters

- IF NOT EXISTS**

Sends a notice instead of an error if tables have identical names. The notice prompts that the table relationship already exists.

- subpartition\_table\_name**

Specifies the name of a level-2 partitioned table.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **column\_name**  
Specifies the name of a column to be created in the new table.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **data\_type**  
Specifies the data type of the column.
- **COLLATE collation**  
Assigns a collation to the column (which must be of a collatable data type). If no collation is specified, the default collation is used. You can run the **SELECT \* FROM pg\_collation;** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.
- **CONSTRAINT constraint\_name**  
Specifies the name of a column or table constraint. The optional constraint clauses specify constraints that new or updated rows must satisfy for an INSERT or UPDATE operation to succeed.  
There are two ways to define constraints:
  - A column constraint is defined as part of a column definition, and it is bound to a particular column.
  - A table constraint is not bound to a particular column but can apply to more than one column.
- **LIKE source\_table [ like\_option ... ]**  
Level-2 partitioned tables do not support this function.
- **WITH ( storage\_parameter [= value] [, ... ] )**  
Specifies an optional storage parameter for a table or an index. Optional parameters are as follows:
  - **FILLFACTOR**  
The fill factor of a table is a percentage from 10 to 100. **100** (complete filling) is the default value. When a smaller fill factor is specified, INSERT operations fill table pages only to the indicated percentage. The remaining space on each page is reserved for updating rows on that page. This gives UPDATE a chance to place the updated copy of a row on the same page, which is more efficient than placing it on a different page. For a table whose entries are never updated, setting the fill factor to **100** (complete filling) is the best choice, but in heavily updated tables a smaller fill factor would be appropriate.  
Value range: 10-100.
  - **ORIENTATION**  
Determines the data storage mode of the table.  
Value range:
    - **COLUMN**: The data will be stored in columns.
    - **ROW** (default value): The data will be stored in rows.

---

**NOTICE**

**ORIENTATION** cannot be modified.

---

– **STORAGE\_TYPE**

Specifies the storage engine type. This parameter cannot be modified once it is set.

Value range:

- **USTORE** indicates that tables support the inplace-update storage engine. Note that the **track\_counts** and **track\_activities** parameters must be enabled when the Ustore table is used. Otherwise, space bloating may occur.
- **ASTORE** indicates that tables support the append-only storage engine.

---

**NOTICE**

To use Ustore tables, you must enable the **track\_counts** and **track\_activities** parameters. Otherwise, space expansion occurs.

---

Default value:

If no table is specified, data is stored in append-only mode by default.

– **segment**

The data is stored in segment-page mode. This parameter supports only row-store tables. Temporary tables and unlogged tables are not supported. Ustore is not supported.

Value range: **on** and **off**

Default value: **off**

• **TABLESPACE tablespace\_name**

Specifies that the new table will be created in the **tablespace\_name** tablespace. If it is not specified, the default tablespace is used.

• **PARTITION BY {RANGE | LIST | HASH} (partition\_key)**

- For **partition\_key**, the partitioning policy supports only one column of partition keys.
- The data types supported by the partition key are the same as those supported by the level-1 partitioned table.

• **SUBPARTITION BY {RANGE | LIST | HASH} (subpartition\_key)**

- For **subpartition\_key**, the partitioning policy supports only one column of partition keys.
- The data types supported by the partition key are the same as those supported by the level-1 partitioned table.

• **{ ENABLE | DISABLE } ROW MOVEMENT**

Specifies whether to enable row movement.

If the tuple value is updated on the partition key during the UPDATE operation, the partition where the tuple is located is altered. Setting this parameter enables error messages to be reported or movement of the tuple between partitions.

Value range:

- **ENABLE** (default value): Row movement is enabled.
- **DISABLE**: Row movement is disabled.

If the row movement is enabled, an error may be reported when UPDATE and DELETE operations are performed concurrently. The causes are as follows:

The old data is marked as deleted under the UPDATE and DELETE operations. If the row movement is enabled, the cross-partition update occurs when the partition key is updated, the kernel marks the old data in the old partition as deleted and adds a data to the new partition. As a result, the new data cannot be found by querying the old data.

If data in the same row is concurrently operated, the cross-partition and non-cross-partition data results have different behaviors in the following three concurrency scenarios: UPDATE and UPDATE concurrency, DELETE and DELETE concurrency, UPDATE and DELETE concurrency.

- a. For non-cross-partition data, no error is reported for the second operation after the first operation is performed.

If the first operation is UPDATE, the latest data can be found and operated after the second operation is performed.

If the first operation is DELETE, the second operation is terminated if the current data is deleted and the latest data cannot be found.

- b. For the cross-partition data result, an error is reported for the second operation after the first operation is performed.

If the first operation is UPDATE, the second operation cannot find the latest data because the new data is in the new partition. Therefore, the second operation fails and an error is reported.

If the first operation is DELETE, performing the second operation can find that the current data is deleted and the latest data cannot be found, but cannot determine whether the operation of deleting the old data is UPDATE or DELETE. If the operation is UPDATE, an error is reported. If the operation is DELETE, the second operation is terminated. To ensure the data correctness, an error is reported.

If the UPDATE and UPDATE concurrency, and UPDATE and DELETE concurrency are performed, the error can be solved only when the operations are performed serially. If the DELETE and DELETE concurrency are performed, the error can be solved by disabling the row movement.

- **NOT NULL**

Specifies that the column value cannot be **NULL**. **ENABLE** can be omitted.

- **NULL**

Specifies that the column value can be **NULL**, which is the default value.

This clause is only provided for compatibility with non-standard SQL databases. It is not recommended.

- **CHECK (condition) [ NO INHERIT ]**

Specifies an expression producing a Boolean result where the INSERT or UPDATE operation of new or updated rows can succeed only when the expression result is true or unknown; otherwise, an error is thrown and the database is not altered.

A CHECK constraint specified as a column constraint should reference only the column's value, while an expression in a table constraint can reference multiple columns.

A constraint marked with **NO INHERIT** will not propagate to child tables.

**ENABLE** can be omitted.

- **DEFAULT default\_expr**

Assigns a default data value to a column. The value can be any variable-free expressions. (Subqueries and cross-references to other columns in the current table are not allowed.) The data type of the default expression must match that of the column.

The default expression will be used in any INSERT operation that does not specify a value for the column. If there is no default value for a column, then the default value is **NULL**.

- **GENERATED ALWAYS AS ( generation\_expr ) STORED**

This clause creates a column as a generated column. The value of the generated column is calculated by **generation\_expr** when data is written (inserted or updated). **STORED** indicates that the value of the generated column is stored as an ordinary column.

 **NOTE**

- The generation expression cannot refer to data other than the current row in any way. The generation expression cannot reference other generation columns or system columns. The generation expression cannot return a result set. No subquery, aggregate function, or window function can be used. The function called by the generation expression can only be an immutable function.
- Default values cannot be specified for generated columns.
- The generated column cannot be used as a part of the partition key.
- Do not specify the generated column and the CASCADE, SET NULL, and SET DEFAULT actions of the ON UPDATE constraint at the same time. Do not specify the generated column and the SET NULL, and SET DEFAULT actions of the ON DELETE constraint at the same time.
- The method of modifying and deleting generated columns is the same as that of ordinary columns. Delete the ordinary column that the generated column depends on. The generated column is automatically deleted. The type of the column on which the generated column depends cannot be changed.
- The generated column cannot be directly written. In the INSERT or UPDATE statement, values cannot be specified for generated columns, but the keyword DEFAULT can be specified.
- The permission control for generated columns is the same as that for ordinary columns.
- **UNIQUE index\_parameters/**  
**UNIQUE ( column\_name [, ... ] ) index\_parameters**  
Specifies that a group of one or more columns of a table can contain only unique values.  
For the purpose of a unique constraint, null is not considered equal.
- **PRIMARY KEY index\_parameters**  
**PRIMARY KEY ( column\_name [, ... ] ) index\_parameters**  
Specifies that a column or columns of a table can contain only unique (non-duplicate) and non-**NULL** values.

Only one primary key can be specified for a table.

- **DEFERRABLE | NOT DEFERRABLE**

They determine whether the constraint is deferrable. A constraint that is not deferrable will be checked immediately after every command. Checking of constraints that are deferrable can be postponed until the end of the transaction using the **SET CONSTRAINTS** command. **NOT DEFERRABLE** is the default value. Currently, only UNIQUE constraints, primary key constraints, and foreign key constraints accept this clause. All the other constraints are not deferrable.

- **INITIALLY IMMEDIATE | INITIALLY DEFERRED**

If a constraint is deferrable, this clause specifies the default time to check the constraint.

- If the constraint is **INITIALLY IMMEDIATE** (default value), it is checked after each statement.
- If the constraint is **INITIALLY DEFERRED**, it is checked only at the end of the transaction.

The constraint check time can be altered by executing the **SET CONSTRAINTS** command.

- **USING INDEX TABLESPACE tablespace\_name**

Specifies a tablespace in which an index associated with a UNIQUE or PRIMARY KEY constraint will be created. If this clause is not used, such index will be created in **default\_tablespace**. If **default\_tablespace** is empty, the default tablespace of the database is used.

## Examples

- Example 1: Create level-2 partitioned tables of various combination types.

```
CREATE TABLE list_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY LIST (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES ('201902')
 (
 SUBPARTITION p_201901_a VALUES ('1'),
 SUBPARTITION p_201901_b VALUES ('2')
),
 PARTITION p_201902 VALUES ('201903')
 (
 SUBPARTITION p_201902_a VALUES ('1'),
 SUBPARTITION p_201902_b VALUES ('2')
)
);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201902', '2', '1', 1);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
insert into list_list values('201903', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
select * from list_list;
 month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
```

```
201903 | 1 | 1 | 1
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(6 rows)

drop table list_list;
CREATE TABLE list_hash
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY LIST (month_code) SUBPARTITION BY HASH (dept_code)
(
 PARTITION p_201901 VALUES ('201902')
 (
 SUBPARTITION p_201901_a,
 SUBPARTITION p_201901_b
),
 PARTITION p_201902 VALUES ('201903')
 (
 SUBPARTITION p_201902_a,
 SUBPARTITION p_201902_b
)
);
insert into list_hash values('201902', '1', '1', 1);
insert into list_hash values('201902', '2', '1', 1);
insert into list_hash values('201902', '3', '1', 1);
insert into list_hash values('201903', '4', '1', 1);
insert into list_hash values('201903', '5', '1', 1);
insert into list_hash values('201903', '6', '1', 1);
select * from list_hash;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 4 | 1 | 1
201903 | 5 | 1 | 1
201903 | 6 | 1 | 1
201902 | 2 | 1 | 1
201902 | 3 | 1 | 1
201902 | 1 | 1 | 1
(6 rows)

drop table list_hash;
CREATE TABLE list_range
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY LIST (month_code) SUBPARTITION BY RANGE (dept_code)
(
 PARTITION p_201901 VALUES ('201902')
 (
 SUBPARTITION p_201901_a values less than ('4'),
 SUBPARTITION p_201901_b values less than ('6')
),
 PARTITION p_201902 VALUES ('201903')
 (
 SUBPARTITION p_201902_a values less than ('3'),
 SUBPARTITION p_201902_b values less than ('6')
)
);
insert into list_range values('201902', '1', '1', 1);
insert into list_range values('201902', '2', '1', 1);
insert into list_range values('201902', '3', '1', 1);
insert into list_range values('201903', '4', '1', 1);
```

```
insert into list_range values('201903', '5', '1', 1);
insert into list_range values('201903', '6', '1', 1);
ERROR: inserted partition key does not map to any table partition
select * from list_range;
 month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 4 | 1 | 1
201903 | 5 | 1 | 1
201902 | 1 | 1 | 1
201902 | 2 | 1 | 1
201902 | 3 | 1 | 1
(5 rows)

drop table list_range;
CREATE TABLE range_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY RANGE (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES LESS THAN('201903')
 (
 SUBPARTITION p_201901_a values ('1'),
 SUBPARTITION p_201901_b values ('2')
),
 PARTITION p_201902 VALUES LESS THAN('201904')
 (
 SUBPARTITION p_201902_a values ('1'),
 SUBPARTITION p_201902_b values ('2')
)
);
insert into range_list values('201902', '1', '1', 1);
insert into range_list values('201902', '2', '1', 1);
insert into range_list values('201902', '1', '1', 1);
insert into range_list values('201903', '2', '1', 1);
insert into range_list values('201903', '1', '1', 1);
insert into range_list values('201903', '2', '1', 1);
select * from range_list;
 month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
(6 rows)

drop table range_list;
CREATE TABLE range_hash
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY RANGE (month_code) SUBPARTITION BY HASH (dept_code)
(
 PARTITION p_201901 VALUES LESS THAN('201903')
 (
 SUBPARTITION p_201901_a,
 SUBPARTITION p_201901_b
),
 PARTITION p_201902 VALUES LESS THAN('201904')
 (
 SUBPARTITION p_201902_a,
```



```
SUBPARTITION p_201902_b
)
);
insert into range_hash values('201902', '1', '1', 1);
insert into range_hash values('201902', '2', '1', 1);
insert into range_hash values('201902', '1', '1', 1);
insert into range_hash values('201903', '2', '1', 1);
insert into range_hash values('201903', '1', '1', 1);
insert into range_hash values('201903', '2', '1', 1);
select * from range_hash;
 month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
(6 rows)

drop table range_hash;
CREATE TABLE range_range
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY RANGE (month_code) SUBPARTITION BY RANGE (dept_code)
(
 PARTITION p_201901 VALUES LESS THAN('201903')
 (
 SUBPARTITION p_201901_a VALUES LESS THAN('2'),
 SUBPARTITION p_201901_b VALUES LESS THAN('3')
),
 PARTITION p_201902 VALUES LESS THAN('201904')
 (
 SUBPARTITION p_201902_a VALUES LESS THAN('2'),
 SUBPARTITION p_201902_b VALUES LESS THAN('3')
)
);
insert into range_range values('201902', '1', '1', 1);
insert into range_range values('201902', '2', '1', 1);
insert into range_range values('201902', '1', '1', 1);
insert into range_range values('201903', '2', '1', 1);
insert into range_range values('201903', '1', '1', 1);
insert into range_range values('201903', '2', '1', 1);
select * from range_range;
 month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 2 | 1 | 1
201903 | 1 | 1 | 1
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
(6 rows)

drop table range_range;
CREATE TABLE hash_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY hash (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901
```

```
(
 SUBPARTITION p_201901_a VALUES ('1'),
 SUBPARTITION p_201901_b VALUES ('2')
),
PARTITION p_201902
(
 SUBPARTITION p_201902_a VALUES ('1'),
 SUBPARTITION p_201902_b VALUES ('2')
)
);
insert into hash_list values('201901', '1', '1', 1);
insert into hash_list values('201901', '2', '1', 1);
insert into hash_list values('201901', '1', '1', 1);
insert into hash_list values('201903', '2', '1', 1);
insert into hash_list values('201903', '1', '1', 1);
insert into hash_list values('201903', '2', '1', 1);
select * from hash_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
201901 | 2 | 1 | 1
201901 | 1 | 1 | 1
201901 | 1 | 1 | 1
(6 rows)

drop table hash_list;
CREATE TABLE hash_hash
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY hash (month_code) SUBPARTITION BY hash (dept_code)
(
 PARTITION p_201901
 (
 SUBPARTITION p_201901_a,
 SUBPARTITION p_201901_b
),
 PARTITION p_201902
 (
 SUBPARTITION p_201902_a,
 SUBPARTITION p_201902_b
)
);
insert into hash_hash values('201901', '1', '1', 1);
insert into hash_hash values('201901', '2', '1', 1);
insert into hash_hash values('201901', '1', '1', 1);
insert into hash_hash values('201903', '2', '1', 1);
insert into hash_hash values('201903', '1', '1', 1);
insert into hash_hash values('201903', '2', '1', 1);
select * from hash_hash;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
201901 | 2 | 1 | 1
201901 | 1 | 1 | 1
201901 | 1 | 1 | 1
(6 rows)

drop table hash_hash;
CREATE TABLE hash_range
(
 month_code VARCHAR2 (30) NOT NULL ,
```

```
dept_code VARCHAR2 (30) NOT NULL ,
user_no VARCHAR2 (30) NOT NULL ,
sales_amt int
)
PARTITION BY hash (month_code) SUBPARTITION BY range (dept_code)
(
 PARTITION p_201901
 (
 SUBPARTITION p_201901_a VALUES LESS THAN ('2'),
 SUBPARTITION p_201901_b VALUES LESS THAN ('3')
),
 PARTITION p_201902
 (
 SUBPARTITION p_201902_a VALUES LESS THAN ('2'),
 SUBPARTITION p_201902_b VALUES LESS THAN ('3')
)
);
insert into hash_range values('201901', '1', '1', 1);
insert into hash_range values('201901', '2', '1', 1);
insert into hash_range values('201901', '1', '1', 1);
insert into hash_range values('201903', '2', '1', 1);
insert into hash_range values('201903', '1', '1', 1);
insert into hash_range values('201903', '2', '1', 1);
select * from hash_range;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 1 | 1 | 1
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201901 | 1 | 1 | 1
201901 | 1 | 1 | 1
201901 | 2 | 1 | 1
(6 rows)
```

- Example 2: Specify partitions in a level-2 partitioned table using DML.

```
CREATE TABLE range_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY RANGE (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES LESS THAN('201903')
 (
 SUBPARTITION p_201901_a values ('1'),
 SUBPARTITION p_201901_b values ('2')
),
 PARTITION p_201902 VALUES LESS THAN('201910')
 (
 SUBPARTITION p_201902_a values ('1'),
 SUBPARTITION p_201902_b values ('2')
)
);
-- Insert data to a specified level-1 partition.
insert into range_list partition (p_201901) values('201902', '1', '1', 1);
-- The actual partition is inconsistent with the specified partition. An error is reported.
insert into range_list partition (p_201902) values('201902', '1', '1', 1);
ERROR: inserted partition key does not map to the table partition
DETAIL: N/A.
-- Insert data to a specified level-2 partition.
insert into range_list subpartition (p_201901_a) values('201902', '1', '1', 1);
-- The actual partition is inconsistent with the specified partition. An error is reported.
insert into range_list subpartition (p_201901_b) values('201902', '1', '1', 1);
ERROR: inserted subpartition key does not map to the table subpartition
DETAIL: N/A.
insert into range_list partition for ('201902') values('201902', '1', '1', 1);
insert into range_list subpartition for ('201902','1') values('201902', '1', '1', 1);
```

```
-- Query data in a specified partition.
select * from range_list partition (p_201901);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(4 rows)

select * from range_list subpartition (p_201901_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(4 rows)

select * from range_list partition for ('201902');
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(4 rows)

select * from range_list subpartition for ('201902','1');
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(4 rows)

-- Update data in a specified partition.
update range_list partition (p_201901) set user_no = '2';
select * from range_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 2 | 1
201902 | 1 | 2 | 1
201902 | 1 | 2 | 1
201902 | 1 | 2 | 1
(4 rows)

update range_list subpartition (p_201901_a) set user_no = '3';
select * from range_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 3 | 1
201902 | 1 | 3 | 1
201902 | 1 | 3 | 1
201902 | 1 | 3 | 1
(4 rows)

update range_list partition for ('201902') set user_no = '4';
select * from range_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
(4 rows)

update range_list subpartition for ('201902','2') set user_no = '5';
openGauss=# select *from range_list;
month_code | dept_code | user_no | sales_amt
```

```
-----+-----+-----+-----
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
201902 | 1 | 4 | 1
(4 rows)
select * from range_list;

-- Delete data from a specified partition.
delete from range_list partition (p_201901);
DELETE 4
delete from range_list partition for ('201903');
DELETE 0
delete from range_list subpartition (p_201901_a);
DELETE 0
delete from range_list subpartition for ('201903','2');
DELETE 0

-- Insert data into a specified partition.
insert into range_list partition (p_201901) values('201902', '1', '1', 1) ON DUPLICATE KEY UPDATE
sales_amt = 5;
insert into range_list subpartition (p_201901_a) values('201902', '1', '1', 1) ON DUPLICATE KEY
UPDATE sales_amt = 10;
insert into range_list partition for ('201902') values('201902', '1', '1', 1) ON DUPLICATE KEY UPDATE
sales_amt = 30;
insert into range_list subpartition for ('201902','1') values('201902', '1', '1', 1) ON DUPLICATE KEY
UPDATE sales_amt = 40;
select * from range_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(4 rows)

-- Merge data into a specified partition.
CREATE TABLE newrange_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY RANGE (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES LESS THAN('201903')
 (
 SUBPARTITION p_201901_a values ('1'),
 SUBPARTITION p_201901_b values ('2')
),
 PARTITION p_201902 VALUES LESS THAN('201910')
 (
 SUBPARTITION p_201902_a values ('1'),
 SUBPARTITION p_201902_b values ('2')
)
);
insert into newrange_list values('201902', '1', '1', 1);
insert into newrange_list values('201903', '1', '1', 2);

MERGE INTO range_list partition (p_201901) p
USING newrange_list partition (p_201901) np
ON p.month_code= np.month_code
WHEN MATCHED THEN
 UPDATE SET dept_code = np.dept_code, user_no = np.user_no, sales_amt = np.sales_amt
WHEN NOT MATCHED THEN
 INSERT VALUES (np.month_code, np.dept_code, np.user_no, np.sales_amt);

select * from range_list;
```

| month_code | dept_code | user_no | sales_amt |
|------------|-----------|---------|-----------|
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |

(4 rows)

```
MERGE INTO range_list partition for ('201901') p
USING newrange_list partition for ('201901') np
ON p.month_code= np.month_code
WHEN MATCHED THEN
 UPDATE SET dept_code = np.dept_code, user_no = np.user_no, sales_amt = np.sales_amt
WHEN NOT MATCHED THEN
 INSERT VALUES (np.month_code, np.dept_code, np.user_no, np.sales_amt);
```

| month_code | dept_code | user_no | sales_amt |
|------------|-----------|---------|-----------|
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |

(4 rows)

```
MERGE INTO range_list subpartition (p_201901_a) p
USING newrange_list subpartition (p_201901_a) np
ON p.month_code= np.month_code
WHEN MATCHED THEN
 UPDATE SET dept_code = np.dept_code, user_no = np.user_no, sales_amt = np.sales_amt
WHEN NOT MATCHED THEN
 INSERT VALUES (np.month_code, np.dept_code, np.user_no, np.sales_amt);
```

| month_code | dept_code | user_no | sales_amt |
|------------|-----------|---------|-----------|
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |

(4 rows)

```
MERGE INTO range_list subpartition for ('201901', '1') p
USING newrange_list subpartition for ('201901', '1') np
ON p.month_code= np.month_code
WHEN MATCHED THEN
 UPDATE SET dept_code = np.dept_code, user_no = np.user_no, sales_amt = np.sales_amt
WHEN NOT MATCHED THEN
 INSERT VALUES (np.month_code, np.dept_code, np.user_no, np.sales_amt);
```

| month_code | dept_code | user_no | sales_amt |
|------------|-----------|---------|-----------|
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |
| 201902     | 1         | 1       | 1         |

(4 rows)

- **Example 3: Truncate a level-2 partitioned table.**

```
CREATE TABLE list_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY LIST (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES ('201902')
```

```
(
 SUBPARTITION p_201901_a VALUES ('1'),
 SUBPARTITION p_201901_b VALUES (default)
),
PARTITION p_201902 VALUES ('201903')
(
 SUBPARTITION p_201902_a VALUES ('1'),
 SUBPARTITION p_201902_b VALUES ('2')
)
);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201902', '2', '1', 1);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
insert into list_list values('201903', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
select * from list_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(6 rows)

select * from list_list partition (p_201901);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(3 rows)

alter table list_list truncate partition p_201901;
select * from list_list partition (p_201901);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)

select * from list_list partition (p_201902);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
(3 rows)

alter table list_list truncate partition p_201902;
select * from list_list partition (p_201902);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)

select * from list_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)

insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201902', '2', '1', 1);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
insert into list_list values('201903', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
select * from list_list subpartition (p_201901_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
```

```
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(2 rows)
```

```
alter table list_list truncate subpartition p_201901_a;
select * from list_list subpartition (p_201901_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)
```

```
select * from list_list subpartition (p_201901_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
(1 row)
```

```
alter table list_list truncate subpartition p_201901_b;
select * from list_list subpartition (p_201901_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)
```

```
select * from list_list subpartition (p_201902_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 1 | 1 | 1
(1 row)
```

```
alter table list_list truncate subpartition p_201902_a;
select * from list_list subpartition (p_201902_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)
```

```
select * from list_list subpartition (p_201902_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
(2 rows)
```

```
alter table list_list truncate subpartition p_201902_b;
select * from list_list subpartition (p_201902_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)
```

```
select * from list_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)
```

```
drop table list_list;
```

- **Example 4: Split a level-2 partitioned table.**

```
CREATE TABLE list_list
(
 month_code VARCHAR2 (30) NOT NULL ,
 dept_code VARCHAR2 (30) NOT NULL ,
 user_no VARCHAR2 (30) NOT NULL ,
 sales_amt int
)
PARTITION BY LIST (month_code) SUBPARTITION BY LIST (dept_code)
(
 PARTITION p_201901 VALUES ('201902')
 (
 SUBPARTITION p_201901_a VALUES ('1'),
 SUBPARTITION p_201901_b VALUES (default)
),
 PARTITION p_201902 VALUES ('201903')
```



```
(
 SUBPARTITION p_201902_a VALUES ('1'),
 SUBPARTITION p_201902_b VALUES (default)
);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201902', '2', '1', 1);
insert into list_list values('201902', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
insert into list_list values('201903', '1', '1', 1);
insert into list_list values('201903', '2', '1', 1);
select * from list_list;
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
201903 | 1 | 1 | 1
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(6 rows)

select * from list_list subpartition (p_201901_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(2 rows)

select * from list_list subpartition (p_201901_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
(1 row)

alter table list_list split subpartition p_201901_b values (2) into
(
 subpartition p_201901_b,
 subpartition p_201901_c
);
select * from list_list subpartition (p_201901_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(2 rows)

select * from list_list subpartition (p_201901_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
(1 row)

select * from list_list subpartition (p_201901_c);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)

select * from list_list partition (p_201901);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201902 | 2 | 1 | 1
201902 | 1 | 1 | 1
201902 | 1 | 1 | 1
(3 rows)

select * from list_list subpartition (p_201902_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
```

```
201903 | 1 | 1 | 1
(1 row)

select * from list_list subpartition (p_201902_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
(2 rows)

alter table list_list split subpartition p_201902_b values (3) into
(
 subpartition p_201902_b,
 subpartition p_201902_c
);
select * from list_list subpartition (p_201902_a);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 1 | 1 | 1
(1 row)

select * from list_list subpartition (p_201902_b);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
(0 rows)

select * from list_list subpartition (p_201902_c);
month_code | dept_code | user_no | sales_amt
-----+-----+-----+-----
201903 | 2 | 1 | 1
201903 | 2 | 1 | 1
(2 rows)

drop table list_list;
```

## 7.13.76 CREATE TRIGGER

### Description

Creates a trigger. The trigger will be associated with the specified table or view, and will execute the specified function operations are performed.

### Precautions

- Currently, triggers can be created only on ordinary row-store tables, instead of on temporary tables or unlogged tables.
- If multiple triggers of the same kind are defined for the same event, they will be fired in alphabetical order by name.
- Triggers are usually used for data association and synchronization between multiple tables. SQL execution performance is greatly affected. Therefore, you are advised not to use this statement when a large amount of data needs to be synchronized and performance requirements are high.

### Syntax

```
CREATE [CONSTRAINT] TRIGGER name { BEFORE | AFTER | INSTEAD OF } { event [OR ...] }
ON table_name
[FROM referenced_table_name]
{ NOT DEFERRABLE | [DEFERRABLE] } { INITIALLY IMMEDIATE | INITIALLY DEFERRED } }
[FOR [EACH] { ROW | STATEMENT }]
[WHEN (condition)]
EXECUTE PROCEDURE function_name (arguments);
```

Events include:

```
INSERT
UPDATE [OF column_name [, ...]]
DELETE
TRUNCATE
```

## Parameters

- **CONSTRAINT**

(Optional) Creates a constraint trigger. That is, the trigger is used as a constraint. This is the same as a regular trigger except that the timing of the trigger firing can be adjusted using SET CONSTRAINTS. Constraint triggers must be AFTER ROW triggers.

- **name**

Specifies the name of the trigger to be created. This must be distinct from the name of any other trigger for the same table. The name cannot be schema-qualified — the trigger inherits the schema of its table. For a constraint trigger, this is also the name to use when modifying the trigger's behavior using [SET CONSTRAINTS](#).

Value range: a string, which complies with the [Identifier Naming Conventions](#). A value can contain a maximum of 63 characters.

- **BEFORE**

Specifies that the function is called before the event.

- **AFTER**

Specifies that the function is called after the event. A constraint trigger can only be specified as **AFTER**.

- **INSTEAD OF**

Specifies that the function is called instead of the event.

- **event**

Specifies the event that will fire the trigger. Values are **INSERT**, **UPDATE**, **DELETE**, and **TRUNCATE**. Multiple events can be specified using **OR**.

For **UPDATE** events, it is possible to specify a list of columns using this syntax:

```
UPDATE OF column_name1 [, column_name2 ...]
```

The trigger will only fire if at least one of the listed columns is mentioned as a target of the UPDATE statement. INSTEAD OF UPDATE events do not support lists of columns. If the column specified by UPDATE OF contains a generated column, the trigger is also fired when the column on which the generated column depends is the target column of the UPDATE statement.

- **table\_name**

Specifies the name of the table for which the trigger is created.

Value range: name of an existing table in the database.

- **referenced\_table\_name**

Specifies the name of another table referenced by the constraint. This option is used for foreign-key constraints. It can only be specified for constraint triggers. Because foreign keys are not supported currently, this option is not recommended for general use.

Value range: name of an existing table in the database.

- **DEFERRABLE | NOT DEFERRABLE**

Specifies the start time of the trigger. It can only be specified for constraint triggers. They determine whether the constraint is deferrable.

For details, see [CREATE TABLE](#).

- **INITIALLY IMMEDIATE | INITIALLY DEFERRED**

If the constraint is deferrable, the two clauses specify the default time to check the constraint. It can only be specified for constraint triggers.

For details, see [CREATE TABLE](#).

- **FOR EACH ROW | FOR EACH STATEMENT**

Specifies the frequency of firing the trigger.

- **FOR EACH ROW** indicates that the trigger should be fired once for every row affected by the trigger event.
- **FOR EACH STATEMENT** indicates that the trigger should be fired just once per SQL statement.

If neither is specified, the default is **FOR EACH STATEMENT**. Constraint triggers can only be marked as **FOR EACH ROW**.

- **condition**

Specifies whether the trigger function will actually be executed. If **WHEN** is specified, the function will be called only when **condition** returns **true**.

In FOR EACH ROW triggers, the WHEN condition can refer to columns of the old and/or new row values by writing **OLD.column\_name** or **NEW.column\_name** respectively. In addition, INSERT triggers cannot refer to OLD, and DELETE triggers cannot refer to NEW.

INSTEAD OF triggers do not support WHEN conditions.

Currently, WHEN expressions cannot contain subqueries.

Note that for constraint triggers, evaluation of the WHEN condition is not deferred, but occurs immediately after the row update operation is performed. If the condition does not evaluate to **true**, then the trigger is not queued for deferred execution.

- **function\_name**

Specifies a user-defined function, which must be declared as taking no parameters and returning type trigger. This is executed when a trigger fires.

- **arguments**

Specifies an optional comma-separated list of parameters to be provided to the function when the trigger is executed. The parameters are literal string constants. Simple names and numeric constants can also be written here, but they will all be converted to strings. Check the description of the implementation language of the trigger function to find out how these parameters can be accessed within the function.

#### **NOTE**

The following details trigger types:

- INSTEAD OF triggers must be marked as FOR EACH ROW and can be defined only on views.
- BEFORE and AFTER triggers on a view must be marked as FOR EACH STATEMENT.
- TRUNCATE triggers must be marked as FOR EACH STATEMENT.

**Table 7-106** Types of triggers supported on tables and views

| When       | Event                | Row-Level      | Statement-Level  |
|------------|----------------------|----------------|------------------|
| BEFORE     | INSERT/UPDATE/DELETE | Tables         | Tables and views |
|            | TRUNCATE             | Not supported. | Tables           |
| AFTER      | INSERT/UPDATE/DELETE | Tables         | Tables and views |
|            | TRUNCATE             | Not supported. | Tables           |
| INSTEAD OF | INSERT/UPDATE/DELETE | Views          | Not supported.   |
|            | TRUNCATE             | Not supported. | Not supported.   |

**Table 7-107** Special variables of PL/pgSQL trigger functions

| Variable        | Description                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------|
| NEW             | New tuple for INSERT and UPDATE operations. This variable is null for DELETE operations.                 |
| OLD             | Old tuple for UPDATE and DELETE operations. This variable is null for INSERT operations.                 |
| TG_NAME         | Trigger name.                                                                                            |
| TG_WHEN         | Trigger timing ( <b>BEFORE</b> , <b>AFTER</b> , or <b>INSTEAD OF</b> ).                                  |
| TG_LEVEL        | Trigger frequency ( <b>ROW</b> or <b>STATEMENT</b> ).                                                    |
| TG_OP           | Trigger event ( <b>INSERT</b> , <b>UPDATE</b> , <b>DELETE</b> , or <b>TRUNCATE</b> ).                    |
| TG_RELID        | OID of the table where the trigger resides.                                                              |
| TG_RELNAME      | Name of the table where the trigger resides. (This variable has been replaced by <b>TG_TABLE_NAME</b> .) |
| TG_TABLE_NAME   | Name of the table where the trigger resides.                                                             |
| TG_TABLE_SCHEMA | Schema of the table where the trigger resides.                                                           |

| Variable  | Description                                    |
|-----------|------------------------------------------------|
| TG_NARGS  | Number of parameters for the trigger function. |
| TG_ARGV[] | List of parameters for the trigger function.   |

## Examples

```
-- Create a source table and a destination table.
openGauss=# CREATE TABLE test_trigger_src_tbl(id1 INT, id2 INT, id3 INT);
openGauss=# CREATE TABLE test_trigger_des_tbl(id1 INT, id2 INT, id3 INT);

-- Create a trigger function.
openGauss=# CREATE OR REPLACE FUNCTION tri_insert_func() RETURNS TRIGGER AS
$$
DECLARE
BEGIN
 INSERT INTO test_trigger_des_tbl VALUES(NEW.id1, NEW.id2, NEW.id3);
 RETURN NEW;
END
$$ LANGUAGE plpgsql;

openGauss=# CREATE OR REPLACE FUNCTION tri_update_func() RETURNS TRIGGER AS
$$
DECLARE
BEGIN
 UPDATE test_trigger_des_tbl SET id3 = NEW.id3 WHERE id1=OLD.id1;
 RETURN OLD;
END
$$ LANGUAGE plpgsql;

openGauss=# CREATE OR REPLACE FUNCTION TRI_DELETE_FUNC() RETURNS TRIGGER AS
$$
DECLARE
BEGIN
 DELETE FROM test_trigger_des_tbl WHERE id1=OLD.id1;
 RETURN OLD;
END
$$ LANGUAGE plpgsql;

-- Create an INSERT trigger.
openGauss=# CREATE TRIGGER insert_trigger
BEFORE INSERT ON test_trigger_src_tbl
FOR EACH ROW
EXECUTE PROCEDURE tri_insert_func();

-- Create an UPDATE trigger.
openGauss=# CREATE TRIGGER update_trigger
AFTER UPDATE ON test_trigger_src_tbl
FOR EACH ROW
EXECUTE PROCEDURE tri_update_func();

-- Create a DELETE trigger.
openGauss=# CREATE TRIGGER delete_trigger
BEFORE DELETE ON test_trigger_src_tbl
FOR EACH ROW
EXECUTE PROCEDURE tri_delete_func();

-- Execute the INSERT event and check the trigger results.
openGauss=# INSERT INTO test_trigger_src_tbl VALUES(100,200,300);
openGauss=# SELECT * FROM test_trigger_src_tbl;
openGauss=# SELECT * FROM test_trigger_des_tbl; // Check whether the trigger operation takes effect.
```

```
-- Execute the UPDATE event and check the trigger results.
openGauss=# UPDATE test_trigger_src_tbl SET id3=400 WHERE id1=100;
openGauss=# SELECT * FROM test_trigger_src_tbl;
openGauss=# SELECT * FROM test_trigger_des_tbl; // Check whether the trigger operation takes effect.

-- Execute the DELETE event and check the trigger results.
openGauss=# DELETE FROM test_trigger_src_tbl WHERE id1=100;
openGauss=# SELECT * FROM test_trigger_src_tbl;
openGauss=# SELECT * FROM test_trigger_des_tbl; // Check whether the trigger operation takes effect.

-- Modify a trigger.
openGauss=# ALTER TRIGGER delete_trigger ON test_trigger_src_tbl RENAME TO delete_trigger_renamed;

-- Disable insert_trigger.
openGauss=# ALTER TABLE test_trigger_src_tbl DISABLE TRIGGER insert_trigger;

-- Disable all triggers on the current table.
openGauss=# ALTER TABLE test_trigger_src_tbl DISABLE TRIGGER ALL;

-- Drop triggers.
openGauss=# DROP TRIGGER insert_trigger ON test_trigger_src_tbl;
openGauss=# DROP TRIGGER update_trigger ON test_trigger_src_tbl;
openGauss=# DROP TRIGGER delete_trigger_renamed ON test_trigger_src_tbl;
```

## Helpful Links

[ALTER TRIGGER](#), [DROP TRIGGER](#), and [ALTER TABLE](#)

## 7.13.77 CREATE TYPE

### Description

Defines a new data type for use in the current database. The user who defines a type becomes its owner. Types are designed only for row-store tables.

The following data types can be created: composite type, base type, shell type, enumerated type, and set type.

- **Composite type**

A composite type is specified by a list of attribute names and data types. If the data type of an attribute is collatable, the attribute's collation rule can also be specified. This is essentially the same as the row type of a table, but using CREATE TYPE avoids the need to create an actual table when all that is wanted is to define a type. A stand-alone composite type is useful as the parameter or return type of a function.

To create a composite type, you must have the USAGE permission on all of its attribute types.

- **Base type**

You can create a base type (scalar type). Generally, these functions must be written in the underlying language.

- **Shell type**

A shell type is simply a placeholder for a type to be defined later; it is created by issuing CREATE TYPE with no parameters except for the type name. Shell types are needed as forward references when base types are created.

- **Enumerated type**

An enumerated type is a list of one or more quoted labels, each of which must be 1 to 63 bytes long.

- **Set type**  
It is similar to an array but has no length limit. It is mainly used in stored procedures.
- A user granted with the CREATE ANY TYPE permission can create types in the public and user schemas.

## Precautions

If a schema name is given then the type is created in the specified schema. Otherwise, it is created in the current schema. The type name must be distinct from the name of any existing type or domain in the same schema. (Because tables have associated data types, the type name must also be distinct from the name of any existing table in the same schema.)

## Syntax

```
CREATE TYPE name AS
 ([attribute_name data_type [COLLATE collation] [, ...]])

CREATE TYPE name (
 INPUT = input_function,
 OUTPUT = output_function
 [, RECEIVE = receive_function]
 [, SEND = send_function]
 [, TYPMOD_IN =
type_modifier_input_function]
 [, TYPMOD_OUT =
type_modifier_output_function]
 [, ANALYZE = analyze_function]
 [, INTERNALLENGTH = { internallength |
VARIABLE }]
 [, PASSEDBYVALUE]
 [, ALIGNMENT = alignment]
 [, STORAGE = storage]
 [, LIKE = like_type]
 [, CATEGORY = category]
 [, PREFERRED = preferred]
 [, DEFAULT = default]
 [, ELEMENT = element]
 [, DELIMITER = delimiter]
 [, COLLATABLE = collatable]
)

CREATE TYPE name

CREATE TYPE name AS ENUM
 (['label' [, ...]])

CREATE TYPE name AS TABLE OF data_type
```

## Parameters

Composite type

- **name**  
Specifies the name (optionally schema-qualified) of the type to be created.
- **attribute\_name**  
Specifies the name of an attribute (column) for the composite type.
- **data\_type**



Specifies the name of an existing data type to become a column of the composite type. You can use **%ROWTYPE** to indirectly reference the type of a table, or **%TYPE** to indirectly reference the type of a column in a table or composite type.

- **collation**

Specifies the name of an existing collation rule to be associated with a column of the composite type. You can run the **select \* from pg\_collation** command to query collation rules from the pg\_collation system catalog. The default collation rule is the row starting with **default** in the query result.

#### Base type

When creating a base type, you can place parameters in any order. The **input\_function** and **output\_function** parameters are required, and other parameters are optional.

- **input\_function**

Specifies the name of a function that converts data from the type's external textual form to its internal form.

The input function may be declared as taking one parameter of type cstring or taking three parameters of types cstring, oid, and integer.

- The cstring-type parameter is the input text as a C string.
- The oid-type parameter is the type's own OID (except for array types, which instead receive their element type's OID).
- The integer-type parameter is typmod of the destination column, if known (–1 will be passed if not known).

The input function must return a value of the data type itself. Usually, an input function should be declared as **STRICT**. Otherwise, when a **NULL** input value is read and the input function is called, the first parameter is **NULL**. The function must still return **NULL** in this case, unless it raises an error. (This case is mainly meant to support domain input functions, which might need to reject **NULL** inputs.)

#### NOTE

The input and output functions can be declared to have results or parameters of the new type, when they have to be created before the new type can be created. The type should first be defined as a shell type, which is a placeholder type that has no attributes except a name and an owner. This is done by issuing the **CREATE TYPE name** statement, with no additional parameters. Then, the I/O functions written in C can be defined as referencing the shell type. Finally, **CREATE TYPE** with a full definition replaces the shell entry with a complete, valid type definition, after which the new type can be used normally.

- **output\_function**

Specifies the name of a function that converts data from the type's internal form to its external textual form.

The output function must be declared as taking one parameter of the new data type. The output function must return type cstring. Output functions are not called for **NULL** values.

- **receive\_function**

(Optional) Specifies the name of a function that converts data from the type's external binary form to its internal form.

If this function is not supplied, the type cannot participate in binary input. The binary representation should be chosen to be cheap to convert to internal form, while being reasonably portable. (For example, the standard integer data types use network byte order as the external binary representation, while the internal representation is in the machine's native byte order.)

**receive\_function** should perform adequate checking to ensure that the value is valid.

The receive function may be declared as taking one parameter of type internal or taking three parameters of types internal, oid, integer.

- The internal-type parameter is a pointer to a StringInfo buffer holding the received byte strings.
- The oid- and integer-type parameters are the same as those of the text input function.

The receive function must return a value of the data type itself. Usually, a receive function should be declared as **STRICT**. Otherwise, when a **NULL** input value is read and the receive function is called, the first parameter is **NULL**. The function must still return **NULL** in this case, unless it raises an error. (This case is mainly meant to support domain receive functions, which might need to reject **NULL** inputs.)

- **send\_function**

(Optional) Specifies the name of a function that converts data from the type's internal form to its external binary form.

If this function is not supplied, the type cannot participate in binary output. The send function must be declared as taking one parameter of the new data type. The send function must return type bytea. Send functions are not called for **NULL** values.

- **type\_modifier\_input\_function**

(Optional) Specifies the name of a function that converts an array of modifiers for a type to its internal format.

- **type\_modifier\_output\_function**

(Optional) Specifies the name of a function that converts the internal format of modifiers for a type to its external text format.

#### NOTE

**type\_modifier\_input\_function** and **type\_modifier\_output\_function** are needed if the type supports modifiers, that is optional constraints attached to a type declaration, such as char(5) or numeric(30,2). GaussDB allows user-defined types to take one or more simple constants or identifiers as modifiers. However, this information must be capable of being packed into a single non-negative integer value for storage in the system catalogs. Declared modifiers are passed to **type\_modifier\_input\_function** in the cstring array format. It must check the values for validity (throwing an error if they are wrong), and if they are correct, return a single non-negative integer value that will be stored as the column "typmod". Type modifiers will be rejected if the type does not have a **type\_modifier\_input\_function**. The **type\_modifier\_output\_function** converts the internal integer typmod value back to the correct form for user display.

**type\_modifier\_output\_function** must return a cstring value that is the exact string to append to the type name; for example, numeric's function might return (30,2). It is allowed to omit the **type\_modifier\_output\_function**, in which case the default display format is just the stored typmod integer value enclosed in parentheses.

- **analyze\_function**

(Optional) Specifies the name of a function that performs statistical analysis for the data type.

By default, ANALYZE will attempt to gather statistics using the type's "equals" and "less-than" operators, if there is a default B-tree operator class for the type. For non-scalar types, this behavior is likely to be unsuitable, so it can be overridden by specifying a custom analysis function. The analysis function must be declared to take one parameter of type internal and return a Boolean result.

- **internallength**

(Optional) Specifies the length in bytes of the new type's internal representation. The default assumption is that it is variable-length.

While the details of the new type's internal representation are only known to the I/O functions and other functions you create to work with the type, there are several attributes of the internal representation that must be declared to GaussDB. Foremost of these is **internallength**. Base data types can be fixed-length, in which case **internallength** is a positive integer, or variable length, indicated by setting **internallength** to **VARIABLE**. (Internally, this is represented by setting **typlen** to **-1**.) The internal representation of all variable-length types must start with a 4-byte integer giving the total length of this value of the type.

- **PASSEDBYVALUE**

(Optional) Indicates that values of this data type are passed by value, rather than by reference. You cannot pass by value types whose internal representation is larger than the size of the Datum type (4 bytes on most machines, 8 bytes on a few).

- **alignment**

(Optional) Specifies the storage alignment requirement of the data type. If specified, it must be **char**, **int2**, **int4**, or **double**; the default is **int4**.

The allowed values equate to alignment on 1, 2, 4, or 8 byte boundaries. Note that variable-length types must have an alignment of at least 4, since they necessarily contain an int4 as their first component.

- **storage**

(Optional) Specifies the storage strategy for the data type.

If specified, it must be **plain**, **external**, **extended**, or **main**. The default is **plain**.

- **plain** specifies that data of the type will always be stored in-line and not compressed. (Only **plain** is allowed for fixed-length types.)
- **extended** specifies that the system will first try to compress a long data value, and will move the value out of the main table row if it is still too long.
- **external** allows the value to be moved out of the main table, but the system will not try to compress it.
- **main** allows compression, but discourages moving the value out of the main table. (Data items with this storage strategy might still be moved out of the main table if there is no other way to make a row fit, but they will be kept in the main table preferentially over **extended** and **external** items.)

All **storage** values other than **plain** imply that the functions of the data type can handle values that have been toasted. The specific other value

given merely determines the default TOAST storage strategy for columns of a toastable data type; users can pick other strategies for individual columns using ALTER TABLE SET STORAGE.

- **like\_type**

(Optional) Specifies the name of an existing data type that the new type will have the same representation as. The values of **internallength**, **passedbyvalue**, **alignment**, and **storage** are copied from that type, unless overridden by explicit specification elsewhere in this CREATE TYPE statement. Specifying representation in this way is especially useful when the underlying implementation of a new type references an existing type.

- **category**

(Optional) Specifies the category code (a single ASCII character) for this type. The default is **U** for a user-defined type. You may also choose other ASCII characters to create custom categories.

- **preferred**

Specifies whether a type is preferred within its type category. If it is, the value will be **TRUE**; otherwise, the value is **FALSE**. This parameter is optional. The default value is **FALSE**. Be very careful about creating a preferred type within an existing type category, as this could cause surprising changes in behavior.

 **NOTE**

The **category** and **preferred** parameters can be used to help control which implicit cast will be applied in ambiguous situations. Each data type belongs to a category named by a single ASCII character, and each type is either preferred or not within its category. The parser will prefer casting to preferred types (but only from other types within the same category) when this rule is helpful in resolving overloaded functions or operators. For types that have no implicit casts to or from any other types, it is sufficient to leave these settings at the defaults. However, for a group of related types that have implicit casts, it is often helpful to mark them all as belonging to a category and select one or two of the most general types as being preferred within the category. The **category** parameter is especially useful when adding a user-defined type to an existing built-in category, such as the numeric or string types. However, it is also possible to create entirely-user-defined type categories. Select any ASCII character other than an uppercase letter to name such a category.

- **default**

(Optional) Specifies the default value for the data type. If this is omitted, the default is null.

A default value can be specified, in case a user wants columns of the data type to default to something other than the null value. Specify the default with the DEFAULT keyword. (Such a default can be overridden by an explicit DEFAULT clause attached to a particular column.)

- **element**

(Optional) Specifies the type of array elements when an array type is created. For example, to define an array of 4-byte integers (int4), specify **ELEMENT = int4**.

- **delimiter**

(Optional) Specifies the delimiter character to be used between values in arrays made of this type.

**delimiter** can be set to a specific character. The default delimiter is the comma (,). Note that the delimiter is associated with the array element type, not the array type itself.

- **collatable**

(Optional) Specifies whether this type's operations can use collation information. If they can, the value will be **TRUE**, else **FALSE** (default).

If **collatable** is **TRUE**, column definitions and expressions of the type may carry collation information through use of the **COLLATE** clause. It is up to the implementations of the functions operating on the type to actually make use of the collation information; this does not happen automatically merely by marking the type collatable.

- **label**

(Optional) Represents the textual label associated with one value of an enumerated type. It is a string of 1 to 63 characters.

 **NOTE**

Whenever a user-defined type is created, GaussDB automatically creates an associated array type whose name consists of the element type's name prepended with an underscore (\_).

## Examples

```
-- Create a composite type, create a table, insert data, and make a query.
openGauss=# CREATE TYPE compfoo AS (f1 int, f2 text);
openGauss=# CREATE TABLE t1_compfoo(a int, b compfoo);
openGauss=# CREATE TABLE t2_compfoo(a int, b compfoo);
openGauss=# INSERT INTO t1_compfoo values(1,(1,'demo'));
openGauss=# INSERT INTO t2_compfoo select * from t1_compfoo;
openGauss=# SELECT (b).f1 FROM t1_compfoo;
openGauss=# SELECT * FROM t1_compfoo t1 join t2_compfoo t2 on (t1.b).f1=(t2.b).f1;

-- Rename the data type.
openGauss=# ALTER TYPE compfoo RENAME TO compfoo1;

-- Change the owner of the user-defined type compfoo1 to usr1.
openGauss=# CREATE USER usr1 PASSWORD '*****';
openGauss=# ALTER TYPE compfoo1 OWNER TO usr1;

-- Change the schema of the user-defined type compfoo1 to usr1.
openGauss=# ALTER TYPE compfoo1 SET SCHEMA usr1;

-- Add a new attribute to the data type.
openGauss=# ALTER TYPE usr1.compfoo1 ADD ATTRIBUTE f3 int;

-- Drop the compfoo1 type.
openGauss=# DROP TYPE usr1.compfoo1 cascade;

-- Drop related tables and users.
openGauss=# DROP TABLE t1_compfoo;
openGauss=# DROP TABLE t2_compfoo;
openGauss=# DROP SCHEMA usr1;
openGauss=# DROP USER usr1;

-- Create an enumerated type.
openGauss=# CREATE TYPE bugstatus AS ENUM ('create', 'modify', 'closed');

-- Add a label.
openGauss=# ALTER TYPE bugstatus ADD VALUE IF NOT EXISTS 'regress' BEFORE 'closed';

-- Rename a label.
openGauss=# ALTER TYPE bugstatus RENAME VALUE 'create' TO 'new';

-- Create a set type.
openGauss=# CREATE TYPE compfoo_table AS TABLE OF compfoo;
```

## Helpful Links

[ALTER TYPE](#) and [DROP TYPE](#)

## 7.13.78 CREATE USER

### Description

Creates a user.

### Precautions

- A user created using the CREATE USER statement has the LOGIN permission by default.
- When you run the **CREATE USER** command to create a user, the system creates a schema with the same name as the user in the database where the command is executed.
- The owner of an object created by a system administrator in a schema with the same name as a common user is the common user, not the system administrator.

### Syntax

```
CREATE USER user_name [[WITH] option [...]] [ENCRYPTED | UNENCRYPTED] { PASSWORD | IDENTIFIED BY } { 'password' [EXPIRED] | DISABLE };
```

The option clause is used to configure information, including permissions and properties.

```
{SYSADMIN | NOSYSADMIN}
| {MONADMIN | NOMONADMIN}
| {OPRADMIN | NOOPRADMIN}
| {POLADMIN | NOPOLADMIN}
| {AUDITADMIN | NOAUDITADMIN}
| {CREATEDB | NOCREATEDB}
| {USEFT | NOUSEFT}
| {CREATEROLE | NOCREATEROLE}
| {INHERIT | NOINHERIT}
| {LOGIN | NOLOGIN}
| {REPLICATION | NOREPLICATION}
| {VCADMIN | NOVCADMIN}
| {PERSISTENCE | NOPERSISTENCE}
| CONNECTION LIMIT connlimit
| VALID BEGIN 'timestamp'
| VALID UNTIL 'timestamp'
| USER GROUP 'groupuser'
| PERM SPACE 'spacelimit'
| TEMP SPACE 'tmpspacelimit'
| SPILL SPACE 'spillspacelimit'
| NODE GROUP logic_cluster_name
| IN ROLE role_name [, ...]
| IN GROUP role_name [, ...]
| ROLE role_name [, ...]
| ADMIN role_name [, ...]
| USER role_name [, ...]
| SYSID uid
| DEFAULT TABLESPACE tablespace_name
| PROFILE DEFAULT
| PROFILE profile_name
| PGUSER
```

## Parameters

- **user\_name**

Username.

Value range: a string that follows the [Identifier Naming Conventions](#) with a maximum of 63 characters.

- **password**

Specifies the login password.

The new password must:

- Contain at least eight characters. This is the default length.
- Differ from the username or the username spelled backward.
- Contain at least three of the following character types: uppercase characters, lowercase characters, digits, and special characters (limited to ~!@#\$%^&\*()-\_+=\|[{ }];;<.>/?).
- The password can also be a ciphertext character string that meets the format requirements. This mode is mainly used to import user data. You are advised not to use it directly. If a ciphertext password is used, you need to know the plaintext corresponding to the ciphertext password and ensure a complex plaintext password. The database does not verify the complexity of the ciphertext password, so you should ensure the password security.
- When creating a user, enclose the user password in single quotation marks.

Value range: a string.

For details about other parameter values of CREATE USER, see [CREATE ROLE](#).

## Examples

```
-- Create user jim whose login password is *****.
openGauss=# CREATE USER jim PASSWORD '*****';

-- Alternatively, you can run the following statement:
openGauss=# CREATE USER kim IDENTIFIED BY '*****';

-- To create a user with the CREATEDB permission, add the CREATEDB keyword.
openGauss=# CREATE USER dim CREATEDB PASSWORD '*****';

-- Change the login password of user jim from ***** to *****.
openGauss=# ALTER USER jim IDENTIFIED BY '*****' REPLACE '*****';

-- Add the CREATEROLE permission to jim.
openGauss=# ALTER USER jim CREATEROLE;

-- Set enable_seqscan to on. (The setting will take effect in the next session.)
openGauss=# ALTER USER jim SET enable_seqscan TO on;

-- Reset the enable_seqscan parameter for jim.
openGauss=# ALTER USER jim RESET enable_seqscan;

-- Lock jim.
openGauss=# ALTER USER jim ACCOUNT LOCK;

-- Drop users.
openGauss=# DROP USER kim CASCADE;
openGauss=# DROP USER jim CASCADE;
openGauss=# DROP USER dim CASCADE;
```

## Helpful Links

[ALTER USER](#), [CREATE ROLE](#), and [DROP USER](#)

## 7.13.79 CREATE USER MAPPING

### Description

Defines a new mapping from a user to a foreign server.

### Precautions

If the **password** option is displayed, ensure that the **usermapping.key.cipher** and **usermapping.key.rand** files exist in the *\$GAUSSHOME/bin* directory of each node in GaussDB. If the two files do not exist, use the *gs\\_guc* tool to generate them and use the *gs\\_ssh* tool to release them to the *\$GAUSSHOME/bin* directory on each node in GaussDB.

When multi-layer quotation marks are used for sensitive columns (such as **password**) in **OPTIONS**, the semantics is different from that in the scenario where quotation marks are not used. Therefore, sensitive columns are not identified for anonymization.

### Syntax

```
CREATE USER MAPPING FOR { user_name | USER | CURRENT_USER | PUBLIC }
 SERVER server_name
 [OPTIONS (option 'value' [, ...])]
```

### Parameters

- **user\_name**  
Specifies the name of an existing user to map to a foreign server.  
**CURRENT\_USER** and **USER** match the name of the current user. When **PUBLIC** is specified, a public mapping is created and used when no mapping for a particular user is available.
- **server\_name**  
Specifies the name of the existing server for which a user mapping will be created.
- **OPTIONS ( { option\_name ' value ' } [ , ... ] )**  
Specifies options for user mapping. These options typically define the actual username and password for this mapping. The option name must be unique. The allowed option names and values are related to the foreign data wrapper of the server.



 NOTE

- User passwords are encrypted and stored in the system catalog **PG\_USER\_MAPPING**. During the encryption, **usermapping.key.cipher** and **usermapping.key.rand** are used as the encryption password file and encryption factor. Before using the tool for the first time, run the following command to create the two files, save the files to the **\$GAUSSHOME/bin** directory on each node, and ensure that you have the read permission on the files. **gs\_ssh** helps you quickly place files in the specified directory of each node.  
`gs_ssh -c "gs_guc generate -o usermapping -S default -D $GAUSSHOME/bin"`
- If the **-S** parameter is set to default, a password is randomly generated. You can also specify a password for the **-S** parameter to ensure the security and uniqueness of the generated password file. You do not need to save or memorize the password. For details about other parameters, see the description of the **gs\_guc** tool in the "Tool Reference".

## Examples

```
-- Create a role.
openGauss=# CREATE ROLE bob PASSWORD '*****';

-- Create a foreign server.
openGauss=# CREATE SERVER my_server FOREIGN DATA WRAPPER log_fdw;

-- Create a user mapping.
openGauss=# CREATE USER MAPPING FOR bob SERVER my_server OPTIONS (user 'bob', password '*****');

-- Modify the user mapping.
openGauss=# ALTER USER MAPPING FOR bob SERVER my_server OPTIONS (SET password '*****');

-- Drop the user mapping.
openGauss=# DROP USER MAPPING FOR bob SERVER my_server;

-- Drop the foreign server.
openGauss=# DROP SERVER my_server;

-- Drop the role.
openGauss=# DROP ROLE bob;
```

## Helpful Links

[ALTER USER MAPPING](#) and [DROP USER MAPPING](#)

## 7.13.80 CREATE VIEW

### Description

Creates a view. A view is a virtual table, not a base table. Only view definition is stored in the database and view data is not. The data is stored in a base table. If data in the base table changes, the data in the view changes accordingly. In this sense, a view is like a window through which users can know their interested data and data changes in the database.

### Precautions

A user granted with the CREATE ANY TABLE permission can create views in the public and user schemas.

## Syntax

```
CREATE [OR REPLACE] [TEMP | TEMPORARY] VIEW view_name [(column_name [, ...])]
[WITH ({view_option_name [= view_option_value]} [, ...])]
AS query;
```

### NOTE

You can use WITH(security\_barrier) to create a relatively secure view. This prevents attackers from printing hidden base table data by using the RAISE statement of low-cost functions.

After a view is created, you are not allowed to use REPLACE to modify column names in the view or delete the columns.

## Parameters

- **OR REPLACE**  
Redefines the view if it already exists.
- **TEMP | TEMPORARY**  
Creates a temporary view.
- **view\_name**  
Specifies the name (optionally schema-qualified) of the view to be created.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Specifies an optional list of names to be used for columns of the view. If not given, the column names are deduced from the query.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **view\_option\_name [= view\_option\_value]**  
Specifies an optional parameter for a view.  
Currently, **view\_option\_name** supports only the **security\_barrier** parameter.  
Use this parameter when the view provides row-level security.  
Value range: Boolean type, **TRUE** or **FALSE**.
- **query**  
Specifies a SELECT or VALUES statement that will provide the columns and rows of the view.

---

### NOTICE

If **query** contains a clause specifying the partition of a partitioned table, the OID of the specified partition is hardcoded to the system catalog when the view is created. If a DDL operation that will change the OID of the specified partition is used, for example, DROP, SPLIT, or MERGE, the view is unavailable. In this case, you need to create a view.

---

## Examples

```
-- Create a view consisting of columns whose spcname is pg_default.
openGauss=# CREATE VIEW myView AS
SELECT * FROM pg_tablespace WHERE spcname = 'pg_default';

-- Query the view.
```

```
openGauss=# SELECT * FROM myView ;

-- Drop the myView view.
openGauss=# DROP VIEW myView;
```

## Helpful Links

[ALTER VIEW](#) and [DROP VIEW](#)

## 7.13.81 CREATE WEAK PASSWORD DICTIONARY

### Description

Inserts one or more weak passwords into the **gs\_global\_config** table.

### Precautions

- Only the initial user and users with the SYSADMIN or CREATEROLE permission can execute this syntax.
- Passwords in the weak password dictionary are stored in the **gs\_global\_config** system catalog.
- The weak password dictionary is empty by default. You can use this syntax to add one or more weak passwords.
- When a user attempts to execute this syntax to insert a weak password that already exists in the **gs\_global\_config** table, only one weak password is retained in the table.

### Syntax

```
CREATE WEAK PASSWORD DICTIONARY
[WITH VALUES] ({'weak_password'} [, ...]);
```

### Parameters

**weak\_password**

Specifies a weak password.

Value range: a character string

### Examples

```
-- Insert a single weak password into the gs_global_config system catalog.
openGauss=# CREATE WEAK PASSWORD DICTIONARY WITH VALUES ('password1');

-- Insert multiple weak passwords into the gs_global_config system catalog.
openGauss=# CREATE WEAK PASSWORD DICTIONARY WITH VALUES ('password2'),('password3');

-- Clear all weak passwords in the gs_global_config system catalog.
openGauss=# DROP WEAK PASSWORD DICTIONARY;

-- View existing weak passwords.
openGauss=# SELECT * FROM gs_global_config WHERE NAME LIKE 'weak_password';
```

## Helpful Links

[DROP WEAK PASSWORD DICTIONARY](#)

## 7.13.82 CURSOR

### Function

**CURSOR** defines a cursor to retrieve a small number of rows at a time out of a larger query.

To process SQL statements, the stored procedure process assigns a memory segment to store context association. Cursors are handles or pointers pointing to context regions. With cursors, stored procedures can control alterations in context regions.

### Precautions

- **CURSOR** is used only in transaction blocks.
- Generally, **CURSOR** and **SELECT** both have text returns. Since data is stored in binary format in the system, the system needs to convert the data from the binary format to the text format. If data is returned in text format, client applications need to convert the data back to the binary format for processing. **FETCH** implements conversion between binary data and text data.
- Binary cursors should be used carefully. Text usually occupies larger space than binary data. A binary cursor returns internal binary data, which is easier to operate. A text cursor returns text, which is easier to retrieve and therefore reduces workload on the client. As an example, if a query returns a value of one from an integer column, you would get a string of 1 with a default cursor, whereas with a binary cursor you would get a 4-byte field containing the internal representation of the value (in big-endian byte order).

### Syntax

```
CURSOR cursor_name
[BINARY] [NO SCROLL] [{ WITH | WITHOUT } HOLD]
FOR query ;
```

### Parameter Description

- **cursor\_name**  
Specifies the name of the cursor to be created.  
Value range: a string. It must comply with the naming convention.
- **BINARY**  
Causes the cursor to return data in binary rather than in text format.
- **NO SCROLL**  
Specifies how the cursor retrieves rows.
  - **NO SCROLL**: specifies that the cursor cannot be used to retrieve rows in a nonsequential fashion.
  - Unspecified: Based on the query's execution plan, the system automatically determines whether the cursor can be used to retrieve rows in a nonsequential fashion.
- **WITH HOLD | WITHOUT HOLD**  
Specifies whether the cursor can continue to be used after the transaction that created it successfully commits.

- **WITH HOLD**: The cursor can continue to be used after the transaction that created it successfully commits.
- **WITHOUT HOLD**: The cursor cannot be used outside of the transaction that created it.
- If neither **WITH HOLD** nor **WITHOUT HOLD** is specified, the default is **WITHOUT HOLD**.
- Cross-node transactions (for example, DDL-contained transactions created in the database with multiple DBnode) do not support **WITH HOLD**.
- **query**  
Uses a **SELECT** or **VALUES** clause to specify the rows to be returned by the cursor.  
Value range: **SELECT** or **VALUES** clause

## Examples

See [Examples](#) in **FETCH**.

## Helpful Links

[FETCH](#)

## 7.13.83 DEALLOCATE

### Description

**DEALLOCATE** deallocates a previously prepared statement. If you do not explicitly deallocate a prepared statement, it is deallocated when the session ends.

The **PREPARE** keyword is always ignored.

### Precautions

None

### Syntax

```
DEALLOCATE [PREPARE] { name | ALL };
```

### Parameters

- **name**  
Specifies the name of the prepared statement to be deallocated.
- **ALL**  
Deallocates all prepared statements.

## Examples

None

## 7.13.84 DECLARE

### Description

**DECLARE** defines a cursor to retrieve a small number of rows at a time out of a larger query and can be the start of an anonymous block.

This section describes usage of cursors. The usage of anonymous blocks is available in [BEGIN](#).

To process SQL statements, the stored procedure process assigns a memory segment to store context association. Cursors are handles or pointers pointing to context regions. With cursors, stored procedures can control alterations in context regions.

Generally, **CURSOR** and **SELECT** both have text returns. Since data is stored in binary format in the system, the system needs to convert the data from the binary format to the text format. If data is returned in text format, client applications need to convert the data back to the binary format for processing. **FETCH** implements conversion between binary data and text data.

### Precautions

- **CURSOR** is used only in transaction blocks.
- Binary cursors should be used carefully. Text usually occupies larger space than binary data. A binary cursor returns internal binary data, which is easier to operate. A text cursor returns text, which is easier to retrieve and therefore reduces workload on the client. As an example, if a query returns a value of one from an integer column, you would get a string of 1 with a default cursor, whereas with a binary cursor you would get a 4-byte field containing the internal representation of the value (in big-endian byte order).

### Syntax

- Define a cursor.  

```
DECLARE cursor_name [BINARY] [NO SCROLL]
 CURSOR [{ WITH | WITHOUT } HOLD] FOR query ;
```
- Enable an anonymous block.  

```
[DECLARE [declare_statements]]
BEGIN
 execution_statements
END;
/
```

### Parameters

- **cursor\_name**  
Specifies the name of the cursor to be created.  
Value range: a string. It must comply with the naming convention.
- **BINARY**  
Causes the cursor to return data in binary rather than in text format.
- **NO SCROLL**  
Specifies how the cursor retrieves rows.

- **NO SCROLL**: specifies that the cursor cannot be used to retrieve rows in a nonsequential fashion.
- Unspecified: Based on the query's execution plan, the system automatically determines whether the cursor can be used to retrieve rows in a nonsequential fashion.
- **WITH HOLD**  
**WITHOUT HOLD**  
Specifies whether the cursor can continue to be used after the transaction that created it successfully commits.
  - **WITH HOLD**: The cursor can continue to be used after the transaction that created it successfully commits.
  - **WITHOUT HOLD**: The cursor cannot be used outside of the transaction that created it.
  - If neither **WITH HOLD** nor **WITHOUT HOLD** is specified, the default is **WITHOUT HOLD**.
- **query**  
Uses a **SELECT** or **VALUES** clause to specify the rows to be returned by the cursor.  
Value range: **SELECT** or **VALUES** clause
- **declare\_statements**  
Declares a variable, including its name and type, for example, **sales\_cnt int**.
- **execution\_statements**  
Specifies the statement to be executed in an anonymous block.  
Value range: an existing function name

## Examples

For details about how to define a cursor, see [Examples](#) in "FETCH."

## Helpful Links

[BEGIN](#), [FETCH](#)

## 7.13.85 DELETE

### Description

Deletes rows that satisfy the WHERE clause from the specified table. If the WHERE clause is absent, the effect is to delete all rows in the table. The result is a valid, but an empty table.

### Precautions

The owner of a table, users granted with the DELETE permission on the table, or users granted with the DELETE ANY TABLE permission can delete data from the table. System administrators have the permission to delete data from the table by default, as well as the SELECT permission on any table in the USING clause or whose values are read in **condition**.

## Syntax

```
[WITH [RECURSIVE] with_query [, ...]]
DELETE [/*+ plan_hint */] [FROM] [ONLY] table_name [partition_clause] [*] [[AS] alias]
 [USING using_list]
 [WHERE condition | WHERE CURRENT OF cursor_name]
 [LIMIT { row_count }]
 [RETURNING { * | { output_expr [[AS] output_name] } [, ...] }];
```

Format of **with\_query**:

```
with_query_name [(column_name [, ...])] AS [[NOT] MATERIALIZED]
({select | values | insert | update | delete})
```

## Parameters

- **WITH [ RECURSIVE ] with\_query [, ...]**

Specifies one or more subqueries that can be referenced by name in the main query, which is equivalent to a temporary table.

If **RECURSIVE** is specified, it allows a SELECT subquery to reference itself by name.

  - **with\_query\_name** specifies the name of the result set generated by a subquery. Such names can be used to access the result sets of subqueries in a query.
  - **column\_name** specifies the column name displayed in the subquery result set.
  - Each subquery can be a SELECT, VALUES, INSERT, UPDATE, or DELETE statement.
  - You can use **MATERIALIZED** or **NOT MATERIALIZED** to modify the CTE.
  - If **MATERIALIZED** is specified, the WITH query will be materialized, and a copy of the subquery result set is generated. The copy is directly queried at the reference point. Therefore, the WITH subquery cannot be jointly optimized with the SELECT statement trunk (for example, predicate pushdown and equivalence class transfer). In this scenario, you can use **NOT MATERIALIZED** for modification. If the WITH query can be executed as a subquery inline, the preceding optimization can be performed.
  - If the user does not explicitly declare the materialized attribute, comply with the following rules: If the CTE is referenced only once in the trunk statement to which it belongs and semantically supports inline execution, it will be rewritten as subquery inline execution. Otherwise, the materialized execution will be performed in CTE Scan mode.
- **plan\_hint clause**

Follows the DELETE keyword in the */\*+ \*/* format. It is used to optimize the plan of a DELETE statement block. For details, see [Hint-based Tuning](#). In each statement, only the first */\*+ plan\_hint \*/* comment block takes effect as a hint. Multiple hints can be written.
- **ONLY**

If **ONLY** is specified before the table name, matching rows are deleted from the named table only. If **ONLY** is not specified, matching rows are also deleted from any tables inheriting from the named table.
- **table\_name**

Specifies the name (optionally schema-qualified) of the target table.



Value range: an existing table name.

- **partition\_clause**

Deletes a specified partition.

PARTITION { ( partition\_name ) | FOR ( partition\_value [, ...] ) } |

SUBPARTITION { ( subpartition\_name ) | FOR ( subpartition\_value [, ...] ) }

For details about the keywords, see [SELECT](#).

For details, see [CREATE TABLE SUBPARTITION](#).

- **alias**

Specifies the alias of the target table.

Value range: a string that complies with the [Identifier Naming Conventions](#).

- **using\_list**

Specifies a USING clause.

- **condition**

Specifies an expression that returns a Boolean value. Only rows for which this expression returns **true** will be deleted. You are advised not to use numeric types such as int for **condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and 0 is implicitly converted to **false**), which may cause unexpected results.

- **WHERE CURRENT OF cursor\_name**

This parameter is reserved.

- **LIMIT clause**

For details about the keywords, see [SELECT](#).

- **output\_expr**

Specifies an expression to be computed and returned by the DELETE statement after each row is deleted. The expression can use any column names of the table. Write \* to return all columns.

- **output\_name**

Specifies a name to use for a returned column.

Value range: a string that complies with the [Identifier Naming Conventions](#).

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer_address table.
openGauss=# CREATE TABLE tpcds.customer_address
(
 ca_address_sk INTEGER NOT NULL,
 ca_address_id CHARACTER(16) NOT NULL,
 ca_street_number INTEGER ,
 ca_street_name CHARACTER (20)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.customer_address VALUES (1, 'AAAAAAAAABAAAAAAA', '18', 'Jackson'),
(10000, 'AAAAAAAACAAAAAAA', '362', 'Washington 6th'),(15000, 'AAAAAAAADAAAAAAA', '585', 'Dogwood
Washington');

-- Create the tpcds.customer_address_bak table.
openGauss=# CREATE TABLE tpcds.customer_address_bak AS TABLE tpcds.customer_address;
```

```
-- Delete employees whose ca_address_sk is smaller than 14888 from the tpcds.customer_address_bak table.
openGauss=# DELETE FROM tpcds.customer_address_bak WHERE ca_address_sk < 14888;

-- Delete all data from the tpcds.customer_address_bak table.
openGauss=# DELETE FROM tpcds.customer_address_bak;

-- Drop the tpcds.customer_address_bak table.
openGauss=# DROP TABLE tpcds.customer_address_bak;

-- Drop the tpcds.customer_address table.
openGauss=# DROP TABLE tpcds.customer_address;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Suggestions

- **delete**  
To delete all records in a table, use the truncate syntax.

## 7.13.86 DO

### Description

Executes an anonymous code block.

A code block is regarded as a function body without parameters. The return value type is void. It is parsed and executed a single time.

### Precautions

- The procedural language to be used must already have been installed into the current database by means of CREATE LANGUAGE. PL/pgSQL is installed by default, but other languages are not.
- To use an untrusted language, you must have the USAGE permission on the programming language or the SYSADMIN permission.

### Syntax

```
DO [LANGUAGE lang_name] code;
```

### Parameters

- **lang\_name**  
Specifies the name of the procedural language the code is written in. If omitted, the default language is PL/pgSQL.
- **code**  
Specifies the programming language code that can be executed. The value must be a character string.

### Examples

```
-- Create the webuser user.
openGauss=# CREATE USER webuser PASSWORD 'xxxxxxxxx';

-- Grant all permissions on all views in the tpcds schema to the webuser user.
```

```
openGauss=# DO $$DECLARE r record;
BEGIN
 FOR r IN SELECT c.relname table_name,n.nspname table_schema FROM pg_class c,pg_namespace n
 WHERE c.relnamespace = n.oid AND n.nspname = 'tpcds' AND relkind IN ('r','v')
 LOOP
 EXECUTE 'GRANT ALL ON ' || quote_ident(r.table_schema) || '.' || quote_ident(r.table_name) || ' TO
webuser';
 END LOOP;
END$$;

-- Drop the webuser user.
openGauss=# DROP USER webuser CASCADE;
```

## 7.13.87 DROP AGGREGATE

### Function

**DROP AGGREGATE** deletes an aggregate function.

### Precautions

**DROP AGGREGATE** deletes an existing aggregate function. Only the owner of the aggregate function can run this command.

### Syntax

```
DROP AGGREGATE [IF EXISTS] name (argtype [, ...]) [CASCADE | RESTRICT]
```

### Parameter Description

- **IF EXISTS**  
Do not throw an error if the specified aggregation does not exist. A notice is issued in this case.
- **name**  
Existing aggregate function name (optionally schema-qualified).
- **argtype**  
Input data type of the aggregate function. To reference a zero-parameter aggregate function, use \* to replace the input data type list.
- **CASCADE**  
Cascade deletes objects that depend on the aggregate function.
- **RESTRICT**  
Refuses to delete the aggregate function if any objects depend on it. This is a default processing.

### Examples

Delete the aggregate function **myavg** of the integer type:

```
DROP AGGREGATE myavg(integer);
```

### Compatibility

The SQL standard does not provide the **DROP AGGREGATE** statement.

## 7.13.88 DROP AUDIT POLICY

### Description

Drops an audit policy.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

### Syntax

```
DROP AUDIT POLICY [IF EXISTS] policy_name;
```

### Parameters

- **policy\_name**  
Specifies the audit policy name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).

### Examples

See [Examples](#) in "CREATE AUDIT POLICY."

### Helpful Links

[ALTER AUDIT POLICY](#) and [CREATE AUDIT POLICY](#)

## 7.13.89 DROP CAST

### Description

Drops a type conversion.

### Precautions

DROP CAST drops a previously defined type conversion.

To drop a type conversion, you must have permissions on the source or destination data type. This is the same permission as creating a type conversion.

### Syntax

```
DROP CAST [IF EXISTS] (source_type AS target_type) [CASCADE | RESTRICT]
```

### Parameters

- **IF EXISTS**  
Do not throw an error if the specified conversion does not exist. A notice is issued in this case.
- **source\_type**  
Source data type in the type conversion.

- **target\_type**  
Type of the target data in the type conversion.
- **CASCADE**  
**RESTRICT**  
These keys have no effect because there is no dependency on type conversion.

## Examples

Drop the conversion from text to int.

```
DROP CAST (text AS int);
```

## Compatibility

DROP CAST complies with the SQL standard.

## 7.13.90 DROP DATABASE

### Description

Drops a database.

### Precautions

- Only the database owner or a user granted the DROP permission can run the **DROP DATABASE** command. System administrators have this permission by default.
- The preinstalled POSTGRES, TEMPLATE0, and TEMPLATE1 databases are protected and therefore cannot be deleted. To check databases in the current service, run the gsql statement **\l**.
- If any users are connected to the database, the database cannot be deleted.
- **DROP DATABASE** cannot be executed within a transaction block.
- If **DROP DATABASE** fails and is rolled back, run **DROP DATABASE IF EXISTS** again.

---

#### NOTICE

**DROP DATABASE** cannot be undone.

---

### Syntax

```
DROP DATABASE [IF EXISTS] database_name ;
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified database does not exist.
- **database\_name**  
Specifies the name of the database to be deleted.

Value range: an existing database name

## Examples

See [Examples](#) in CREATE DATABASE.

## Helpful Links

[CREATE DATABASE](#)

## Suggestions

- drop database  
Do not delete databases during transactions.

# 7.13.91 DROP DIRECTORY

## Function

**DROP Directory** deletes a synonym.

## Precautions

When **enable\_access\_server\_directory** is set to **off**, only the initial user is allowed to delete directory objects. When **enable\_access\_server\_directory** is set to **on**, a user with the SYSADMIN permission, the owner of the directory object, a user who is granted with the DROP permission of the directory, or a user who inherits the **gs\_role\_directory\_drop** permission of the built-in role can delete directory objects.

## Syntax

```
DROP DIRECTORY [IF EXISTS] directory_name;
```

## Parameter Description

- **directory\_name**  
Specifies the name of the directory to be deleted.  
Value range: an existing directory name

## Examples

```
-- Create a directory.
openGauss=# CREATE OR REPLACE DIRECTORY dir as '/tmp/';

-- Delete a directory.
openGauss=# DROP DIRECTORY dir;
```

## Helpful Links

[CREATE DIRECTORY](#) and [ALTER DIRECTORY](#)

## 7.13.92 DROP FUNCTION

### Description

Drops an existing function.

### Precautions

- If a function involves operations on temporary tables, DROP FUNCTION cannot be used.
- Only the function owner or a user granted with the DROP permission can run the **DROP FUNCTION** command. System administrators have this permission by default.

### Syntax

```
DROP FUNCTION [IF EXISTS] function_name
[(([argname] [argmode] argtype) [, ...])) [CASCADE | RESTRICT] ;
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified function does not exist.
- **function\_name**  
Specifies the name of the function to be dropped.  
Value range: an existing function name.
- **argmode**  
Specifies the parameter mode of the function.
- **argname**  
Specifies the parameter name of the function.
- **argtype**  
Specifies the parameter type of the function.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops the objects that depend on the function.
  - **RESTRICT**: refuses to drop the function if any objects depend on it. This is the default action.

### Examples

For details, see [Examples](#).

### Helpful Links

[ALTER FUNCTION](#) and [CREATE FUNCTION](#)

## 7.13.93 DROP GLOBAL CONFIGURATION

### Description

Drops parameter values from the gs\_global\_config system catalog.

## Precautions

- Only the initial database user can run this command.
- The parameter name cannot be **weak\_password** or **undostoragetype**.

## Syntax

```
DROP GLOBAL CONFIGURATION name [, ...];
```

## Parameters

- **name**  
The parameter must exist in the gs\_global\_config system catalog. If you drop a parameter that does not exist, an error will be reported.

## Helpful Links

[ALTER GLOBAL CONFIGURATION](#)

## 7.13.94 DROP GROUP

### Description

Drops a user group.

DROP GROUP is an alias for DROP ROLE.

### Precautions

DROP GROUP is an API of the GaussDB management tool. You are advised not to use this API, because doing so affects GaussDB.

### Syntax

```
DROP GROUP [IF EXISTS] group_name [, ...];
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified role does not exist.
- **group\_name**  
Specifies the name of the role to be dropped.  
Value range: an existing role name.

### Examples

See [Examples](#) in "CREATE ROLE."

### Helpful Links

[CREATE GROUP](#), [ALTER GROUP](#), and [DROP ROLE](#)



## 7.13.95 DROP INDEX

### Description

Drops an index.

### Precautions

Only the index owner, a user of a schema where the index resides, or a user who has the INDEX permission on the table where the index resides can run the **DROP INDEX** command. System administrators have this permission by default.

For a global temporary table, if a session has initialized a global temporary table object (including creating a global temporary table and inserting data into the global temporary table for the first time), other sessions cannot drop indexes from the table.

### Syntax

```
DROP INDEX [CONCURRENTLY] [IF EXISTS]
index_name [, ...] [CASCADE | RESTRICT];
```

### Parameters

- **CONCURRENTLY**  
Drops an index without locking it. When an index is dropped, other statements cannot access the table on which the index depends. With this option, the statement does not lock the table during index deletion.  
This parameter allows only one index name and does not support CASCADE.  
The **DROP INDEX** statement can be run within a transaction, but **DROP INDEX CONCURRENTLY** cannot.
- **IF EXISTS**  
Reports a notice instead of an error if the specified index does not exist.
- **index\_name**  
Specifies the name of the index to be dropped.  
Value range: name of an existing index.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops the objects that depend on the index.
  - **RESTRICT**: refuses to drop the index if any objects depend on it. This is the default action.

### Examples

See [Examples](#) in "CREATE INDEX."

### Helpful Links

[ALTER INDEX](#) and [CREATE INDEX](#)

## 7.13.96 DROP LANGUAGE

This version does not support this syntax.

## 7.13.97 DROP MASKING POLICY

### Description

Drops a masking policy.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

### Syntax

```
DROP MASKING POLICY [IF EXISTS] policy_name;
```

### Parameters

- **policy\_name**  
Specifies the audit policy name, which must be unique.  
Value range: a string that complies with the [Identifier Naming Conventions](#).  
The value is an existing policy name.

### Examples

```
-- Drop a masking policy.
openGauss=# DROP MASKING POLICY IF EXISTS maskpol1;

-- Drop a group of masking policies.
openGauss=# DROP MASKING POLICY IF EXISTS maskpol1, maskpol2, maskpol3;
```

### Helpful Links

[ALTER MASKING POLICY](#) and [CREATE MASKING POLICY](#)

## 7.13.98 DROP MATERIALIZED VIEW

### Function

Forcibly deletes an existing materialized view from the database.

### Precautions

The owner of a materialized view, owner of the schema of the materialized view, users granted with the DROP permission on the materialized view, or users granted with the DROP ANY TABLE permission can run the **DROP MATERIALIZED VIEW** command. By default, the system administrator has the permission to run the command.

### Syntax

```
DROP MATERIALIZED VIEW [IF EXISTS] mv_name [, ...] [CASCADE | RESTRICT];
```

## Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the specified materialized view does not exist.
- **mv\_name**  
Name of the materialized view to be deleted.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically deletes the objects that depend on the materialized view.
  - **RESTRICT**: refuses to delete the materialized view if any objects depend on it. This is the default action.

## Examples

```
-- Create a table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int);

-- Create a materialized view named my_mv.
openGauss=# CREATE MATERIALIZED VIEW my_mv AS SELECT * FROM my_table;

-- Delete the materialized view named my_mv.
openGauss=# DROP MATERIALIZED VIEW my_mv;

-- Delete the table.
openGauss=# DROP TABLE my_table;
```

## Helpful Links

[ALTER MATERIALIZED VIEW](#), [CREATE INCREMENTAL MATERIALIZED VIEW](#), [CREATE MATERIALIZED VIEW](#), [CREATE TABLE](#), [REFRESH INCREMENTAL MATERIALIZED VIEW](#), and [REFRESH MATERIALIZED VIEW](#)

## 7.13.99 DROP OPERATOR

### Description

Drops an operator object.

### Precautions

DROP OPERATOR drops the definition of an operator. To use DROP OPERATOR, you must be the owner of the operator.

### Syntax

```
DROP OPERATOR [IF EXISTS] name ({ lefttype | NONE } , { righttype | NONE }) [CASCADE | RESTRICT]
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the operator does not exist.
- **name**  
Name (optionally schema-qualified) of an existing operator.

- **lefttype**  
Type of the left operand of an operator. If there is no left operand, **NONE** is used.
- **righttype**  
Type of the right operand of an operator. If there is no right operand, **NONE** is used.
- **CASCADE**  
Automatically drops objects that depend on an operator (such as views that use the operator), and then drops all objects that depend on those objects.
- **RESTRICT**  
Refuses to drop any object that depends on an operator. This is the default value.

## Helpful Links

[ALTER OPERATOR](#) and [CREATE OPERATOR](#)

## 7.13.100 DROP OWNED

### Description

Drops the database objects owned by a database role.

### Precautions

- The role's permissions on all objects in the current database and shared objects (databases and tablespaces) will be revoked.
- DROP OWNED is often used to prepare for removing one or more roles. Because DROP OWNED affects only the objects in the current database, you need to run this statement in each database that contains the objects owned by the role to be removed.
- Using the **CASCADE** option may cause this statement to recursively remove objects owned by other users.
- The databases and tablespaces owned by the role will not be removed.

### Syntax

```
DROP OWNED BY name [, ...] [CASCADE | RESTRICT];
```

### Parameters

- **name**  
Specifies the role name.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops the objects that depend on the objects to be dropped.
  - **RESTRICT**: refuses to drop the objects if other objects depend on them. This is the default action.

## Helpful Links

[REASSIGN OWNED](#) and [DROP ROLE](#)

### 7.13.101 DROP PACKAGE

#### Description

Drops an existing package or package body.

#### Precautions

After the package body is dropped, the stored procedures and functions in the package become invalid at the same time.

#### Syntax

```
DROP PACKAGE [IF EXISTS] package_name;
DROP PACKAGE BODY [IF EXISTS] package_name;
```

#### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified stored procedure does not exist.
- **package\_name**  
Specifies the name of the package to be dropped.  
Value range: an existing package name.

#### Examples

```
-- Create a package.
openGauss=# CREATE OR REPLACE PACKAGE PCK1
IS
a int;
END pck1;
/
CREATE PACKAGE

-- Drop the package.
openGauss=# DROP PACKAGE PCK1;
DROP PACKAGE
```

## Helpful Links

[ALTER PACKAGE](#) and [CREATE PACKAGE](#)

### 7.13.102 DROP PROCEDURE

#### Description

Drops a stored procedure.

#### Syntax

```
DROP PROCEDURE [IF EXISTS] procedure_name ;
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified stored procedure does not exist.
- **procedure\_name**  
Specifies the name of the stored procedure to be dropped.  
Value range: an existing stored procedure name.

## Helpful Links

[CREATE PROCEDURE](#)

## 7.13.103 DROP RESOURCE LABEL

### Description

Drops a resource label.

### Precautions

Only users with the POLADMIN or SYSADMIN permission, or the initial user can perform this operation.

### Syntax

```
DROP RESOURCE LABEL [IF EXISTS] label_name[, ...];
```

### Parameters

- **label\_name**  
Specifies the resource label name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).

### Examples

```
-- Drop a resource label.
openGauss=# DROP RESOURCE LABEL IF EXISTS res_label1;

-- Drop a group of resource labels.
openGauss=# DROP RESOURCE LABEL IF EXISTS res_label1, res_label2, res_label3;
```

## Helpful Links

[ALTER RESOURCE LABEL](#) and [CREATE RESOURCE LABEL](#)

## 7.13.104 DROP ROLE

### Description

Drops a role.

## Syntax

```
DROP ROLE [IF EXISTS] role_name [, ...];
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified role does not exist.
- **role\_name**  
Specifies the name of the role to be dropped.  
Value range: an existing role name.

## Examples

See [Examples](#) in "CREATE ROLE."

## Helpful Links

[CREATE ROLE](#), [ALTER ROLE](#), and [SET ROLE](#)

## 7.13.105 DROP ROW LEVEL SECURITY POLICY

### Description

Drops a row-level security policy from a table.

### Precautions

Only the table owner or administrators can drop a row-level security policy from the table.

## Syntax

```
DROP [ROW LEVEL SECURITY] POLICY [IF EXISTS] policy_name ON table_name [CASCADE | RESTRICT]
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified row-level security policy does not exist.
- **policy\_name**  
Specifies the name of the row-level security policy to be dropped.
- **table\_name**  
Specifies the name of the table containing the row-level security policy.
- **CASCADE | RESTRICT**  
Currently, no objects depend on row-level security policies. Therefore, **CASCADE** is equivalent to **RESTRICT**, and they are reserved to ensure backward compatibility.

## Examples

```
-- Create data table all_data.
openGauss=# CREATE TABLE all_data(id int, role varchar(100), data varchar(100));
```

```
-- Create a row-level security policy.
openGauss=# CREATE ROW LEVEL SECURITY POLICY all_data_rls ON all_data USING(role =
CURRENT_USER);

-- Drop a row-level security policy.
openGauss=# DROP ROW LEVEL SECURITY POLICY all_data_rls ON all_data;

-- Drop the all_data table.
openGauss=# DROP TABLE all_data;
```

## Helpful Links

[ALTER ROW LEVEL SECURITY POLICY](#) and [CREATE ROW LEVEL SECURITY POLICY](#)

## 7.13.106 DROP RULE

### Description

Drops a rewriting rule.

### Syntax

```
DROP RULE [IF EXISTS] name ON table_name [CASCADE | RESTRICT]
```

### Parameters

- **IF EXISTS**  
If the rule does not exist, a **NOTICE** is thrown.
- **name**  
Name of an existing rule to be dropped.
- **table\_name**  
Name of the table to which the rule applies.
- **CASCADE**  
Automatically cascade drops objects that depend on this rule.
- **RESTRICT**  
Refuses to drop the rule if any objects depend on it. This is the default value.

### Examples

```
-- Create the def_test table and the def_view_test view for creating rules.
openGauss=# CREATE TABLE def_test (
c1 int4 DEFAULT 5,
c2 text DEFAULT 'initial_default'
);
openGauss=# CREATE VIEW def_view_test AS SELECT * FROM def_test;
-- Create the rule def_view_test_ins.
openGauss=# CREATE RULE def_view_test_ins AS
openGauss=# ON INSERT TO def_view_test
openGauss=# DO INSTEAD INSERT INTO def_test SELECT new.*;
-- Drop the rule def_view_test_ins.
openGauss=# DROP RULE def_view_test_ins ON def_view_test;
-- Drop the def_test table and the def_view_test view.
openGauss=# DROP VIEW def_view_test;
openGauss=# DROP TABLE def_test;
```



## 7.13.107 DROP SCHEMA

### Description

Drops a schema from the current database.

### Precautions

Only the schema owner or a user granted with the DROP permission can run the **DROP SCHEMA** command. System administrators have this permission by default.

### Syntax

```
DROP SCHEMA [IF EXISTS] schema_name [, ...] [CASCADE | RESTRICT];
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified schema does not exist.
- **schema\_name**  
Specifies the schema name.  
Value range: an existing schema name.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops all the objects contained in the schema.
  - **RESTRICT**: refuses to drop the schema if the schema contains objects. This is the default action.

#### NOTICE

- Schemas beginning with **pg\_temp** or **pg\_toast\_temp** are for internal use. Do not drop them. Otherwise, unexpected consequences may be incurred.
- The schema currently being used cannot be dropped. To drop it, switch to another schema first.

### Examples

See [Examples](#) in "CREATE SCHEMA."

### Helpful Links

[ALTER SCHEMA](#) and [CREATE SCHEMA](#)

## 7.13.108 DROP SEQUENCE

### Description

Drops a sequence from the current database.

## Precautions

- Only the owner of a sequence, the owner of the schema of the sequence, or users granted with the DROP permission on the sequence can drop the sequence. By default, system administrators have the permission to drop the sequence.
- If the LARGE identifier is used when a sequence is created, the LARGE identifier must be used when the sequence is dropped.

## Syntax

```
DROP [LARGE] SEQUENCE [IF EXISTS] {[schema.]sequence_name} [, ...] [CASCADE | RESTRICT];
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified sequence does not exist.
- **sequence\_name**  
Specifies the name of the sequence to be dropped.
- **CASCADE**  
Automatically drops the objects that depend on the sequence.
- **RESTRICT**  
Refuses to drop the sequence if any objects depend on it. This is the default action.

## Examples

```
-- Create an ascending sequence named serial, starting from 101.
openGauss=# CREATE SEQUENCE serial START 101;

-- Drop a sequence.
openGauss=# DROP SEQUENCE serial;
```

## Helpful Links

[ALTER SEQUENCE](#) and [DROP SEQUENCE](#)

## 7.13.109 DROP SERVER

### Description

Drops an existing data server.

### Precautions

Only the server owner or a user granted the DROP permission can run the **DROP SERVER** command. System administrators have this permission by default.

### Syntax

```
DROP SERVER [IF EXISTS] server_name [{ CASCADE | RESTRICT }] ;
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified data server does not exist.
- **server\_name**  
Specifies the name of the data server to be dropped.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops the objects that depend on the data server.
  - **RESTRICT**: refuses to drop the server if any objects depend on it. This is the default action.

## Examples

```
-- Create a server.
openGauss=# create server my_server foreign data wrapper log_fdw;
CREATE SERVER

-- Drop my_server.
openGauss=# DROP SERVER my_server;
DROP SERVER
```

## Helpful Links

[ALTER SERVER](#) and [CREATE SERVER](#)

## 7.13.110 DROP SYNONYM

### Description

Drops a synonym.

### Precautions

Only the owner of a synonym or a system administrator has the **DROP SYNONYM** permission.

### Syntax

```
DROP SYNONYM [IF EXISTS] synonym_name [CASCADE | RESTRICT];
```

### Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the specified synonym does not exist.
- **synonym\_name**  
Specifies the name (optionally schema-qualified) of the synonym to be deleted.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically deletes the objects (such as views) that depend on the synonym.
  - **RESTRICT**: refuses to delete the synonym if any objects depend on it. This is the default action.

## Examples

See [Examples](#) in section "CREATE SYNONYM."

## Helpful Links

[ALTER SYNONYM](#), [CREATE SYNONYM](#)

# 7.13.111 DROP TABLE

## Description

Drops a table.

## Precautions

- DROP TABLE forcibly drops the specified table and the indexes depending on the table. After the table is dropped, the functions and stored procedures that need to use this table cannot be executed. Dropping a partitioned table also drops all partitions in the table.
- The owner of a table, the owner of the schema of the table, users granted with the DROP permission on the table, or users granted with the DROP ANY TABLE permission can drop the specified table. System administrators have the permission to drop the specified table by default.

## Syntax

```
DROP TABLE [IF EXISTS]
{ [schema.]table_name } [, ...] [CASCADE | RESTRICT] [PURGE];
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified table does not exist.
- **schema**  
Specifies the schema name.
- **table\_name**  
Specifies the table name.
- **CASCADE | RESTRICT**
  - **CASCADE**: The objects that depend on the table, such as views, triggers, and indexes, can be dropped cascadingly. Note that joined table objects cannot be dropped cascadingly.
  - **RESTRICT**: refuses to drop the table if any objects depend on it. This is the default action.
- **PURGE**  
Specifies that even if the recycle bin function is enabled, the table is physically dropped instead of being moved to the recycle bin.

## Examples

See [Examples](#) in "CREATE TABLE."

## Helpful Links

[ALTER TABLE](#) and [CREATE TABLE](#)

## 7.13.112 DROP TABLESPACE

### Description

Drops a tablespace.

### Precautions

- Only the tablespace owner or a user granted the DROP permission can run the **DROP TABLESPACE** command. A system administrator has this permission by default.
- The tablespace to be deleted should not contain any database objects. Otherwise, an error will be reported.
- **DROP TABLESPACE** cannot be rolled back and therefore cannot be run in transaction blocks.
- During execution of **DROP TABLESPACE**, database queries by other sessions using **\db** may fail and need to be reattempted.
- If **DROP TABLESPACE** fails to be executed, run **DROP TABLESPACE IF EXISTS**.

### Syntax

```
DROP TABLESPACE [IF EXISTS] tablespace_name;
```

### Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified tablespace does not exist.
- **tablespace\_name**  
Specifies the name of the tablespace to be deleted.  
Value range: an existing tablespace name

### Examples

See [Examples](#) in **CREATE TABLESPACE**.

## Helpful Links

[ALTER TABLESPACE](#) and [CREATE TABLESPACE](#)

### Suggestions

- drop tablespace  
Do not delete tablespaces during transactions.

## 7.13.113 DROP TRIGGER

### Function

**DROP TRIGGER** deletes a trigger.

### Precautions

Only the owner of a trigger or a system administrator has the **DROP TRIGGER** permission.

### Syntax

```
DROP TRIGGER [IF EXISTS] trigger_name ON table_name [CASCADE | RESTRICT];
```

### Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the specified trigger does not exist.
- **trigger\_name**  
Specifies the name of the trigger to be deleted.  
Value range: an existing trigger name
- **table\_name**  
Specifies the name of the table containing the trigger.  
Value range: name of the table containing the trigger
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically deletes the objects that depend on the trigger.
  - **RESTRICT**: refuses to delete the trigger if any objects depend on it. This is the default action.

### Examples

For details, see [Examples](#) in [CREATE TRIGGER](#).

### Helpful Links

[CREATE TRIGGER](#), [ALTER TRIGGER](#), and [ALTER TABLE](#)

## 7.13.114 DROP TYPE

### Function

**DROP TYPE** deletes a user-defined data type.

### Precautions

Only the type owner or a user granted with the DROP permission can run the **DROP TYPE** command. The system administrator has this permission by default.

## Syntax

```
DROP TYPE [IF EXISTS] name [, ...] [CASCADE | RESTRICT]
```

## Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the specified type does not exist.
- **name**  
Specifies the name (optionally schema-qualified) of the type to be deleted.
- **CASCADE**  
Automatically deletes the objects (such as fields, functions, and operators) that depend on the type.  
**RESTRICT**  
Refuses to delete the type if any objects depend on it. This is the default action.

## Examples

See [Examples](#) in **CREATE TYPE**.

## Helpful Links

[CREATE TYPE](#) and [ALTER TYPE](#)

## 7.13.115 DROP USER

### Description

Drops a user and the schema with the same name as the user.

### Precautions

- **CASCADE** is used to drop the objects (excluding databases) that depend on the user. It cannot drop locked objects unless the objects are unlocked or the processes locking the objects are terminated.
- In GaussDB, the **enable\_kill\_query** configuration parameter exists in the **postgresql.conf** file. This parameter affects **CASCADE**.
  - If **enable\_kill\_query** is **on** and **CASCADE** is used, the statement automatically kills the processes locking dependent objects and then drops the specified user.
  - If **enable\_kill\_query** is **off** and **CASCADE** is used, the statement waits until the processes locking dependent objects stop and then drops the specified user.
- If the dependent objects are other databases or reside in other databases, manually drop them before dropping the user from the current database. **DROP USER** cannot drop objects across databases.
- Before dropping a user, you need to drop all the objects owned by the user and revoke the user's permissions on other objects. Alternatively, you can specify **CASCADE** to drop the objects owned by the user and the granted permissions.

## Syntax

```
DROP USER [IF EXISTS] user_name [, ...] [CASCADE | RESTRICT];
```

## Parameters

- **IF EXISTS**  
Reports a notice instead of an error if the specified user does not exist.
- **user\_name**  
Specifies the name of the user to be dropped.  
Value range: an existing username.
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically drops objects that depend on the user and revokes the permissions granted to the user.
  - **RESTRICT**: refuses to drop a user if the user has any dependent objects or has been granted permissions on other objects. This is the default value.

### NOTE

In GaussDB, the **enable\_kill\_query** configuration parameter exists in the **postgresql.conf** file. This parameter affects **CASCADE**.

- If **enable\_kill\_query** is **on** and **CASCADE** is used, the statement automatically kills the processes locking dependent objects and then drops the specified user.
- If **enable\_kill\_query** is **off** and **CASCADE** is used, the statement waits until the processes locking dependent objects stop and then drops the specified user.

## Examples

See [Examples](#) in "CREATE USER."

## Helpful Links

[ALTER USER](#) and [CREATE USER](#)

## 7.13.116 DROP USER MAPPING

### Function

**DROP USER MAPPING** drops a user mapping for a foreign server.

### Syntax

```
DROP USER MAPPING [IF EXISTS] FOR { user_name | USER | CURRENT_USER | PUBLIC } SERVER
server_name;
```

### Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the user mapping does not exist.
- **user\_name**  
Specifies username of the mapping.



CURRENT\_USER and USER match the name of the current user. PUBLIC is used to match all current and future usernames in the system.

- **server\_name**  
Specifies name of the server to which the user is mapped.

## Helpful Links

[ALTER USER MAPPING](#) and [CREATE USER MAPPING](#)

## 7.13.117 DROP VIEW

### Function

**DROP VIEW** forcibly deletes a view from the database.

### Precautions

The owner of a view, owner of the schema of the view, users granted with the DROP permission on the view, or users granted with the DROP ANY TABLE permission can run the **DROP VIEW** command. By default, the system administrator has the permission to run the command.

### Syntax

```
DROP VIEW [IF EXISTS] view_name [, ...] [CASCADE | RESTRICT];
```

### Parameter Description

- **IF EXISTS**  
Reports a notice instead of an error if the specified view does not exist.
- **view\_name**  
Specifies the name of the view to be deleted.  
Value range: an existing view name
- **CASCADE | RESTRICT**
  - **CASCADE**: automatically deletes the objects (such as other views) that depend on the view.
  - **RESTRICT**: refuses to delete the view if any objects depend on it. This is the default action.

### Examples

See [Examples](#) in **CREATE VIEW**.

## Helpful Links

[ALTER VIEW](#) and [CREATE VIEW](#)

## 7.13.118 DROP WEAK PASSWORD DICTIONARY

### Description

Clears all weak passwords in gs\_global\_config.

### Precautions

Only the initial user and users with the SYSADMIN or CREATEROLE permission can execute this syntax.

### Syntax

```
DROP WEAK PASSWORD DICTIONARY;
```

### Examples

See [Examples](#) in "CREATE WEAK PASSWORD DICTIONARY."

### Helpful Links

[CREATE WEAK PASSWORD DICTIONARY](#)

## 7.13.119 EXECUTE

### Description

Executes a prepared statement. Because a prepared statement exists only in the lifetime of the session, the prepared statement must be created earlier in the current session by using the PREPARE statement.

### Precautions

If the PREPARE statement declares some parameters when the prepared statement is created, the parameter set passed to the EXECUTE statement must be compatible. Otherwise, an error occurs.

### Syntax

```
EXECUTE name [(parameter [, ...])] ;
```

### Parameters

- **name**  
Specifies the name of the prepared statement to be executed.
- **parameter**  
Specifies a parameter of the prepared statement. It must be of the same data type as that specified parameter in creating and generating the prepared statement.

### Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;
```

```
-- Create the reason table.
openGauss=# CREATE TABLE tpcds.reason (
 CD_DEMO_SK INTEGER NOT NULL,
 CD_GENDER character(16) ,
 CD_MARITAL_STATUS character(100)
);

-- Insert data.
openGauss=# INSERT INTO tpcds.reason VALUES(51, 'AAAAAAAADDAAAAAA', 'reason 51');

-- Create the reason_t1 table.
openGauss=# CREATE TABLE tpcds.reason_t1 AS TABLE tpcds.reason;

-- Create a prepared statement for an INSERT statement and execute the prepared statement.
openGauss=# PREPARE insert_reason(integer,character(16),character(100)) AS INSERT INTO
tpcds.reason_t1 VALUES($1,$2,$3);

openGauss=# EXECUTE insert_reason(52, 'AAAAAAAADDAAAAAA', 'reason 52');

-- Drop the reason and reason_t1 tables.
openGauss=# DROP TABLE tpcds.reason;
openGauss=# DROP TABLE tpcds.reason_t1;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.120 EXPLAIN

### Description

Shows the execution plan of an SQL statement.

The execution plan shows how the tables referenced by the SQL statement will be scanned, for example, by plain sequential scan or index scan. If multiple tables are referenced, the execution plan also shows what join algorithms will be used to bring together the required rows from each input table.

The most critical part of the display is the estimated statement execution cost, which is the plan generator's guess at how long it will take to run the statement.

The **ANALYZE** option causes the statement to be actually executed, not only planned. The total elapsed time expended within each plan node (in milliseconds) and total number of rows it actually returned are added to the display. This is useful for determining whether the plan generator's estimated value is close to the actual one.

### Precautions

- The statement is actually executed when the **ANALYZE** option is used. If you want to use EXPLAIN ANALYZE on an INSERT, UPDATE, DELETE, CREATE TABLE AS, or EXECUTE statement without letting the statement affect your data, use the following approach:

```
START TRANSACTION;
EXPLAIN ANALYZE ...;
ROLLBACK;
```

- The **DETAIL**, **NODES**, and **NUM\_NODES** parameters cannot be used in standalone mode. If the parameters are used, the following error is reported:

```
openGauss=# CREATE TABLE student(id int, name char(20));
CREATE TABLE
openGauss=# EXPLAIN (NODES true) INSERT INTO student VALUES(5,'a'),('6','b');
ERROR: unrecognized EXPLAIN option "nodes"
```

```
openGauss=# EXPLAIN (NUM_NODES true) INSERT INTO student VALUES(5,'a'),(6,'b');
ERROR: unrecognized EXPLAIN option "num_nodes"
```

## Syntax

- Display the execution plan of an SQL statement, which supports multiple options and has no requirements for the order of options.

```
EXPLAIN [(option [, ...])] statement;
```

The syntax of the option clause is as follows:

```
ANALYZE [boolean] |
 ANALYZE [boolean] |
 VERBOSE [boolean] |
 COSTS [boolean] |
 CPU [boolean] |
 DETAIL [boolean] | (Unavailable in the centralized mode)
 NODES [boolean] | (Unavailable in the centralized mode)
 NUM_NODES [boolean] | (Unavailable in the centralized mode)
 BUFFERS [boolean] |
 TIMING [boolean] |
 PLAN [boolean] |
 FORMAT { TEXT | XML | JSON | YAML }
```

- Display the execution plan of an SQL statement, where options are in order.

```
EXPLAIN { [ANALYZE | ANALYSE] [VERBOSE] | PERFORMANCE } statement;
```

## Parameters

- **statement**

Specifies the SQL statement to explain.

- **ANALYZE boolean | ANALYSE boolean**

Specifies whether to display actual run times and other statistics. When two parameters are used at the same time, the latter one in option takes effect.

Value range:

- **TRUE** (default): displays them.
- **FALSE**: does not display them.

- **VERBOSE boolean**

Specifies whether to display additional information regarding the plan.

Value range:

- **TRUE** (default): displays it.
- **FALSE**: does not display it.

- **COSTS boolean**

Specifies whether to display the estimated total cost of each plan node, estimated number of rows, estimated width of each row.

Value range:

- **TRUE** (default): displays them.
- **FALSE**: does not display them.

- **CPU boolean**

Specifies whether to display CPU usage.

Value range:

- **TRUE** (default): displays it.
- **FALSE**: does not display it.

- **DETAIL boolean** (Unavailable in the centralized mode)  
Displays information about database nodes.  
Value range:
  - **TRUE** (default): displays it.
  - **FALSE**: does not display it.
- **NODES boolean** (Unavailable in the centralized mode)  
Specifies whether to display information about the nodes executed by query.  
Value range:
  - **TRUE** (default): displays it.
  - **FALSE**: does not display it.
- **NUM\_NODES boolean** (Unavailable in the centralized mode)  
Specifies whether to display the number of executing nodes.  
Value range:
  - **TRUE** (default): displays it.
  - **FALSE**: does not display it.
- **BUFFERS boolean**  
Specifies whether to display buffer usage.  
Value range:
  - **TRUE**: displays it.
  - **FALSE** (default): does not display it.
- **TIMING boolean**  
Specifies whether to display the actual startup time and time spent on the output node.  
Value range:
  - **TRUE** (default): displays them.
  - **FALSE**: does not display them.
- **PLAN**  
Specifies whether to store the execution plan in **PLAN\_TABLE**. If this parameter is set to **on**, the execution plan is stored in **PLAN\_TABLE** and not displayed on the screen. Therefore, this parameter cannot be used together with other parameters when it is set to **on**.  
Value range:
  - **ON** (default): The execution plan is stored in **PLAN\_TABLE** and not printed on the screen. If the plan is stored successfully, "EXPLAIN SUCCESS" is returned.
  - **OFF**: The execution plan is not stored but is printed on the screen.
- **FORMAT**  
Specifies the output format.  
Value range: **TEXT**, **XML**, **JSON**, and **YAML**  
Default value: **TEXT**
- **PERFORMANCE**  
Prints all relevant information in execution.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer_address table.
openGauss=# CREATE TABLE tpcds.customer_address
(
 ca_address_sk INTEGER NOT NULL,
 ca_address_id CHARACTER(16) NOT NULL
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.customer_address VALUES (5000, 'AAAAAAAABAAAAAAA'),(10000,
'AAAAAAAACAAAAAAA');

-- Create the tpcds.customer_address_p1 table.
openGauss=# CREATE TABLE tpcds.customer_address_p1 AS TABLE tpcds.customer_address;

-- Change the value of explain_perf_mode to normal.
openGauss=# SET explain_perf_mode=normal;

-- Display an execution plan for simple queries in the table.
openGauss=# EXPLAIN SELECT * FROM tpcds.customer_address_p1;
QUERY PLAN

Data Node Scan (cost=0.00..0.00 rows=0 width=0)
Node/s: All dbnodes
(2 rows)

-- Generate an execution plan in JSON format (with explain_perf_mode being normal).
openGauss=# EXPLAIN(FORMAT JSON) SELECT * FROM tpcds.customer_address_p1;
QUERY PLAN

[
 {
 "Plan": {
 "Node Type": "Data Node Scan",+
 "Startup Cost": 0.00, +
 "Total Cost": 0.00, +
 "Plan Rows": 0, +
 "Plan Width": 0, +
 "Node/s": "All dbnodes" +
 }
 }
]
(1 row)

-- If there is an index and we use a query with an indexable WHERE condition, a different plan may be
displayed.
openGauss=# EXPLAIN SELECT * FROM tpcds.customer_address_p1 WHERE ca_address_sk=10000;
QUERY PLAN

Data Node Scan (cost=0.00..0.00 rows=0 width=0)
Node/s: dn_6005_6006
(2 rows)

-- Generate an execution plan in YAML format (with explain_perf_mode being normal).
openGauss=# EXPLAIN(FORMAT YAML) SELECT * FROM tpcds.customer_address_p1 WHERE
ca_address_sk=10000;
QUERY PLAN

- Plan:
 Node Type: "Data Node Scan"+
 Startup Cost: 0.00 +
 Total Cost: 0.00 +
 Plan Rows: 0 +
 Plan Width: 0 +
 Node/s: "dn_6005_6006"
(1 row)
```

```
-- Here is an example of a query plan with cost estimates suppressed:
openGauss=# EXPLAIN(COSTS FALSE)SELECT * FROM tpcds.customer_address_p1 WHERE
ca_address_sk=10000;
QUERY PLAN

Data Node Scan
Node/s: dn_6005_6006
(2 rows)

-- Here is an example of a query plan for a query using an aggregate function:
openGauss=# EXPLAIN SELECT SUM(ca_address_sk) FROM tpcds.customer_address_p1 WHERE
ca_address_sk<10000;
QUERY PLAN

Aggregate (cost=18.19..14.32 rows=1 width=4)
-> Streaming (type: GATHER) (cost=18.19..14.32 rows=3 width=4)
Node/s: All dbnodes
-> Aggregate (cost=14.19..14.20 rows=3 width=4)
-> Seq Scan on customer_address_p1 (cost=0.00..14.18 rows=10 width=4)
Filter: (ca_address_sk < 10000)
(6 rows)

-- Create a level-2 partitioned table.
openGauss=# CREATE TABLE range_list
openGauss=# (
openGauss(# month_code VARCHAR2 (30) NOT NULL ,
openGauss(# dept_code VARCHAR2 (30) NOT NULL ,
openGauss(# user_no VARCHAR2 (30) NOT NULL ,
openGauss(# sales_amt int
openGauss(#)
openGauss=# PARTITION BY RANGE (month_code) SUBPARTITION BY LIST (dept_code)
openGauss=# (
openGauss(# PARTITION p_201901 VALUES LESS THAN('201903')
openGauss(# (
openGauss(# SUBPARTITION p_201901_a values ('1'),
openGauss(# SUBPARTITION p_201901_b values ('2')
openGauss(#),
openGauss(# PARTITION p_201902 VALUES LESS THAN('201910')
openGauss(# (
openGauss(# SUBPARTITION p_201902_a values ('1'),
openGauss(# SUBPARTITION p_201902_b values ('2')
openGauss(#)
openGauss(#);
CREATE TABLE

-- Run a query statement containing a level-2 partitioned table.
-- Iterations and Sub Iterations specifies the numbers of level-1 and level-2 partitions that are traversed,
respectively.
-- Selected Partitions specifies which level-1 partitions are actually scanned. Selected Subpartitions: (p:s)
indicates that s level-2 partitions under the pth level-1 partition are actually scanned. If all level-2 partitions
under the level-1 partition are scanned, the value of s is ALL.
openGauss=# EXPLAIN SELECT * FROM range_list WHERE dept_code = '1';
QUERY PLAN

Partition Iterator (cost=0.00..13.81 rows=2 width=238)
Iterations: 2, Sub Iterations: 2
-> Partitioned Seq Scan on range_list (cost=0.00..13.81 rows=2 width=238)
Filter: ((dept_code)::text = '1'::text)
Selected Partitions: 1..2
Selected Subpartitions: 1:1, 2:1
(6 rows)

-- Drop the tpcds.customer_address_p1 table.
openGauss=# DROP TABLE tpcds.customer_address_p1;

-- Drop the tpcds.customer_address table.
openGauss=# DROP TABLE tpcds.customer_address;
```

```
-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[ANALYZE](#) | [ANALYZE](#)

## 7.13.121 EXPLAIN PLAN

### Description

EXPLAIN PLAN saves information about an execution plan into the **PLAN\_TABLE** table. Different from the EXPLAIN statement, EXPLAIN PLAN only saves plan information and does not print information on the screen.

### Syntax

```
EXPLAIN PLAN
[SET STATEMENT_ID = name]
FOR statement ;
```

### Parameters

- **name**  
Specifies a query tag.  
Value range: a string.
- **PLAN**: saves plan information into **PLAN\_TABLE**. If information is stored successfully, "EXPLAIN SUCCESS" is returned.
- **STATEMENT\_ID**: tags each query. The tag information will be stored in **PLAN\_TABLE**.

#### NOTE

If the EXPLAIN PLAN statement does not contain **SET STATEMENT\_ID**, **STATEMENT\_ID** is empty by default. In addition, the value of **STATEMENT\_ID** cannot exceed 30 bytes. Otherwise, an error will be reported.

### Precautions

- EXPLAIN PLAN cannot be executed on a database node.
- Plan information cannot be collected for SQL statements that failed to be executed.
- Data in **PLAN\_TABLE** is in a session-level lifecycle. Sessions are isolated from users, and therefore users can only view the data of the current session and current user.

### Example 1

You can perform the following steps to collect execution plans of SQL statements by running **EXPLAIN PLAN**:

**Step 1** Run the **EXPLAIN PLAN** statement.



 **NOTE**

After the **EXPLAIN PLAN** statement is executed, plan information is automatically stored in **PLAN\_TABLE**. INSERT, UPDATE, and ANALYZE cannot be performed on **PLAN\_TABLE**.

For details about **PLAN\_TABLE**, see [PLAN\\_TABLE](#).

```
-- Create tables foo1 and foo2.
openGauss=# CREATE TABLE foo1(f1 int, f2 text, f3 text[]);
openGauss=# CREATE TABLE foo2(f1 int, f2 text, f3 text[]);

-- Run EXPLAIN PLAN.
openGauss=# EXPLAIN PLAN SET STATEMENT_ID = 'TPCH-Q4' FOR SELECT f1, count(*) FROM foo1 WHERE
f1 > 1 AND f1 < 3 AND EXISTS (SELECT * FROM foo2) GROUP BY f1;
```

**Step 2** Query **PLAN\_TABLE**.

```
openGauss=# SELECT * FROM plan_table;
```

**Step 3** Delete data from **PLAN\_TABLE**.

```
openGauss=# DELETE FROM plan_table WHERE STATEMENT_ID = 'TPCH-Q4';
openGauss=# DROP TABLE foo1;
openGauss=# DROP TABLE foo2;
```

----End

## 7.13.122 FETCH

### Description

FETCH retrieves rows using a previously created cursor.

A cursor has an associated position, which is used by FETCH. The cursor position can be before the first row of the query result, on any particular row of the result, or after the last row of the result.

- When created, a cursor is positioned before the first row.
- After fetching some rows, the cursor is positioned on the row most recently retrieved.
- If FETCH runs off the end of the available rows then the cursor is left positioned after the last row, or before the first row if fetching backward.
- FETCH ALL or FETCH BACKWARD ALL will always leave the cursor positioned after the last row or before the first row.

### Precautions

- If the cursor is declared with **NO SCROLL**, backward fetches like FETCH BACKWARD are not allowed.
- The forms **NEXT**, **PRIOR**, **FIRST**, **LAST**, **ABSOLUTE**, and **RELATIVE** fetch a single row after moving the cursor appropriately. If there is no such row, an empty result is returned, and the cursor is left positioned before the first row (backward fetch) or after the last row (forward fetch) as appropriate.
- The forms using FORWARD and BACKWARD retrieve the indicated number of rows moving in the forward or backward direction, leaving the cursor positioned on the last-returned row or after (backward fetch)/before (forward fetch) all rows if **count** exceeds the number of rows available.
- **RELATIVE 0**, **FORWARD 0**, and **BACKWARD 0** all request fetching the current row without moving the cursor, that is, re-fetching the most recently fetched

row. This action will succeed unless the cursor is positioned before the first row or after the last row. If the cursor is positioned before the first row or after the last row, no row is returned.

## Syntax

```
FETCH [direction { FROM | IN }] cursor_name;
```

The direction clause specifies optional parameters.

```
NEXT
| PRIOR
| FIRST
| LAST
| ABSOLUTE count
| RELATIVE count
| count
| ALL
| FORWARD
| FORWARD count
| FORWARD ALL
| BACKWARD
| BACKWARD count
| BACKWARD ALL
```

## Parameters

- **direction**

Defines the fetch direction.

Value range:

- **NEXT** (default value)  
Fetches the next row.
- **PRIOR**  
Fetches the prior row.
- **FIRST**  
Fetches the first row of the query (same as **ABSOLUTE 1**).
- **LAST**  
Fetches the last row of the query (same as **ABSOLUTE -1**).
- **ABSOLUTE count**  
Fetches the *count*<sup>th</sup> row of the query.

**ABSOLUTE** fetches are not any faster than navigating to the desired row with a relative move: the underlying implementation must traverse all the intermediate rows anyway.

Value range: a possibly-signed integer.

- If *count* is positive, the *count*<sup>th</sup> row of the query will be fetched.
- If *count* is negative, the **abs(count)**<sup>th</sup> row from the end of the query result will be fetched.
- If *count* is set to **0**, the cursor is positioned before the first row.
- **RELATIVE count**  
Fetches the *count*<sup>th</sup> succeeding row or the *count*<sup>th</sup> prior row if count is negative.

Value range: a possibly-signed integer.

- If *count* is positive, the *count*th succeeding row will be fetched.
- If **count** is negative, the *abs(count)*th prior rows will be fetched.
- If the current row contains no data, **RELATIVE 0** returns null.
- count  
Fetches the next *count* rows (same as **FORWARD count**).
- ALL  
Fetches all remaining rows (same as **FORWARD ALL**).
- FORWARD  
Fetches the next row (same as **NEXT**).
- FORWARD count  
Fetches the *count* succeeding rows or *count* prior rows if *count* is negative.
- FORWARD ALL  
Fetches all remaining rows.
- BACKWARD  
Fetches the prior row (same as **PRIOR**).
- BACKWARD count  
Fetches the prior *count* rows (scanning backwards).  
Value range: a possibly-signed integer.
  - If *count* is positive, the prior *count* rows will be fetched.
  - If *count* is a negative, the succeeding *abs (count)* rows will be fetched.
  - **BACKWARD 0** re-fetches the current row, if any.
- BACKWARD ALL  
Fetches all prior rows (scanning backwards).
- **{ FROM | IN } cursor\_name**  
Specifies the cursor name using the keyword **FROM** or **IN**.  
Value range: an existing cursor name.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer_address table.
openGauss=# CREATE TABLE tpcds.customer_address
(
 ca_address_sk INTEGER NOT NULL,
 ca_address_id CHARACTER(16) NOT NULL,
 ca_street_number INTEGER ,
 ca_street_name CHARACTER (20)
);

-- Insert multiple records into the table.
```

```
openGauss=# INSERT INTO tpceds.customer_address VALUES (1, 'AAAAAAAAABAAAAAAA', '18', 'Jackson'),(2,
'AAAAAAAAACAAAAAAA', '362', 'Washington 6th'),(3, 'AAAAAAAADAAAAAAA', '585', 'Dogwood Washington');

-- (For the SELECT statement, traverse a table using a cursor.) Start a transaction.
openGauss=# START TRANSACTION;

-- Set up cursor1.
openGauss=# CURSOR cursor1 FOR SELECT * FROM tpceds.customer_address ORDER BY 1;

-- Fetch the first three rows in cursor1.
openGauss=# FETCH FORWARD 3 FROM cursor1;
ca_address_sk | ca_address_id | ca_street_number | ca_street_name
-----+-----+-----+-----
1 | AAAAAAAAAABAAAAAAA | 18 | Jackson
2 | AAAAAAAAAACAAAAAAA | 362 | Washington 6th
3 | AAAAAAADAAAAAAA | 585 | Dogwood Washington
(3 rows)

-- Close the cursor and commit the transaction.
openGauss=# CLOSE cursor1;

-- End the transaction.
openGauss=# END;

-- (For the VALUES clause, traverse the clause using a cursor.) Start a transaction.
openGauss=# START TRANSACTION;

-- Set up cursor2.
openGauss=# CURSOR cursor2 FOR VALUES(1,2),(0,3) ORDER BY 1;

-- Fetch the first two rows in cursor2.
openGauss=# FETCH FORWARD 2 FROM cursor2;
column1 | column2
-----+-----
0 | 3
1 | 2
(2 rows)

-- Close the cursor and commit the transaction.
openGauss=# CLOSE cursor2;

-- End the transaction.
openGauss=# END;

-- (WITH HOLD cursor) Start a transaction.
openGauss=# START TRANSACTION;

-- Create a cursor using WITH HOLD.
openGauss=# DECLARE cursor1 CURSOR WITH HOLD FOR SELECT * FROM tpceds.customer_address ORDER
BY 1;

-- Fetch the first two rows in cursor1.
openGauss=# FETCH FORWARD 2 FROM cursor1;
ca_address_sk | ca_address_id | ca_street_number | ca_street_name
-----+-----+-----+-----
1 | AAAAAAAAAABAAAAAAA | 18 | Jackson
2 | AAAAAAAAAACAAAAAAA | 362 | Washington 6th
(2 rows)

-- End the transaction.
openGauss=# END;

-- Fetch the next row in cursor1.
openGauss=# FETCH FORWARD 1 FROM cursor1;
ca_address_sk | ca_address_id | ca_street_number | ca_street_name
-----+-----+-----+-----
3 | AAAAAAADAAAAAAA | 585 | Dogwood Washington
(1 row)
```

```
-- Close the cursor.
openGauss=# CLOSE cursor1;

-- Drop the tpcds.customer_address table.
openGauss=# DROP TABLE tpcds.customer_address;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[CLOSE](#) and [MOVE](#)

## 7.13.123 GRANT

### Description

Grants permissions to roles and users.

GRANT is used in the following scenarios:

- **Granting system permissions to roles or users**

System permissions are also called user attributes, including SYSADMIN, CREATEDB, CREATEROLE, AUDITADMIN, MONADMIN, OPRADMIN, POLADMIN, INHERIT, REPLICATION, VCADMIN, and LOGIN.

They can be specified only by the CREATE ROLE or ALTER ROLE statement. The SYSADMIN permissions can be granted and revoked using GRANT ALL PRIVILEGE and REVOKE ALL PRIVILEGE, respectively. System permissions cannot be inherited by a user from a role, and cannot be granted using PUBLIC.

- **Granting database object permissions to roles or users**

Grant permissions on a database object (table, view, column, database, function, schema, or tablespace) to a role or user.

GRANT gives specific permissions on a database object to one or more roles. These permissions are added to those already granted, if any.

The keyword PUBLIC indicates that the permissions are to be granted to all roles, including those that might be created later. **PUBLIC** can be thought of as an implicitly defined group that always includes all roles. Any particular role will have the sum of permissions granted directly to it, permissions granted to any role it is presently a member of, and permissions granted to **PUBLIC**.

If **WITH GRANT OPTION** is specified, the recipient of the permission can in turn grant it to others. Without a grant option, the recipient cannot do that. This option cannot be granted to **PUBLIC**, which is a unique GaussDB attribute.

GaussDB grants the permissions on objects of certain types to PUBLIC users. By default, permissions on tables, columns, sequences, foreign data sources, foreign servers, schemas, and tablespaces are not granted to PUBLIC users, but the following permissions are granted to PUBLIC users: CONNECT and CREATE TEMP TABLE permissions on databases, EXECUTE permission on functions, and USAGE permission on languages and data types (including domains). An object owner can revoke the default permissions granted to PUBLIC users and grant permissions to other users as needed. For security

purposes, you are advised to create an object and set its permissions in the same transaction so that other users do not have time windows to use the object. In addition, you can restrict the permissions of the PUBLIC user group by referring to "Permission Management" in *Security Hardening Guide*. These default permissions can be modified using the **ALTER DEFAULT PRIVILEGES** command.

By default, an object owner has all permissions on the object. For security purposes, the owner can discard some permissions. However, the ALTER, DROP, COMMENT, INDEX, VACUUM, and re-grantable permissions of the object are inherent permissions implicitly owned by the owner.

- **Granting the permissions of one role or user to others**

Grant the permissions of one role or user to others. In this case, every role or user can be regarded as a set of one or more database permissions.

If **WITH ADMIN OPTION** is specified, the recipients can in turn grant the permissions to other roles or users or revoke the permissions they have granted to other roles or users. If recipients' permissions are changed or revoked later, the grantees' permissions will also change.

Database system administrators can grant or revoke permissions to or from any roles or users. Roles with the CREATEROLE permission can grant or revoke permissions to or from non-SYSADMIN roles.

- **Granting ANY permissions to roles or users**

Grant ANY permissions to a specified role or user. For details about the value range of the ANY permissions, see the syntax. If **WITH ADMIN OPTION** is specified, the grantee can grant the ANY permissions to or revoke them from other roles or users. The ANY permissions can be inherited by a role but cannot be granted to PUBLIC users. When separation of duties is disabled, the initial user and system administrators can grant the ANY permissions to or revoke them from any role or user.

Currently, the following ANY permissions are supported: CREATE ANY TABLE, ALTER ANY TABLE, DROP ANY TABLE, SELECT ANY TABLE, INSERT ANY TABLE, UPDATE ANY TABLE, DELETE ANY TABLE, CREATE ANY SEQUENCE, CREATE ANY INDEX, CREATE ANY FUNCTION, EXECUTE ANY FUNCTION, CREATE ANY PACKAGE, EXECUTE ANY PACKAGE, and CREATE ANY TYPE. For details about the ANY permission scope, see [Table 7-108](#).

## Precautions

- It is not allowed to grant the ANY permissions to PUBLIC users or revoke the ANY permissions from PUBLIC users.
- The ANY permissions are database permissions and are valid only for database objects that are granted with the permissions. For example, SELECT ANY TABLE only allows a user to view all user table data in the current database, but the user does not have the permission to view user tables in other databases.
- The ANY permissions and the original permissions do not affect each other.
- If a user is granted with the CREATE ANY TABLE permission, the owner of a table created in a schema with the same name as the user is the creator of the schema. When the user performs other operations on the table, the user needs to be granted with the corresponding operation permission.

- Exercise caution when granting the CREATE ANY FUNCTION or CREATE ANY PACKAGE permission to users to prevent other users from using DEFINER functions or PACKAGE for privilege escalation.

## Syntax

- Grant the table or view access permission to a user or role.  

```
GRANT { { SELECT | INSERT | UPDATE | DELETE | TRUNCATE | REFERENCES | ALTER | DROP | COMMENT | INDEX | VACUUM } [, ...]
 | ALL [PRIVILEGES] }
ON { [TABLE] table_name [, ...]
 | ALL TABLES IN SCHEMA schema_name [, ...] }
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the column access permission to a user or role.  

```
GRANT { { { SELECT | INSERT | UPDATE | REFERENCES | COMMENT } (column_name [, ...]) } [, ...]
 | ALL [PRIVILEGES] (column_name [, ...]) }
ON [TABLE] table_name [, ...]
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the sequence access permission to a specified role or user. The **LARGE** column is optional. The assignment statement does not distinguish whether the sequence is **LARGE**.  

```
GRANT { { SELECT | UPDATE | USAGE | ALTER | DROP | COMMENT } [, ...]
 | ALL [PRIVILEGES] }
ON { [[LARGE] SEQUENCE] sequence_name [, ...]
 | ALL SEQUENCES IN SCHEMA schema_name [, ...] }
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the database access permission to a user or role.  

```
GRANT { { CREATE | CONNECT | TEMPORARY | TEMP | ALTER | DROP | COMMENT } [, ...]
 | ALL [PRIVILEGES] }
ON DATABASE database_name [, ...]
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the domain access permission to a user or role.  

```
GRANT { USAGE | ALL [PRIVILEGES] }
ON DOMAIN domain_name [, ...]
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
-  **NOTE**
- In the current version, the domain access permission cannot be granted.
- Grant the foreign data source access permission to a user or role.  

```
GRANT { USAGE | ALL [PRIVILEGES] }
ON FOREIGN DATA WRAPPER fdw_name [, ...]
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the foreign server access permission to a user or role.  

```
GRANT { { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON FOREIGN SERVER server_name [, ...]
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the function access permission to a user or role.  

```
GRANT { { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON { FUNCTION {function_name ([{[argmode] [arg_name] arg_type} [, ...]) } } [, ...]
 | ALL FUNCTIONS IN SCHEMA schema_name [, ...] }
TO { [GROUP] role_name | PUBLIC } [, ...]
[WITH GRANT OPTION];
```
  - Grant the procedural procedure access permission to a user or role.

```
GRANT { { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON { PROCEDURE {proc_name ([{ argmode } [arg_name] arg_type } [, ...]) } } [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

- Grant the procedural language access permission to a user or role.

```
GRANT { USAGE | ALL [PRIVILEGES] }
 ON LANGUAGE lang_name [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

- Grant the large object access permission to a specified user or role.

```
GRANT { { SELECT | UPDATE } [, ...] | ALL [PRIVILEGES] }
 ON LARGE OBJECT loid [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

#### NOTE

In the current version, the large object access permission cannot be granted.

- Grant the schema access permission to a user or role.

```
GRANT { { CREATE | USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON SCHEMA schema_name [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

#### NOTE

When you grant table or view permissions to other users, you also need to grant the USAGE permission on the schema that the tables and views belong to. Without the USAGE permission, the users with table or view permissions can only see the object names, but cannot access them. This syntax cannot be used to grant the permission to create tables in schemas with the same name, but you can use the syntax for granting permission of a role to another user or role to achieve the same effect.

- Grant the tablespace access permission to a user or role.

```
GRANT { { CREATE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON TABLESPACE tablespace_name [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

- Grant the type access permission to a user or role.

```
GRANT { { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON TYPE type_name [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

#### NOTE

In the current version, the type access permission cannot be granted.

- Grant the directory permission to a role.

```
GRANT { { READ | WRITE | ALTER | DROP } [, ...] | ALL [PRIVILEGES] }
 ON DIRECTORY directory_name [, ...]
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

- Grant the package permission to a role.

```
GRANT { { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON { PACKAGE package_name [, ...]
 | ALL PACKAGES IN SCHEMA schema_name [, ...] }
 TO { [GROUP] role_name | PUBLIC } [, ...]
 [WITH GRANT OPTION];
```

- Grant a role's permissions to another user or role.

```
GRANT role_name [, ...]
 TO role_name [, ...]
 [WITH ADMIN OPTION];
```

- Grant the SYSADMIN permission to a role.



```
GRANT ALL { PRIVILEGES | PRIVILEGE }
TO role_name;
```

- Grant the ANY permissions to another user or role.

```
GRANT { CREATE ANY TABLE | ALTER ANY TABLE | DROP ANY TABLE | SELECT ANY TABLE | INSERT
ANY TABLE | UPDATE ANY TABLE |
DELETE ANY TABLE | CREATE ANY SEQUENCE | CREATE ANY INDEX | CREATE ANY FUNCTION |
EXECUTE ANY FUNCTION |
CREATE ANY PACKAGE | EXECUTE ANY PACKAGE | CREATE ANY TYPE } [, ...]
TO [GROUP] role_name [, ...]
[WITH ADMIN OPTION];
```

## Parameters

GRANT permissions are classified as follows:

- **SELECT**  
Allows SELECT from any column, or the specific columns listed, of the specified table, view, or sequence. The SELECT permission on the corresponding field is also required for UPDATE or DELETE.
- **INSERT**  
Allows INSERT of a new row into a table.
- **UPDATE**  
Allows UPDATE of any column of a table. Generally, UPDATE also requires the SELECT permission to query which rows need to be updated. SELECT ... FOR UPDATE and SELECT ... FOR SHARE also require this permission on at least one column, in addition to the SELECT permission.
- **DELETE**  
Allows DELETE of a row from a table. Generally, DELETE also requires the SELECT permission to query which rows need to be deleted.
- **TRUNCATE**  
Allows TRUNCATE on a table.
- **REFERENCES**  
Allows creation of a foreign key constraint referencing a table. This permission is required on both referencing and referenced tables.
- **CREATE**
  - For databases, allows new schemas to be created within the database.
  - For schemas, allows new objects to be created within the schema. To rename an existing object, you must own the object and have the CREATE permission on the schema of the object.
  - For tablespaces, allows tables to be created within the tablespace, and allows databases and schemas to be created that have the tablespace as their default tablespace.
- **CONNECT**  
Allows the grantee to connect to the database.
- **EXECUTE**  
Allows calling a function, including use of any operators that are implemented on top of the function.
- **USAGE**

- For procedural languages, allows use of the language for the creation of functions in that language.
- For schemas, allows access to objects contained in the schema. Without this permission, it is still possible to see the object names.
- For sequences, allows use of the nextval function.
- **ALTER**  
Allows users to modify the attributes of a specified object, excluding the owner and schema of the object.
- **DROP**  
Allows users to drop specified objects.
- **COMMENT**  
Allows users to define or modify comments of a specified object.
- **INDEX**  
Allows users to create indexes on specified tables, manage indexes on the tables, and perform REINDEX and CLUSTER operations on the tables.
- **VACUUM**  
Allows users to perform ANALYZE and VACUUM operations on specified tables.
- **ALL PRIVILEGES**  
Grants all available permissions to a user or role at a time. Only a system administrator has the GRANT ALL PRIVILEGES permission.

GRANT parameters are as follows:

- **role\_name**  
Specifies the username.
- **table\_name**  
Specifies the table name.
- **column\_name**  
Specifies the column name.
- **schema\_name**  
Specifies the schema name.
- **database\_name**  
Specifies the database name.
- **function\_name**  
Specifies the function name.
- **procedure\_name**  
Specifies the stored procedure name.
- **sequence\_name**  
Specifies the sequence name.
- **domain\_name**  
Specifies the domain type name.
- **fdw\_name**  
Specifies the foreign data wrapper name.

- **lang\_name**  
Specifies the language name.
- **type\_name**  
Specifies the type name.
- **argmode**  
Specifies the parameter mode.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **arg\_name**  
Specifies the parameter name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **arg\_type**  
Specifies the parameter type.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **loid**  
Specifies the identifier of the large object that includes this page.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **tablespace\_name**  
Specifies the tablespace name.
- **directory\_name**  
Specifies a directory name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **WITH GRANT OPTION**  
If **WITH GRANT OPTION** is specified, the recipient of the permission can in turn grant it to others. Without a grant option, the recipient cannot do that. Grant options cannot be granted to **PUBLIC**.

When a non-owner of an object attempts to grant permissions on the object:

- The statement will fail outright if the user has no permissions whatsoever on the object.
- As long as some permission is available, the statement will proceed, but it will grant only those permissions for which the user has grant options.
- The GRANT ALL PRIVILEGES forms will issue a warning message if no grant options are held, while the other forms will issue a warning if grant options for any of the permissions specifically named in the statement are not held.

#### NOTE

Database system administrators can access all objects regardless of object permission settings. This is comparable to the permissions of **root** in a Unix system. As with **root**, it is unwise to operate as a system administrator except when necessary.

- **WITH ADMIN OPTION**  
If **WITH ADMIN OPTION** is specified for a role, the grantee can grant the role to other roles or users or revoke the role from other roles or users.  
For the ANY permissions, if **WITH ADMIN OPTION** is specified, the grantee can grant the ANY permissions to or revoke them from other roles or users.

**Table 7-108** ANY permissions

| System Permission    | Description                                                                                                                                                                                                                                                              |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CREATE ANY TABLE     | Users can create tables or views in the public and user schemas. The users must be granted with the permission to create sequences to create a table that contains serial columns.                                                                                       |
| ALTER ANY TABLE      | Users' ALTER permission on tables or views in the public and user schemas. If the users want to modify the unique index of a table to add a primary key constraint or unique constraint to the table, the users must be granted with the index permission for the table. |
| DROP ANY TABLE       | Users' DROP permission on tables or views in the public and user schemas.                                                                                                                                                                                                |
| SELECT ANY TABLE     | Users' SELECT permission on tables or views in the public and user schemas, which is still subject to row-level security.                                                                                                                                                |
| UPDATE ANY TABLE     | Users' UPDATE permission on tables or views in the public and user schemas, which is still subject to row-level security.                                                                                                                                                |
| INSERT ANY TABLE     | Users' INSERT permission on tables or views in the public and user schemas.                                                                                                                                                                                              |
| DELETE ANY TABLE     | Users' DELETE permission on tables or views in the public and user schemas, which is still subject to row-level security.                                                                                                                                                |
| CREATE ANY FUNCTION  | Users can create functions or stored procedures in the user schemas.                                                                                                                                                                                                     |
| EXECUTE ANY FUNCTION | Users' EXECUTE permission on functions or stored procedures in the public and user schemas.                                                                                                                                                                              |
| CREATE ANY PACKAGE   | Users can create packages in the public and user schemas.                                                                                                                                                                                                                |
| EXECUTE ANY PACKAGE  | Users' EXECUTE permission on packages in the public and user schemas.                                                                                                                                                                                                    |
| CREATE ANY TYPE      | Users can create types in the public and user schemas.                                                                                                                                                                                                                   |
| CREATE ANY SEQUENCE  | Users can create sequences in the public and user schemas.                                                                                                                                                                                                               |
| CREATE ANY INDEX     | Users can create indexes in the public and user schemas. The users must be granted with the permission to create tablespaces to create a partitioned table index in a tablespace.                                                                                        |

 NOTE

If a user is granted with any ANY permission, the user has the USAGE permission on the public and user schemas but does not have the USAGE permission on the system schemas except **public** listed in [Table 13-1](#).

## Examples

### Example: Granting system permissions to a user or role

Create the **joe** user and grant the SYSADMIN permissions to it.

```
openGauss=# CREATE USER joe PASSWORD '*****';
openGauss=# GRANT ALL PRIVILEGES TO joe;
```

Then **joe** has the SYSADMIN permission.

### Example: Granting object permissions to a user or role

1. Revoke the SYSADMIN permission from the **joe** user. Grant the usage permission of the **tpcds** schema and all permissions on the **tpcds.reason** table to **joe**.

```
openGauss=# CREATE SCHEMA tpcds;

openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk INTEGER NOT NULL,
 r_reason_id CHAR(16) NOT NULL,
 r_reason_desc VARCHAR(20)
);

openGauss=# REVOKE ALL PRIVILEGES FROM joe;
openGauss=# GRANT USAGE ON SCHEMA tpcds TO joe;
openGauss=# GRANT ALL PRIVILEGES ON tpcds.reason TO joe;
```

Then **joe** has all permissions on the **tpcds.reason** table, including CREATE, RETRIEVE, UPDATE, and DELETE.

2. Grant the retrieve permission of **r\_reason\_sk**, **r\_reason\_id**, and **r\_reason\_desc** columns and the UPDATE permission of the **r\_reason\_desc** column in the **tpcds.reason** table to **joe**.

```
openGauss=# GRANT select (r_reason_sk,r_reason_id,r_reason_desc),update (r_reason_desc) ON
tpcds.reason TO joe;
```

Then **joe** has the retrieve permission on the **r\_reason\_sk**, **r\_reason\_id**, and **r\_reason\_desc** columns in the **tpcds.reason** table. To enable **joe** to grant these permissions to other users, execute the following statement:

```
openGauss=# GRANT select (r_reason_sk, r_reason_id) ON tpcds.reason TO joe WITH GRANT OPTION;
```

Grant the database connection permission and the permission to create schemas in GaussDB to user **joe**, and allow user **joe** to grant this permission to other users.

```
openGauss=# CREATE DATABASE openGauss;
openGauss=# GRANT create,connect on database openGauss TO joe WITH GRANT OPTION;
```

Create the **tpcds\_manager** role, grant the access and object creation permissions of the **tpcds** schema to **tpcds\_manager**, but do not allow **tpcds\_manager** to grant these permissions to others.

```
openGauss=# CREATE ROLE tpcds_manager PASSWORD '*****';
openGauss=# GRANT USAGE,CREATE ON SCHEMA tpcds TO tpcds_manager;
```

Grant all permissions on the **tpcds\_tbsp** tablespace to **joe**, but do not allow **joe** to grant these permissions to others.

```
openGauss=# CREATE TABLESPACE tpcds_tbspc RELATIVE LOCATION 'tablespace/tablespace_1';
openGauss=# GRANT ALL ON TABLESPACE tpcds_tbspc TO joe;
```

### Example: Granting the permissions of one user or role to others

1. Create the **manager** role, grant **joe**'s permissions to **manager**, and allow **manager** to grant these permissions to others.

```
openGauss=# CREATE ROLE manager PASSWORD '*****';
openGauss=# GRANT joe TO manager WITH ADMIN OPTION;
```

2. Create the **senior\_manager** user and grant **manager**'s permissions to it.

```
openGauss=# CREATE ROLE senior_manager PASSWORD '*****';
openGauss=# GRANT manager TO senior_manager;
```

3. Revoke permissions and drop users.

```
openGauss=# REVOKE joe FROM manager;
openGauss=# REVOKE manager FROM senior_manager;
openGauss=# DROP USER manager;
```

### Example: Revoking permissions and dropping roles and users

```
openGauss=# REVOKE ALL PRIVILEGES ON tpcds.reason FROM joe;
openGauss=# REVOKE ALL PRIVILEGES ON SCHEMA tpcds FROM joe;
openGauss=# REVOKE ALL ON TABLESPACE tpcds_tbspc FROM joe;
openGauss=# DROP TABLESPACE tpcds_tbspc;
openGauss=# REVOKE USAGE,CREATE ON SCHEMA tpcds FROM tpcds_manager;
openGauss=# DROP ROLE tpcds_manager;
openGauss=# DROP ROLE senior_manager;
openGauss=# DROP USER joe CASCADE;
openGauss=# DROP TABLE tpcds.reason;
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[REVOKE](#) and [ALTER DEFAULT PRIVILEGES](#)

## 7.13.124 INSERT

### Description

Inserts new rows into a table.

### Precautions

- You must have the INSERT permission on a table to insert data into it. If a user is granted the INSERT ANY TABLE permission, the user has the USAGE permission on all schemas except system schemas and the INSERT permission on tables in these schemas.
- Use of the RETURNING clause requires the SELECT permission on all columns mentioned in RETURNING.
- If ON DUPLICATE KEY UPDATE is used, you must have the INSERT and UPDATE permissions on the table and the SELECT permission on the columns of the UPDATE clause.
- If you use the query clause to insert rows from a query, you need to have the SELECT permission on any table or column used in the query.
- The generated column cannot be directly written. In the INSERT statement, values cannot be specified for generated columns, but the keyword DEFAULT can be specified.

- When you connect to a database compatible to Teradata and **td\_compatible\_truncation** is **on**, a long string will be automatically truncated. If later INSERT statements insert long strings into columns of char- and varchar-typed columns in the target table (not involving foreign tables), the system will truncate the long strings at the maximum length defined in the target table.

#### NOTE

If inserting multi-byte character data (such as Chinese characters) into a database with the character set byte encoding (SQL\_ASCII, LATIN1), and the character data crosses the truncation position, the string is truncated based on its bytes instead of characters. Unexpected result will occur in tail after the truncation. If you want correct truncation result, you are advised to adopt encoding set such as UTF8, which has no character data crossing the truncation position.

## Syntax

```
[WITH [RECURSIVE] with_query [, ...]]
INSERT [/*+ plan_hint */] INTO table_name [partition_clause] [AS alias] [(column_name [, ...])]
 { DEFAULT VALUES
 | VALUES (({ expression | DEFAULT } [, ...]))[, ...]
 | query
 | ON DUPLICATE KEY UPDATE { NOTHING | { column_name = { expression | DEFAULT } } [, ...] [WHERE
condition] } }
 [RETURNING { * | {output_expression [[AS] output_name] }[, ...] }];
```

## Parameters

- WITH [ RECURSIVE ] with\_query [, ...]**

Specifies one or more subqueries that can be referenced by name in the main query, which is equivalent to a temporary table.

If **RECURSIVE** is specified, it allows a SELECT subquery to reference itself by name.

Format of **with\_query**:

```
with_query_name [(column_name [, ...])] AS [[NOT] MATERIALIZED]
({SELECT | VALUES | INSERT | UPDATE | DELETE})
```

– **with\_query\_name** specifies the name of the result set generated by a subquery. Such names can be used to access the result sets of subqueries in a query.

– **column\_name** specifies the column name displayed in the subquery result set.

– Each subquery can be a SELECT, VALUES, INSERT, UPDATE, or DELETE statement.

– You can use **MATERIALIZED** or **NOT MATERIALIZED** to modify the CTE.

– If **MATERIALIZED** is specified, the WITH query will be materialized, and a copy of the subquery result set is generated. The copy is directly queried at the reference point. Therefore, the WITH subquery cannot be jointly optimized with the SELECT statement trunk (for example, predicate pushdown and equivalence class transfer). In this scenario, you can use **NOT MATERIALIZED** for modification. If the WITH query can be executed as a subquery inline, the preceding optimization can be performed.

– If the user does not explicitly declare the materialized attribute, comply with the following rules: If the CTE is referenced only once in the trunk statement to which it belongs and semantically supports inline execution,

it will be rewritten as subquery inline execution. Otherwise, the materialized execution will be performed in CTE Scan mode.

#### NOTE

INSERT ON DUPLICATE KEY UPDATE does not support the WITH or WITH RECURSIVE clauses.

- **plan\_hint** clause

Follows the INSERT keyword in the */\*+ \*/* format. It is used to optimize the plan of an INSERT statement block. For details, see [Hint-based Tuning](#). In each statement, only the first */\*+ plan\_hint \*/* comment block takes effect as a hint. Multiple hints can be written.

- **table\_name**

Specifies the name of the target table where data will be inserted.

Value range: an existing table name.

- **partition\_clause**

Inserts data into a specified partition.

```
PARTITION { (partition_name) | FOR (partition_value [, ...]) } |
SUBPARTITION { (subpartition_name) | FOR (subpartition_value [, ...]) }
```

For details about the keywords, see [SELECT](#).

If the value of the VALUE clause is inconsistent with that of the specified partition, an exception is displayed.

For details, see [CREATE TABLE SUBPARTITION](#).

- **column\_name**

Specifies the name of a column in a table.

- The column name can be qualified with a subcolumn name or array index, if needed.
- Each column not present in the explicit or implicit column list will be filled with a default value, either its declared default value or **NULL** if there is none. Inserting data into only some columns of a composite type leaves the other columns **NULL**.
- The target column names **column\_name** can be listed in any order. If no list of column names is given at all, the default is all the columns of the table in their declared order.
- The target columns are the first *N* column names, if there are only *N* columns provided by the VALUE clause and QUERY.
- The values provided by the VALUE clause and QUERY are joined with the corresponding columns from left to right in the table.

Value range: name of an existing column.

- **expression**

Specifies an expression or a value to assign to the corresponding column.

- In the INSERT ON DUPLICATE KEY UPDATE statement, **expression** can be set to **VALUES(column\_name)** or **EXCLUDED.column\_name**, indicating that the value of **column\_name** corresponding to the conflict row is referenced. Note that **VALUES(column\_name)** cannot be nested in an expression (for example, **VALUES(column\_name)+1**). **EXCLUDED** is not subject to this restriction.



- If single quotation marks are inserted in a column, the single quotation marks need to be used for escape.
- If the expression for any column is not of the correct data type, automatic type conversion will be attempted. If the attempt fails, data insertion fails, and the system returns an error message.
- **DEFAULT**  
Specifies the default value of a column. The value is **NULL** if no default value has been assigned to it.
- **query**  
Specifies a query statement (SELECT statement) that uses the query result as the inserted data.
- **RETURNING**  
Returns the inserted rows. The syntax of the RETURNING list is identical to that of the output list of SELECT. Note that INSERT ON DUPLICATE KEY UPDATE does not support the RETURNING clause.
- **output\_expression**  
Specifies an expression used to calculate the output result of the INSERT statement after each row is inserted.  
Value range: The expression can use any column in the table. You can use the asterisk (\*) to return all columns of the inserted row.
- **output\_name**  
Specifies a name to use for a returned column.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **ON DUPLICATE KEY UPDATE**  
For a table with a unique constraint (UNIQUE INDEX or PRIMARY KEY), if the inserted data violates the unique constraint, the UPDATE clause is executed on the conflicting row to complete the update. For a table without a unique constraint, only the INSERT operation is performed. If the clause of UPDATE is NOTHING, this insertion is ignored. You can use "**EXCLUDE.**" or "**VALUES()**" to select the column corresponding to the source data.
  - Triggers are supported. The execution sequence of triggers is determined by the actual execution process.
    - Executing INSERT will trigger the BEFORE INSERT and AFTER INSERT triggers.
    - Executing **UPDATE** will trigger the BEFORE INSERT, BEFORE UPDATE, and AFTER UPDATE.
    - Executing **UPDATE NOTHING** will trigger the BEFORE INSERT trigger.
  - The unique constraint or primary key of DEFERRABLE is not supported.
  - If a table has multiple unique constraints and the inserted data violates multiple unique constraints, only the first row that has a conflict is updated. (The check sequence is closely related to index maintenance. Generally, the conflict check is performed on the index that is created first.)

- If multiple rows are inserted and these rows conflict with the same row in the table, the system inserts or updates the first row and then updates other rows in sequence.
- Primary keys and unique index columns cannot be updated.
- MOTs are not supported.
- Subquery expressions are supported. The syntax and function of the expressions are the same as those of UPDATE. In a subquery expression, "**EXCLUDED.**" can be used to select the columns corresponding to the source data.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert a record into a table.
openGauss=# INSERT INTO tpcds.reason(r_reason_sk, r_reason_id, r_reason_desc) VALUES (0,
'AAAAAAAAAAAAAAAA', 'reason0');

-- Create the tpcds.reason_t2 table.
openGauss=# CREATE TABLE tpcds.reason_t2
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert a record into a table.
openGauss=# INSERT INTO tpcds.reason_t2(r_reason_sk, r_reason_id, r_reason_desc) VALUES (1,
'AAAAAAAABAAAAAAA', 'reason1');

-- Insert a record into the table, which is equivalent to the previous syntax.
openGauss=# INSERT INTO tpcds.reason_t2 VALUES (2, 'AAAAAAAABAAAAAAA', 'reason2');

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason_t2 VALUES (3, 'AAAAAAAACAAAAAAA', 'reason3'),(4,
'AAAAAAAADAAAAAAA', 'reason4'),(5, 'AAAAAAAAEAAAAAAA', 'reason5');

-- Insert records whose r_reason_sk in the tpcds.reason table is less than 5.
openGauss=# INSERT INTO tpcds.reason_t2 SELECT * FROM tpcds.reason WHERE r_reason_sk <5;

-- Create a unique index for the table.
openGauss=# CREATE UNIQUE INDEX reason_t2_u_index ON tpcds.reason_t2(r_reason_sk);

-- Insert multiple records into the table. If the records conflict, update the r_reason_id column in the
conflict data row to BBBBBBBBCAAAAAAA.
openGauss=# INSERT INTO tpcds.reason_t2 VALUES (5, 'BBBBBBBBCAAAAAAA', 'reason5'),(6,
'AAAAAAAADAAAAAAA', 'reason6') ON DUPLICATE KEY UPDATE r_reason_id = 'BBBBBBBBCAAAAAAA';

-- Drop the tpcds.reason_t2.
openGauss=# DROP TABLE tpcds.reason_t2;

-- Drop the tpcds.reason table.
openGauss=# DROP TABLE tpcds.reason;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Suggestions

### VALUES

When you use the INSERT statement to insert data in batches, you are advised to combine multiple records into one statement to improve data loading performance. Example: **INSERT INTO sections VALUES (30, 'Administration', 31, 1900),(40, 'Development', 35, 2000), (50, 'Development' , 60 , 2001);**

## 7.13.125 LOCK

### Description

LOCK TABLE obtains a table-level lock.

GaussDB always tries to select the lock mode with minimum constraints when automatically requesting a lock for a statement referenced by a table. Use LOCK if users need a more strict lock mode. For example, suppose an application runs a transaction at the Read Committed isolation level and needs to ensure that data in a table remains stable in the duration of the transaction. To achieve this, you could obtain SHARE lock mode over the table before the query. This will prevent concurrent data changes and ensure subsequent reads of the table see a stable view of committed data. It is because the SHARE lock mode conflicts with the ROW EXCLUSIVE lock acquired by writers, and your LOCK TABLE name IN SHARE MODE statement will wait until any concurrent holders of ROW EXCLUSIVE mode locks commit or roll back. Therefore, once you obtain the lock, there are no uncommitted writes outstanding; furthermore none can begin until you release the lock.

### Precautions

- LOCK TABLE is useless outside a transaction block, because the lock would remain held only to the completion of the statement. If LOCK TABLE is out of any transaction block, an error is reported.
- If no lock mode is specified, then **ACCESS EXCLUSIVE**, the most restrictive mode, is used.
- LOCK TABLE ... IN ACCESS SHARE MODE requires the SELECT permission on the target table. All other forms of LOCK require table-level UPDATE and/or the DELETE permission.
- There is no **UNLOCK TABLE** command. Locks are always released at transaction end.
- LOCK TABLE only deals with table-level locks, and so the mode names involving **ROW** are all misnomers. These mode names should generally be read as indicating the intention of the user to acquire row-level locks within the locked table. Also, ROW EXCLUSIVE mode is a shareable table lock. Note that all the lock modes have identical semantics so far as LOCK TABLE is concerned, differing only in the rules about which modes conflict with which. For details about the rules, see [Table 7-109](#).
- If the xc\_maintenance\_mode parameter is not enabled, an error is reported when an ACCESS EXCLUSIVE lock is applied for a system catalog.

## Syntax

```
LOCK [TABLE] {[ONLY] name [, ...]} {name [*]} [, ...]
 [IN {ACCESS SHARE | ROW SHARE | ROW EXCLUSIVE | SHARE UPDATE EXCLUSIVE | SHARE | SHARE
ROW EXCLUSIVE | EXCLUSIVE | ACCESS EXCLUSIVE} MODE]
 [NOWAIT];
```

## Parameters

**Table 7-109** Lock mode conflicts

| Requested Lock Mode/<br>Current Lock Mode | ACCESS SHARE | ROW SHARE | ROW EXCLUSIVE | SHARE UPDATE EXCLUSIVE | SHARE | SHARE ROW EXCLUSIVE | EXCLUSIVE | ACCESS EXCLUSIVE |
|-------------------------------------------|--------------|-----------|---------------|------------------------|-------|---------------------|-----------|------------------|
| ACCESS SHARE                              | -            | -         | -             | -                      | -     | -                   | -         | X                |
| ROW SHARE                                 | -            | -         | -             | -                      | -     | -                   | X         | X                |
| ROW EXCLUSIVE                             | -            | -         | -             | -                      | X     | X                   | X         | X                |
| SHARE UPDATE EXCLUSIVE                    | -            | -         | -             | X                      | X     | X                   | X         | X                |
| SHARE                                     | -            | -         | X             | X                      | -     | X                   | X         | X                |
| SHARE ROW EXCLUSIVE                       | -            | -         | X             | X                      | X     | X                   | X         | X                |
| EXCLUSIVE                                 | -            | X         | X             | X                      | X     | X                   | X         | X                |
| ACCESS EXCLUSIVE                          | X            | X         | X             | X                      | X     | X                   | X         | X                |

**LOCK** parameters are as follows:

- **name**  
Specifies the name (optionally schema-qualified) of an existing table to lock. Tables are locked one-by-one in the order specified in the **LOCK TABLE** command.  
Value range: an existing table name.
- **ONLY**  
If **ONLY** is specified, only that table is locked. If it is not specified, the table and all its sub-tables are locked.
- **ACCESS SHARE**  
Conflicts with the ACCESS EXCLUSIVE lock mode only.  
The SELECT statement acquires a lock of this mode on referenced tables. Typically, any command that reads a table without modifying it acquires this lock mode.
- **ROW SHARE**  
It conflicts with the EXCLUSIVE and ACCESS EXCLUSIVE lock modes.  
**SELECT FOR UPDATE** and **SELECT FOR SHARE** automatically acquire the ROW SHARE lock on the target table and add the ACCESS SHARE lock to other referenced tables on which FOR SHARE and FOR UPDATE have not been performed.
- **ROW EXCLUSIVE**  
Allows concurrent read of a table but does not allow modification of data in the table, which is the same as ROW SHARE. UPDATE, DELETE, and INSERT automatically acquire this lock on the target table and add the ACCESS SHARE lock to other referenced tables. Generally, all statements that modify table data acquire the ROW EXCLUSIVE lock for tables.
- **SHARE UPDATE EXCLUSIVE**  
Protects a table against concurrent schema changes and VACUUM runs.  
The **VACUUM** (without **FULL**), **ANALYZE**, and **CREATE INDEX CONCURRENTLY** commands automatically request this lock.
- **SHARE**  
Allows concurrent queries of a table but does not allow modification of the table.  
CREATE INDEX (without **CONCURRENTLY**) automatically requests this lock.
- **SHARE ROW EXCLUSIVE**  
Protects a table against concurrent data changes, and is self-exclusive so that only one session can hold it at a time.  
No SQL statements automatically acquire this lock mode.
- **EXCLUSIVE**  
Allows concurrent queries of the target table but does not allow any other operations.  
This mode allows only concurrent ACCESS SHARE locks; that is, only reads from the table can proceed in parallel with a transaction holding this lock mode.  
No SQL statements automatically acquire this lock mode on user tables. However, it will be acquired on some system catalogs in case of some operations.

- **ACCESS EXCLUSIVE**

Guarantees that the holder is the only transaction accessing the table in any way.

The ALTER TABLE, DROP TABLE, TRUNCATE, and REINDEX statements automatically request this lock.

This is also the default lock mode for **LOCK TABLE** commands that do not specify a mode explicitly.

- **NOWAIT**

Specifies that LOCK TABLE does not wait for any conflicting locks to be released. If the lock cannot be obtained immediately, the command exits and an error message is displayed.

If a table-level lock is obtained without specifying **NOWAIT** and other mutex locks exist, the system waits for other locks to be released.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk INTEGER NOT NULL,
 r_reason_id CHAR(16) NOT NULL,
 r_reason_desc INTEGER
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason VALUES (1, 'AAAAAAAAABAAAAAAA', '18'),(5,
'AAAAAAAACAAAAAAA', '362'),(7, 'AAAAAAAADAAAAAAA', '585');

-- Obtain a SHARE ROW EXCLUSIVE lock on a primary key table when going to perform a DELETE
operation.
openGauss=# CREATE TABLE tpcds.reason_t1 AS TABLE tpcds.reason;

openGauss=# START TRANSACTION;

openGauss=# LOCK TABLE tpcds.reason_t1 IN SHARE ROW EXCLUSIVE MODE;

openGauss=# DELETE FROM tpcds.reason_t1 WHERE r_reason_desc IN(SELECT r_reason_desc FROM
tpcds.reason_t1 WHERE r_reason_sk < 6);

openGauss=# DELETE FROM tpcds.reason_t1 WHERE r_reason_sk = 7;

openGauss=# COMMIT;

-- Drop the tpcds.reason_t1 table.
openGauss=# DROP TABLE tpcds.reason_t1;

-- Drop the table.
openGauss=# DROP TABLE tpcds.reason;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.126 MERGE INTO

### Description

MERGE INTO conditionally matches data in a target table with that in a source table. If data matches, **UPDATE** is executed on the target table; if data does not

match, **INSERT** is executed. You can use this syntax to run **UPDATE** and **INSERT** at a time for convenience

## Precautions

You have the **INSERT** and **UPDATE** permissions on the target table and the **SELECT** permission on the source table.

## Syntax

```
MERGE [/*+ plan_hint */] INTO table_name [partition_clause] [[AS] alias]
USING { { table_name | view_name } | subquery } [[AS] alias]
ON (condition)
[
 WHEN MATCHED THEN
 UPDATE SET { column_name = { expression | subquery | DEFAULT } |
 (column_name [, ...]) = ({ expression | subquery | DEFAULT } [, ...]) } [, ...]
 [WHERE condition]
]
[
 WHEN NOT MATCHED THEN
 INSERT { DEFAULT VALUES |
 [(column_name [, ...])] VALUES ({ expression | subquery | DEFAULT } [, ...]) [, ...] [WHERE condition] }
];
```

## Parameters

- **plan\_hint** clause  
Follows the **MERGE** keyword in the */\*+ \*/* format. It is used to optimize the plan of a **MERGE** statement block. For details, see [Hint-based Tuning](#). In each statement, only the first */\*+ plan\_hint \*/* comment block takes effect as a hint. Multiple hints can be written.
- **INTO** clause  
Specifies the target table that is being updated or has data being inserted.
- **table\_name**  
Specifies the name of the target table.
- **partition\_clause**  
Performs **MERGE** operations on a specified partition.  

```
PARTITION { (partition_name) | FOR (partition_value [, ...]) } |
SUBPARTITION { (subpartition_name) | FOR (subpartition_value [, ...]) }
```

  
For details about the keywords, see [SELECT](#).  
If the value of the **VALUE** clause is inconsistent with that of the specified partition, an exception is displayed.  
For details, see [CREATE TABLE SUBPARTITION](#).
- **alias**  
Specifies the alias of the target table.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **USING** clause  
Specifies the source table, which can be a table, view, or subquery.
- **ON** clause  
Specifies the condition used to match data between the source and target tables. Columns in the condition cannot be updated.

- **WHEN MATCHED** clause  
Performs UPDATE if data in the source table matches that in the target table based on the condition.  
System catalogs and system columns cannot be updated.
- **WHEN NOT MATCHED** clause  
Performs INSERT if data in the source table does not match that in the target table based on the condition.  
An INSERT clause can contain only one **VALUES**.  
The order of WHEN MATCHED and WHEN NOT MATCHED clauses can be reversed. One of them can be used by default, but they cannot be both used at one time. Two **WHEN MATCHED** or **WHEN NOT MATCHED** clauses cannot be specified at the same time.
- **DEFAULT**  
Specifies the default value of a column.  
The value is **NULL** if no default value has been assigned to it.
- **WHERE** condition  
Specifies the conditions for the UPDATE and INSERT clauses. The two clauses will be executed only when the conditions are met. The default value can be used. System columns cannot be referenced in WHERE condition. You are advised not to use numeric types such as int for **condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and **0** is implicitly converted to **false**), which may cause unexpected results.

## Examples

```
-- Create the target table products and source table newproducts, and insert data into them.
openGauss=# CREATE TABLE products
(
 product_id INTEGER,
 product_name VARCHAR2(60),
 category VARCHAR2(60)
);

openGauss=# INSERT INTO products VALUES (1501, 'vivitar 35mm', 'electrncs');
openGauss=# INSERT INTO products VALUES (1502, 'olympus is50', 'electrncs');
openGauss=# INSERT INTO products VALUES (1600, 'play gym', 'toys');
openGauss=# INSERT INTO products VALUES (1601, 'lamaze', 'toys');
openGauss=# INSERT INTO products VALUES (1666, 'harry potter', 'dvd');

openGauss=# CREATE TABLE newproducts
(
 product_id INTEGER,
 product_name VARCHAR2(60),
 category VARCHAR2(60)
);

openGauss=# INSERT INTO newproducts VALUES (1502, 'olympus camera', 'electrncs');
openGauss=# INSERT INTO newproducts VALUES (1601, 'lamaze', 'toys');
openGauss=# INSERT INTO newproducts VALUES (1666, 'harry potter', 'toys');
openGauss=# INSERT INTO newproducts VALUES (1700, 'wait interface', 'books');

-- Run MERGE INTO.
openGauss=# MERGE INTO products p
 USING newproducts np
 ON (p.product_id = np.product_id)
 WHEN MATCHED THEN
 UPDATE SET p.product_name = np.product_name, p.category = np.category WHERE p.product_name !=
```



```
'play gym'
WHEN NOT MATCHED THEN
 INSERT VALUES (np.product_id, np.product_name, np.category) WHERE np.category = 'books';
MERGE 4

-- Query updates.
openGauss=# SELECT * FROM products ORDER BY product_id;
 product_id | product_name | category
-----+-----+-----
 1501 | vivitar 35mm | electrncs
 1502 | olympus camera | electrncs
 1600 | play gym | toys
 1601 | lamaze | toys
 1666 | harry potter | toys
 1700 | wait interface | books
(6 rows)

-- Drop the table.
openGauss=# DROP TABLE products;
openGauss=# DROP TABLE newproducts;
```

## 7.13.127 MOVE

### Description

MOVE repositions a cursor without retrieving any data. MOVE works exactly like the **FETCH** statement, except it only repositions the cursor and does not return rows.

### Syntax

```
MOVE [direction [FROM | IN]] cursor_name;
```

The direction clause specifies optional parameters.

```
NEXT
| PRIOR
| FIRST
| LAST
| ABSOLUTE count
| RELATIVE count
| count
| ALL
| FORWARD
| FORWARD count
| FORWARD ALL
| BACKWARD
| BACKWARD count
| BACKWARD ALL
```

### Parameters

The parameters of **MOVE** and **FETCH** commands are the same. For details, see [Parameters](#) in "FETCH."

#### NOTE

On successful completion, the **MOVE** command returns a "**MOVE count**" tag, where *count* indicates the number of rows that the **FETCH** command with the same parameters would have returned (possibly zero).

### Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;
```

```
-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk INTEGER NOT NULL,
 r_reason_id CHAR(16) NOT NULL,
 r_reason_desc VARCHAR(40)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason VALUES (1, 'AAAAAAAABAAAAAAA', 'Xxxxxxxx'),(2,
'AAAAAAAACAAAAAAA', ' Xxxxxxxx'),(3, 'AAAAAAAADAAAAAAA', ' Xxxxxxxx'),(4, 'AAAAAAAEEAAAAAAA',
'Not the product that was ordered'),(5, 'AAAAAAAFAAAAAAA', 'Parts missing'),(6, 'AAAAAAGAAAAAAA',
'Does not work with a product that I have'),(7, 'AAAAAAAHAAAAAAA', 'Gift exchange');

-- Start a transaction.
openGauss=# START TRANSACTION;

-- Define a cursor cursor1.
openGauss=# CURSOR cursor1 FOR SELECT * FROM tpcds.reason;

-- Skip the first three rows of cursor1.
openGauss=# MOVE FORWARD 3 FROM cursor1;

-- Fetch the first four rows from cursor1.
openGauss=# FETCH 4 FROM cursor1;
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----
4 | AAAAAAAEEAAAAAAA | Not the product that was
ordered
5 | AAAAAAFAAAAAAAA | Parts missing
6 | AAAAAAGAAAAAAA | Does not work with a product that I
have
7 | AAAAAAHAAAAAAA | Gift
exchange
(4 rows)

-- Close the cursor.
openGauss=# CLOSE cursor1;

-- End the transaction.
openGauss=# END;

-- Drop the table.
openGauss=# DROP TABLE tpcds.reason;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[CLOSE](#) and [FETCH](#)

## 7.13.128 PREPARE

### Description

Creates a prepared statement.

A prepared statement is a performance optimizing object on the server. When **PREPARE** is executed, the specified query is parsed, analyzed, and rewritten. When **EXECUTE** is executed, the prepared statement is planned and executed. This avoids repetitive parsing and analysis. After the PREPARE statement is created, it exists throughout the database session. Once it is created (even if in a transaction

block), it will not be deleted when a transaction is rolled back. It can only be deleted by explicitly calling **DEALLOCATE** or automatically deleted when the session ends.

## Syntax

```
PREPARE name [(data_type [, ...])] AS statement;
```

## Parameters

- **name**  
Specifies the name of a prepared statement. It must be unique in the session.
- **data\_type**  
Specifies the type of an argument.
- **statement**  
Specifies a SELECT, INSERT, UPDATE, DELETE, MERGE INTO, or VALUES statement.

## Examples

See [Examples](#) in "EXECUTE."

## Helpful Links

[DEALLOCATE](#)

# 7.13.129 PREPARE TRANSACTION

## Description

Prepares the current transaction for two-phase commit.

After this statement, the transaction is no longer associated with the current session; instead, its state is fully stored on disk, and there is a high probability that it can be committed successfully, even if a database crash occurs before the commit is requested.

Once prepared, a transaction can later be committed or rolled back with **COMMIT PREPARED** or **ROLLBACK PREPARED**, respectively. Those statements can be issued from any session, not only the one that executed the original transaction.

From the point of view of the issuing session, PREPARE TRANSACTION is not unlike a ROLLBACK statement: after executing it, there is no active current transaction, and the effects of the prepared transaction are no longer visible. (The effects will become visible again if the transaction is committed.)

If the PREPARE TRANSACTION statement fails for any reason, it becomes a ROLLBACK and the current transaction is canceled.

## Precautions

- The transaction function is maintained automatically by the database, and should be not visible to users.

- When running the **PREPARE TRANSACTION** command, increase the value of **max\_prepared\_transactions** in configuration file **postgresql.conf**. You are advised to set it to a value not less than that of **max\_connections** so that one pending prepared transaction is available for each session.

## Syntax

```
PREPARE TRANSACTION transaction_id;
```

## Parameters

- transaction\_id**  
Specifies an arbitrary identifier that later identifies this transaction for COMMIT PREPARED or ROLLBACK PREPARED. The identifier must be different from those for current prepared transactions.  
Value range: The identifier must be written as a string literal, and must be less than 200 bytes long.

## Examples

```
-- Start.
openGauss=# BEGIN;
BEGIN

-- Prepare a transaction whose identifier is trans_test.
openGauss=# PREPARE TRANSACTION 'trans_test';
PREPARE TRANSACTION
```

## Helpful Links

[COMMIT PREPARED](#) and [ROLLBACK PREPARED](#)

## 7.13.130 PURGE

### Description

The PURGE statement can be used to:

- Clear tables or indexes from the recycle bin and release all space related to the objects.
- Clear the recycle bin.
- Clear the objects of a specified tablespace in the recycle bin.

### Precautions

- The PURGE operation supports tables (**PURGE TABLE**), indexes (**PURGE INDEX**), and recycle bins (**PURGE RECYCLEBIN**).
- The permission requirements for performing the PURGE operation are as follows:
  - PURGE TABLE:** The user must be the owner of the table and must have the USAGE permission on the schema to which the table belongs. System administrators have this permission by default.

- **PURGE INDEX:** The user must be the owner of the index and have the USAGE permission on the schema to which the index belongs. By default, system administrators have this permission.
- **PURGE RECYCLEBIN:** Common users can clear only the objects owned by themselves in the recycle bin. In addition, the user must have the USAGE permission of the schema to which the objects belong. By default, system administrators can clear all objects in the recycle bin.

## Syntax

```
PURGE { TABLE [schema_name.]table_name
 | INDEX index_name
 | RECYCLEBIN
 }
```

## Parameters

- **schema\_name**  
Schema name
- **TABLE [ schema\_name. ] table\_name**  
Clears a specified table in the recycle bin.
- **INDEX index\_name**  
Clears a specified index in the recycle bin.
- **RECYCLEBIN**  
Clears the objects in the recycle bin.

## Examples

```
-- Create the reason_table_space tablespace.
openGauss=# CREATE TABLESPACE REASON_TABLE_SPACE1 owner tpcds RELATIVE location 'tablespace/
tsp_reason1';
-- Create the tpcds.reason_t1 table in the tablespace.
openGauss=# CREATE TABLE tpcds.reason_t1
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
) tablespace reason_table_space1;
-- Create the tpcds.reason_t2 table in the tablespace.
openGauss=# CREATE TABLE tpcds.reason_t2
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
) tablespace reason_table_space1;
-- Create the tpcds.reason_t3 table in the tablespace.
openGauss=# CREATE TABLE tpcds.reason_t3
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
) tablespace reason_table_space1;
-- Create an index on the tpcds.reason_t1 table.
openGauss=# CREATE INDEX index_t1 on tpcds.reason_t1(r_reason_id);
openGauss=# DROP TABLE tpcds.reason_t1;
openGauss=# DROP TABLE tpcds.reason_t2;
openGauss=# DROP TABLE tpcds.reason_t3;
-- View the recycle bin.
openGauss=# SELECT rcyname,rcyoriginname,rcytablespace FROM GS_RECYCLEBIN;
rcyname | rcyoriginname | rcytablespace
```

```
-----+-----+-----
BIN$16409$2CEE988==$0 | reason_t1 | 16408
BIN$16412$2CF2188==$0 | reason_t2 | 16408
BIN$16415$2CF2EC8==$0 | reason_t3 | 16408
BIN$16418$2CF3EC8==$0 | index_t1 | 0
(4 rows)
-- Purge the table.
openGauss=# PURGE TABLE tpcds.reason_t3;
openGauss=# SELECT rcynname,rcyoriginname,rcytablespace FROM GS_RECYCLEBIN;
 rcynname | rcyoriginname | rcytablespace
-----+-----+-----
BIN$16409$2CEE988==$0 | reason_t1 | 16408
BIN$16412$2CF2188==$0 | reason_t2 | 16408
BIN$16418$2CF3EC8==$0 | index_t1 | 0
(3 rows)
-- Purge the index.
openGauss=# PURGE INDEX tindex_t1;
openGauss=# SELECT rcynname,rcyoriginname,rcytablespace FROM GS_RECYCLEBIN;
 rcynname | rcyoriginname | rcytablespace
-----+-----+-----
BIN$16409$2CEE988==$0 | reason_t1 | 16408
BIN$16412$2CF2188==$0 | reason_t2 | 16408
(2 rows)
-- Purge all objects in the recycle bin.
openGauss=# PURGE recyclebin;
openGauss=# SELECT rcynname,rcyoriginname,rcytablespace FROM GS_RECYCLEBIN;
 rcynname | rcyoriginname | rcytablespace
-----+-----+-----
(0 rows)
```

## 7.13.131 REASSIGN OWNED

### Description

Changes the owner of the database object.

REASSIGN OWNED requires that the system change the owner of all the database objects owned by an old role to a new role.

### Precautions

- REASSIGN OWNED is often executed before role deletion.
- To run **REASSIGN OWNED**, you must have the permissions of the original and target roles.

### Syntax

```
REASSIGN OWNED BY old_role [, ...] TO new_role;
```

### Parameters

- **old\_role**  
Specifies the role name of the old owner.
- **new\_role**  
Specifies the role name of the new owner. Note: Only the initial user can use the REASSIGN OWNED syntax to change the owner to the initial user.

### Examples

None

## 7.13.132 REFRESH INCREMENTAL MATERIALIZED VIEW

### Function

**REFRESH INCREMENTAL MATERIALIZED VIEW** refreshes a materialized view in incremental mode.

### Precautions

- Incremental refresh supports only fast-refresh materialized views.
- To refresh a materialized view, you must have the SELECT permission on the base table.

### Syntax

```
REFRESH INCREMENTAL MATERIALIZED VIEW mv_name;
```

### Parameter Description

- **mv\_name**  
Name of the materialized view to be refreshed.

### Examples

```
-- Create an ordinary table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int);

-- Create a fast-refresh materialized view.
openGauss=# CREATE INCREMENTAL MATERIALIZED VIEW my_imv AS SELECT * FROM my_table;

-- Write data to the base table.
openGauss=# INSERT INTO my_table VALUES(1,1),(2,2);

-- Fast refresh the fast-refresh materialized view my_imv.
openGauss=# REFRESH INCREMENTAL MATERIALIZED VIEW my_imv;

-- Delete the fast-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW my_imv;

-- Delete the my_table table.
openGauss=# DROP TABLE my_table;
```

### Helpful Links

[ALTER MATERIALIZED VIEW](#), [CREATE INCREMENTAL MATERIALIZED VIEW](#), [CREATE MATERIALIZED VIEW](#), [CREATE TABLE](#), [DROP MATERIALIZED VIEW](#), and [REFRESH MATERIALIZED VIEW](#)

## 7.13.133 REFRESH MATERIALIZED VIEW

### Function

**REFRESH MATERIALIZED VIEW** refreshes materialized views in full refresh mode.

### Precautions

- Full refreshing can be performed on both complete- and fast-refresh materialized views.

- To refresh a materialized view, you must have the SELECT permission on the base table.

## Syntax

```
REFRESH MATERIALIZED VIEW mv_name;
```

## Parameter Description

- **mv\_name**  
Name of the materialized view to be refreshed.

## Examples

```
-- Create an ordinary table.
openGauss=# CREATE TABLE my_table (c1 int, c2 int);

-- Create a complete-refresh materialized view.
openGauss=# CREATE MATERIALIZED VIEW my_mv AS SELECT * FROM my_table;

-- Create a fast-refresh materialized view.
openGauss=# CREATE INCREMENTAL MATERIALIZED VIEW my_imv AS SELECT * FROM my_table;

-- Write data to the base table.
openGauss=# INSERT INTO my_table VALUES(1,1),(2,2);

-- Completely refresh the complete-refresh materialized view my_mv.
openGauss=# REFRESH MATERIALIZED VIEW my_mv;

-- Completely refresh the fast-refresh materialized view my_imv.
openGauss=# REFRESH MATERIALIZED VIEW my_imv;

-- Delete the fast-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW my_imv;

-- Delete the complete-refresh materialized view.
openGauss=# DROP MATERIALIZED VIEW my_mv;

-- Delete the my_table table.
openGauss=# DROP TABLE my_table;
```

## Helpful Links

[ALTER MATERIALIZED VIEW](#), [CREATE INCREMENTAL MATERIALIZED VIEW](#), [CREATE MATERIALIZED VIEW](#), [CREATE TABLE](#), [DROP MATERIALIZED VIEW](#), and [REFRESH INCREMENTAL MATERIALIZED VIEW](#)

## 7.13.134 REINDEX

### Description

Rebuilds an index using the data stored in the index's table, replacing the old copy of the index.

There are several scenarios in which REINDEX can be used:

- An index has become corrupted, and no longer contains valid data.
- An index has become "bloated", that is, it contains many empty or nearly-empty pages.



- You have altered a storage parameter (such as a fill factor) for an index, and wish that the change takes full effect.
- An index build with the **CONCURRENTLY** option failed, leaving an "invalid" index.

## Precautions

**REINDEX DATABASE** and **REINDEX SYSTEM** cannot be executed in transaction blocks. Currently, REINDEX operations cannot be performed on materialized views.

## Syntax

- Rebuild an ordinary index.  
`REINDEX { INDEX | [INTERNAL] TABLE | DATABASE | SYSTEM } name [ FORCE ];`
- Rebuild an index partition.  
`REINDEX { INDEX | [INTERNAL] TABLE } name  
PARTITION partition_name [ FORCE ];`

## Parameters

- **INDEX**  
Rebuilds the specified index.
- **TABLE**  
Rebuilds all indexes of a specified table. The TOAST table (if any) of the table is reindexed as well. If an index in the table has been invalidated by running **alter unusable**, the index cannot be rebuilt.
- **DATABASE**  
Rebuilds all indexes within the current database.
- **SYSTEM**  
Rebuilds all indexes on system catalogs within the current database. Indexes on user tables are not processed.
- **name**  
Specifies the name of the index, table, or database whose index needs to be rebuilt. Tables and indexes can be schema-qualified.  
  

 **NOTE**

REINDEX DATABASE and REINDEX SYSTEM can rebuild indexes for only the current database. Therefore, **name** must be the same as the current database name.
- **FORCE**  
Specifies an invalid option, which will be ignored.
- **partition\_name**  
Specifies the name of the partition or index partition to be rebuilt.  
Value range:
  - If REINDEX INDEX is used, specify the name of an index partition.
  - If REINDEX TABLE is used, specify the name of a partition.

**NOTICE**

**REINDEX DATABASE** and **REINDEX SYSTEM** cannot be executed in transaction blocks.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.customer table.
openGauss=# CREATE TABLE tpcds.customer
(
 c_customer_sk INTEGER NOT NULL,
 c_customer_id CHAR(16) NOT NULL
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.customer VALUES (1, 'AAAAAAAAABAAAAAA'),(5,
'AAAAAAAAACAAAAAA'),(10, 'AAAAAADAAAAAA');

-- Create a row-store table tpcds.customer_t1 and create an index on the c_customer_sk column in the
table.
openGauss=# CREATE TABLE tpcds.customer_t1
(
 c_customer_sk integer not null,
 c_customer_id char(16) not null,
 c_current_demo_sk integer ,
 c_current_hdemo_sk integer ,
 c_current_addr_sk integer ,
 c_first_shipto_date_sk integer ,
 c_first_sales_date_sk integer ,
 c_salutation char(10) ,
 c_first_name char(20) ,
 c_last_name char(30) ,
 c_preferred_cust_flag char(1) ,
 c_birth_day integer ,
 c_birth_month integer ,
 c_birth_year integer ,
 c_birth_country varchar(20) ,
 c_login char(13) ,
 c_email_address char(50) ,
 c_last_review_date char(10)
)
WITH (orientation = row);

openGauss=# CREATE INDEX tpcds_customer_index1 ON tpcds.customer_t1 (c_customer_sk);

openGauss=# INSERT INTO tpcds.customer_t1 SELECT * FROM tpcds.customer WHERE c_customer_sk < 10;

-- Rebuild a single index.
openGauss=# REINDEX INDEX tpcds.tpcds_customer_index1;

-- Rebuild all indexes in the tpcds.customer_t1 table:
openGauss=# REINDEX TABLE tpcds.customer_t1;

-- Drop the tpcds.customer_t1 table.
openGauss=# DROP TABLE tpcds.customer_t1;

-- Drop the table.
openGauss=# DROP TABLE tpcds.customer;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Suggestions

- **INTERNAL TABLE**  
This scenario is used for fault recovery. You are advised not to perform concurrent operations.
- **DATABASE**  
You are not allowed to re-index a database in a transaction.
- **SYSTEM**  
You are not allowed to re-index system catalogs in transactions.

## 7.13.135 RELEASE SAVEPOINT

### Description

RELEASE SAVEPOINT deletes a savepoint previously defined in the current transaction.

Deleting a savepoint makes it unavailable as a rollback point, but it has no other user visible behavior. It does not undo the effects of statements executed after the savepoint was established. To do that, use ROLLBACK TO SAVEPOINT. Deleting a savepoint when it is no longer needed allows the system to recycle some resources earlier than transaction end.

RELEASE SAVEPOINT also deletes all savepoints that were established after the named savepoint was established.

### Precautions

- Releasing a savepoint name that was not previously defined will cause an error.
- It is not possible to release a savepoint when the transaction is in an aborted state.
- If multiple savepoints have the same name, only the one that was most recently defined is released.

### Syntax

```
RELEASE [SAVEPOINT] savepoint_name;
```

### Parameters

**savepoint\_name**

Specifies the name of the savepoint you want to delete.

### Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create a table.
openGauss=# CREATE TABLE tpcds.table1(a int);

-- Start a transaction.
openGauss=# START TRANSACTION;
```

```
-- Insert data.
openGauss=# INSERT INTO tpcds.table1 VALUES (3);

-- Create a savepoint.
openGauss=# SAVEPOINT my_savepoint;

-- Insert data.
openGauss=# INSERT INTO tpcds.table1 VALUES (4);

-- Delete the savepoint.
openGauss=# RELEASE SAVEPOINT my_savepoint;

-- Commit the transaction.
openGauss=# COMMIT;

-- Query the table content, which should contain both 3 and 4.
openGauss=# SELECT * FROM tpcds.table1;

-- Drop the table.
openGauss=# DROP TABLE tpcds.table1;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[SAVEPOINT](#) and [ROLLBACK TO SAVEPOINT](#)

## 7.13.136 RESET

### Description

RESET restores GUC parameters to their default values. The default values are parameter default values compiled in the **postgresql.conf** configuration file.

**RESET** is an alternative spelling for:

```
SET configuration_parameter TO DEFAULT;
```

### Precautions

RESET and SET have the same transaction behavior. Their impact will be rolled back.

### Syntax

```
RESET {configuration_parameter | CURRENT_SCHEMA | TIME_ZONE | TRANSACTION ISOLATION LEVEL |
SESSION AUTHORIZATION | ALL };
```

### Parameters

- **configuration\_parameter**  
Specifies the name of a GUC parameter.  
Value range: GUC parameters. You can view them by running the **SHOW ALL** command.
- **CURRENT\_SCHEMA**  
Specifies the current schema.

- **TIME ZONE**  
Specifies the time zone.
- **TRANSACTION ISOLATION LEVEL**  
Specifies the transaction isolation level.
- **SESSION AUTHORIZATION**  
Specifies the session authorization.
- **ALL**  
Specifies all GUC parameters.

## Examples

```
-- Reset timezone to the default value.
openGauss=# RESET timezone;

-- Set all parameters to their default values.
openGauss=# RESET ALL;
```

## Helpful Links

[SET](#) and [SHOW](#)

## 7.13.137 REVOKE

### Description

REVOKE revokes permissions from one or more roles.

### Precautions

If a non-owner user of an object attempts to the REVOKE permission on the object, the statement is executed based on the following rules:

- If the user has no permissions whatsoever on the object, the statement will fail outright.
- If an authorized user has some permissions, only the permissions with authorization options are revoked.
- If the authorized user does not have the authorization option, the REVOKE ALL PRIVILEGES form will issue an error message. For other forms of statements, if the permission specified in the statement does not have the corresponding authorization option, the statement will issue a warning.

### Syntax

- Revoke the permission on a specified table or view.  

```
REVOKE [GRANT OPTION FOR]
 { { SELECT | INSERT | UPDATE | DELETE | TRUNCATE | REFERENCES | ALTER | DROP | COMMENT |
 INDEX | VACUUM }, ... }
 | ALL [PRIVILEGES] }
ON { [TABLE] table_name [, ...]
 | ALL TABLES IN SCHEMA schema_name [, ...] }
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```
- Revoke the permission on a specified field in a table.

```
REVOKE [GRANT OPTION FOR]
{ { SELECT | INSERT | UPDATE | REFERENCES | COMMENT } (column_name [, ...]) [, ...]
| ALL [PRIVILEGES] (column_name [, ...]) }
ON [TABLE] table_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified sequence. The **LARGE** column is optional. The recycling statement does not distinguish whether the sequence is **LARGE**.

```
REVOKE [GRANT OPTION FOR]
{ { SELECT | UPDATE | ALTER | DROP | COMMENT } [, ...]
| ALL [PRIVILEGES] }
ON { [[LARGE] SEQUENCE] sequence_name [, ...]
| ALL SEQUENCES IN SCHEMA schema_name [, ...] }
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified database.

```
REVOKE [GRANT OPTION FOR]
{ { CREATE | CONNECT | TEMPORARY | TEMP | ALTER | DROP | COMMENT } [, ...]
| ALL [PRIVILEGES] }
ON DATABASE database_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified domain.

```
REVOKE [GRANT OPTION FOR]
{ USAGE | ALL [PRIVILEGES] }
ON DOMAIN domain_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified directory.

```
REVOKE [GRANT OPTION FOR]
{ { READ | WRITE | ALTER | DROP } [, ...] | ALL [PRIVILEGES] }
ON DIRECTORY directory_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified external data source.

```
REVOKE [GRANT OPTION FOR]
{ USAGE | ALL [PRIVILEGES] }
ON FOREIGN DATA WRAPPER fdw_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified external server.

```
REVOKE [GRANT OPTION FOR]
{ { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON FOREIGN SERVER server_name [, ...]
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified function.

```
REVOKE [GRANT OPTION FOR]
{ { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON { FUNCTION {function_name ([{ [argmode] [arg_name] arg_type } [, ...]) } } [, ...]
| ALL FUNCTIONS IN SCHEMA schema_name [, ...] }
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified stored procedure.

```
REVOKE [GRANT OPTION FOR]
{ { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
ON { PROCEDURE {proc_name ([{ [argmode] [arg_name] arg_type } [, ...]) } } [, ...]
| ALL PROCEDURE IN SCHEMA schema_name [, ...] }
FROM { [GROUP] role_name | PUBLIC } [, ...]
[CASCADE | RESTRICT];
```

- Revoke the permission on a specified procedural language.

```
REVOKE [GRANT OPTION FOR]
 { USAGE | ALL [PRIVILEGES] }
 ON LANGUAGE lang_name [, ...]
 FROM { [GROUP] role_name | PUBLIC } [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the permission on a specified large object.

```
REVOKE [GRANT OPTION FOR]
 { { SELECT | UPDATE } [, ...] | ALL [PRIVILEGES] }
 ON LARGE OBJECT loid [, ...]
 FROM { [GROUP] role_name | PUBLIC } [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the permission on a specified schema.

```
REVOKE [GRANT OPTION FOR]
 { { CREATE | USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON SCHEMA schema_name [, ...]
 FROM { [GROUP] role_name | PUBLIC } [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the permission on a specified tablespace.

```
REVOKE [GRANT OPTION FOR]
 { { CREATE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON TABLESPACE tablespace_name [, ...]
 FROM { [GROUP] role_name | PUBLIC } [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the permission on a specified type.

```
REVOKE [GRANT OPTION FOR]
 { { USAGE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON TYPE type_name [, ...]
 FROM { [GROUP] role_name | PUBLIC } [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the permission on a package object.

```
REVOKE [GRANT OPTION FOR]
 { { EXECUTE | ALTER | DROP | COMMENT } [, ...] | ALL [PRIVILEGES] }
 ON PACKAGE package_name [, ...]
 FROM {[GROUP] role_name | PUBLIC} [, ...]
 [CASCADE | RESTRICT];
```

- Revoke permissions from a role.

```
REVOKE [ADMIN OPTION FOR]
 role_name [, ...] FROM role_name [, ...]
 [CASCADE | RESTRICT];
```

- Revoke the SYSADMIN permission from a role.

```
REVOKE ALL { PRIVILEGES | PRIVILEGE } FROM role_name;
```

- Revoke the ANY permissions.

```
REVOKE [ADMIN OPTION FOR]
 { CREATE ANY TABLE | ALTER ANY TABLE | DROP ANY TABLE | SELECT ANY TABLE | INSERT ANY
 TABLE | UPDATE ANY TABLE |
 DELETE ANY TABLE | CREATE ANY SEQUENCE | CREATE ANY INDEX | CREATE ANY FUNCTION |
 EXECUTE ANY FUNCTION |
 CREATE ANY PACKAGE | EXECUTE ANY PACKAGE | CREATE ANY TYPE } [, ...]
 FROM [GROUP] role_name [, ...];
```

## Parameters

The keyword PUBLIC indicates an implicitly defined group that has all roles.

For details about permission types and parameters, see [Parameters](#) in "GRANT".

Permissions of a role include the permissions directly granted to the role, permissions inherited from the parent role, and permissions granted to PUBLIC users. Therefore, revoking the SELECT permission on an object from PUBLIC users does not necessarily mean that the SELECT permission on the object has been revoked from all roles, because the SELECT permission directly granted to roles

and inherited from parent roles remains. Similarly, if the SELECT permission is revoked from a user but is not revoked from PUBLIC users, the user can still use SELECT.

If **GRANT OPTION FOR** is specified, the permission cannot be granted to others, but permission itself is not revoked.

If user A holds the UPDATE permission (with **WITH GRANT OPTION**) on a table and has granted it to user B, the permission that user B holds is called dependent permission. When user A's permission or grant option is revoked, **CASCADE** must be declared to revoke all dependent permissions.

A user can only revoke permissions that were granted directly by that user. For example, if user A has granted a permission with grant option (**WITH ADMIN OPTION**) to user B, and user B has in turn granted it to user C, then user A cannot revoke the permission directly from C. However, user A can revoke the grant option held by user B and use **CASCADE**. In this way, the permission of user C is automatically revoked. For another example, if both user A and user B have granted the same permission to C, A can revoke A's own grant but not B's grant, so C will still effectively have the permission.

If the role executing **REVOKE** holds permissions indirectly via more than one role membership path, it is unspecified which containing role will be used to execute the statement. In such cases, it is the best practice to use SET ROLE to become the specific role, and then execute **REVOKE**. Failure to do so may lead to deleting permissions not intended to delete, or not deleting any permissions at all.

## Examples

See [Examples](#) in "GRANT."

## Helpful Links

[GRANT](#)

## 7.13.138 ROLLBACK

### Description

Rolls back the current transaction and backs out all updates in the transaction.

If a fault occurs during the running of a transaction, the transaction cannot be executed. The system cancels all completed operations on the database in the transaction, and the database status returns to the time when the transaction starts.

### Precautions

If a **ROLLBACK** statement is executed out of a transaction, no error occurs, but a notice is displayed.

### Syntax

```
ROLLBACK [WORK | TRANSACTION];
```



## Parameters

### WORK | TRANSACTION

Specifies the optional keyword that more clearly illustrates the syntax.

## Examples

```
-- Start a transaction.
openGauss=# START TRANSACTION;

-- Back out all changes.
openGauss=# ROLLBACK;
```

## Helpful Links

[COMMIT](#) | [END](#)

## 7.13.139 ROLLBACK PREPARED

### Function

Cancels a transaction ready for two-phase committing.

### Precautions

- The function is only available in maintenance mode (when GUC parameter **xc\_maintenance\_mode** is **on**). Exercise caution when enabling the mode. It is used by maintenance engineers for troubleshooting. Common users should not use the mode.
- Only the user that initiates a transaction or the system administrator can roll back the transaction.
- The transaction function is maintained automatically by the database, and should be not visible to users.

### Syntax

```
ROLLBACK PREPARED transaction_id ;
```

### Parameter Description

#### transaction\_id

Specifies the identifier of the transaction to be committed. The identifier must be different from those for current prepared transactions.

## Examples

```
-- Start.
openGauss=# BEGIN;
BEGIN

-- Prepare a transaction whose identifier is trans_test.
openGauss=# PREPARE TRANSACTION 'trans_test';
PREPARE TRANSACTION

-- Cancel the transaction whose identifier is trans_test.
```

```
openGauss=# ROLLBACK PREPARED 'trans_test';
ROLLBACK PREPARED
```

## Helpful Links

[COMMIT PREPARED](#) and [PREPARE TRANSACTION](#)

## 7.13.140 ROLLBACK TO SAVEPOINT

### Description

ROLLBACK TO SAVEPOINT rolls back to a savepoint. It implicitly deletes all savepoints that were established after the named savepoint.

Rolls back all statements that were executed after the savepoint was established. The savepoint remains valid and can be rolled back to again later, if needed.

### Precautions

- Specifying a savepoint name that has not been established is an error.
- Cursors have somewhat non-transactional behavior with respect to savepoints. Any cursor that is opened inside a savepoint will be closed when the savepoint is rolled back. If a previously opened cursor is affected by a FETCH statement inside a savepoint that is later rolled back, the cursor remains at the position that FETCH left it pointing to (that is, the cursor motion caused by FETCH is not rolled back). Closing a cursor is not undone by rolling back, either. A cursor whose execution causes a transaction to abort is put in a cannot-execute state, so while the transaction can be restored using ROLLBACK TO SAVEPOINT, the cursor can no longer be used.
- Use ROLLBACK TO SAVEPOINT to roll back to a savepoint. Use RELEASE SAVEPOINT to delete a savepoint but keep the effects of the commands executed after the savepoint was established.

### Syntax

```
ROLLBACK [WORK | TRANSACTION] TO [SAVEPOINT] savepoint_name;
```

### Parameters

*savepoint\_name*

Rolls back to a savepoint.

### Examples

```
-- Undo the effects of the statements executed after my_savepoint was established:
openGauss=# START TRANSACTION;
openGauss=# SAVEPOINT my_savepoint;
openGauss=# ROLLBACK TO SAVEPOINT my_savepoint;
-- Cursor positions are not affected by savepoint rollback.
openGauss=# DECLARE foo CURSOR FOR SELECT 1 UNION SELECT 2;
openGauss=# SAVEPOINT foo;
openGauss=# FETCH 1 FROM foo;
?column?

1
openGauss=# ROLLBACK TO SAVEPOINT foo;
```

```
openGauss=# FETCH 1 FROM foo;
?column?

2
openGauss=# RELEASE SAVEPOINT my_savepoint;
openGauss=# COMMIT;
```

## Helpful Links

[SAVEPOINT](#) and [RELEASE SAVEPOINT](#)

## 7.13.141 SAVEPOINT

### Description

SAVEPOINT establishes a new savepoint in the current transaction.

A savepoint is a special mark inside a transaction. It allows all statements that are executed after it was established to be rolled back, restoring the transaction state to what it was at the time of the savepoint.

### Precautions

- Use ROLLBACK TO SAVEPOINT to roll back to a savepoint. Use RELEASE SAVEPOINT to delete a savepoint but keep the effects of the commands executed after the savepoint was established.
- Savepoints can only be established when inside a transaction block. Multiple savepoints can be defined in a transaction.
- In the case of an unexpected termination of a thread or process caused by a node or connection failure, or of an error caused by the inconsistency between source and destination table structures in a COPY FROM operation, the transaction cannot be rolled back to the established savepoint. Instead, the entire transaction will be rolled back.
- According to the SQL standard, when a savepoint with the same name is created, the previous savepoint with the same name is automatically deleted. In GaussDB, the old savepoint is retained, but only the latest one is used during rollback or release. If the latest savepoint is released, the previous savepoint will again become accessible to ROLLBACK TO SAVEPOINT and RELEASE SAVEPOINT. In addition, SAVEPOINT fully complies with the SQL standard.

### Syntax

```
SAVEPOINT savepoint_name;
```

### Parameters

savepoint\_name

Specifies the name of the new savepoint.

**NOTICE**

When using SAVEPOINT, you are advised to release SAVEPOINT promptly to avoid too many nested subtransactions. It is recommended that the number of nested subtransactions be less than or equal to 10,000. If the number of nested subtransactions is too large, the current transaction performance may deteriorate.

## Examples

```
-- Create a table.
openGauss=# CREATE TABLE table1(a int);

-- Start a transaction.
openGauss=# START TRANSACTION;

-- Insert data.
openGauss=# INSERT INTO table1 VALUES (1);

-- Create a savepoint.
openGauss=# SAVEPOINT my_savepoint;

-- Insert data.
openGauss=# INSERT INTO table1 VALUES (2);

-- Roll back the savepoint.
openGauss=# ROLLBACK TO SAVEPOINT my_savepoint;

-- Insert data.
openGauss=# INSERT INTO table1 VALUES (3);

-- Commit the transaction.
openGauss=# COMMIT;

-- Query the content of the table. You can see 1 and 3 at the same time, but cannot see 2 because 2 is
rolled back.
openGauss=# SELECT * FROM table1;

-- Drop the table.
openGauss=# DROP TABLE table1;

-- Create a table.
openGauss=# CREATE TABLE table2(a int);

-- Start a transaction.
openGauss=# START TRANSACTION;

-- Insert data.
openGauss=# INSERT INTO table2 VALUES (3);

-- Create a savepoint.
openGauss=# SAVEPOINT my_savepoint;

-- Insert data.
openGauss=# INSERT INTO table2 VALUES (4);

-- Roll back the savepoint.
openGauss=# RELEASE SAVEPOINT my_savepoint;

-- Commit the transaction.
openGauss=# COMMIT;

-- Query the table content, which should contain both 3 and 4.
openGauss=# SELECT * FROM table2;

-- Drop the table.
openGauss=# DROP TABLE table2;
```

## Helpful Links

[RELEASE SAVEPOINT](#) and [ROLLBACK TO SAVEPOINT](#)

## 7.13.142 SELECT

### Description

SELECT retrieves data from tables or views.

Serving as an overlaid filter on database tables, the SELECT statement uses SQL keywords to filter data tables and extract the required data.

### Precautions

- The owner of a table, users granted the SELECT permission on the table, or users granted the SELECT ANY TABLE permission can read data in the table or view. System administrators have the permission to read data in the table or view by default.
- You must have the SELECT permission on each field used in the SELECT statement.
- Using FOR UPDATE, FOR NO KEY UPDATE, FOR SHARE, or FOR KEY SHARE also requires the UPDATE permission in addition to the SELECT permission.

### Syntax

```
• Query data.
[WITH [RECURSIVE] with_query [, ...]]
SELECT [/*+ plan_hint */] [ALL | DISTINCT [ON (expression [, ...])]]
{ * | {expression [[AS] output_name]} [, ...] }
[FROM from_item [, ...]]
[WHERE condition]
[[START WITH condition] CONNECT BY [NOCYCLE] condition [ORDER SIBLINGS BY expression]]
[GROUP BY grouping_element [, ...]]
[HAVING condition [, ...]]
[WINDOW {window_name AS (window_definition)} [, ...]]
[{ UNION | INTERSECT | EXCEPT | MINUS } [ALL | DISTINCT] select]
[ORDER BY {expression [[ASC | DESC | USING operator] | nlssort_expression_clause } [NULLS { FIRST | LAST }]} [, ...]]
[LIMIT { [offset,] count | ALL }]
[OFFSET start [ROW | ROWS]]
[FETCH { FIRST | NEXT } [count] { ROW | ROWS } ONLY]
[{FOR { UPDATE | NO KEY UPDATE | SHARE | KEY SHARE } [OF table_name [, ...]] [NOWAIT | WAIT n]}
[...]];
TABLE { ONLY {(table_name)| table_name} | table_name [*]};
```

 NOTE

In condition and expression, you can use the aliases of expressions in **targetlist** in compliance with the following rules:

- Reference only within the same level.
- Only reference aliases in **targetlist**.
- Reference a prior expression in a subsequent expression.
- Volatile functions cannot be used.
- Window functions cannot be used.
- Aliases cannot be referenced in the condition of JOIN ON.
- An error is reported if **targetlist** contains multiple aliases to be referenced.
- The subquery **with\_query** is as follows:  

```
with_query_name [(column_name [, ...])]
AS [[NOT] MATERIALIZED] ({select | values | insert | update | delete})
```
- The specified query source **from\_item** is as follows:  

```
{[ONLY] table_name [*] [partition_clause] [[AS] alias [(column_alias [, ...])]]
[TABLESAMPLE sampling_method (argument [, ...]) [REPEATABLE (seed)]]
[TIMECAPSULE {TIMESTAMP | CSN} expression]
[(select) [AS] alias [(column_alias [, ...])]]
[with_query_name [[AS] alias [(column_alias [, ...])]]]
[function_name ([argument [, ...]]) [AS] alias [(column_alias [, ...] | column_definition [, ...])]
[function_name ([argument [, ...]]) AS (column_definition [, ...])]
[from_item [NATURAL] join_type from_item [ON join_condition | USING (join_column [, ...])]]
```
- The GROUP clause is as follows:  

```
()
| expression
| (expression [, ...])
| ROLLUP ({ expression | (expression [, ...]) } [, ...])
| CUBE ({ expression | (expression [, ...]) } [, ...])
| GROUPING SETS (grouping_element [, ...])
```
- The specified partition **partition\_clause** in **from\_item** is as follows:  

```
PARTITION { (partition_name) | FOR (partition_value [, ...]) } |
SUBPARTITION { (subpartition_name) | FOR (subpartition_value [, ...]) }
```

 NOTE

Specifying partitions applies only to partitioned tables.

- The sorting order **nlsort\_expression\_clause** is as follows:  

```
NLSSORT (column_name, ' NLS_SORT = { SCHINESE_PINYIN_M | generic_m_ci } ')
```

The second parameter can be **generic\_m\_ci**, which supports only the case-insensitive order for English characters.
- Simplified query syntax, equivalent to SELECT \* FROM *table\_name*.  

```
TABLE { ONLY {(table_name)| table_name} | table_name [*]};
```

## Parameters

- **WITH [ RECURSIVE ] with\_query [, ...]**  

Specifies one or more subqueries that can be referenced by name in the main query, which is equivalent to a temporary table. This subquery statement structure is called the common table expression (CTE) structure. When this structure is used, the execution plan contains the CTE SCAN content.

If **RECURSIVE** is specified, it allows a SELECT subquery to reference itself by name.

Format of **with\_query**:

```
with_query_name [(column_name [, ...])] AS [[NOT] MATERIALIZED] ({select | values | insert |
update | delete})
```

- **with\_query\_name** specifies the name of the result set generated by a subquery. Such names can be used to access the result sets of subqueries in a query.
- **column\_name** specifies the column name displayed in the subquery result set.
- Each subquery can be a SELECT, VALUES, INSERT, UPDATE, or DELETE statement.
- RECURSIVE can appear only after WITH. In the case of multiple CTEs, you only need to declare RECURSIVE at the first CTE.
- You can use **MATERIALIZED** or **NOT MATERIALIZED** to modify the CTE.
  - If **MATERIALIZED** is specified, the WITH query will be materialized, and a copy of the subquery result set is generated. The copy is directly queried at the reference point. Therefore, the WITH subquery cannot be jointly optimized with the SELECT statement trunk (for example, predicate pushdown and equivalence class transfer). In this scenario, you can use **NOT MATERIALIZED** for modification. If the WITH query can be executed as a subquery inline, the preceding optimization can be performed.
  - If the user does not explicitly declare the materialized attribute, comply with the following rules: If the CTE is referenced only once in the SELECT statement trunk to which it belongs and semantically supports inline execution, it will be rewritten as subquery inline execution. Otherwise, the materialized execution will be performed in CTE Scan mode.
- **plan\_hint** clause  
Follows the SELECT keyword in the */\*+<Plan hint> \*/* format. It is used to optimize the plan of a **SELECT** statement block. For details, see [Hint-based Tuning](#). In each statement, only the first */\*+ plan\_hint \*/* comment block takes effect as a hint. Multiple hints can be written.
- **ALL**  
Specifies that all rows that meet the conditions are returned. This is the default behavior and can be omitted.
- **DISTINCT [ ON ( expression [, ...] ) ]**  
Removes all duplicate rows from the result set of the SELECT statement so that each row in the result set is unique.  
Retains only the first row in the set of rows that have the same result calculated on the given expression.

---

**NOTICE**

DISTINCT ON expression is explained with the same rule of ORDER BY. Unless you use ORDER BY to guarantee that the required row appears first, you cannot know what the first row is.

---

- **SELECT** list

Specifies the name of a column in the table to be queried. The value can be a part of the column name or all of the column names. The wildcard (\*) is used to represent the column name.

You may use the *AS output\_name* clause to assign an alias to the output column. The alias is usually used for displaying the output column. The name, value, and type keywords can be used as column aliases.

Column names can be expressed in the following formats:

- Manually input column names which are spaced using commas (,).
- Columns computed in the FROM clause.
- FROM clause

Specifies one or more source tables for SELECT.

The FROM clause can contain the following elements:

- table\_name

Specifies the name of a table or view. The schema name can be added before the table name or view name, for example, *schema\_name.table\_name*.

- alias

Gives a temporary alias to a table to facilitate the reference by other queries.

An alias is used for brevity or to eliminate ambiguity for self-joins. If an alias is provided, it completely hides the actual name of the table.

- TABLESAMPLE sampling\_method ( argument [, ...] ) [ REPEATABLE ( seed ) ]

The TABLESAMPLE clause following *table\_name* specifies that the specified *sampling\_method* should be used to retrieve the subset of rows in the table.

The optional REPEATABLE clause specifies the number of seeds used to generate random numbers in the sampling method. The seed value can be any non-null constant value. If the table was not changed during the query, the two queries having the same seed and *argument* values will select the same sampling in this table. However, different seed values usually generate different samples. If **REPEATABLE** is not specified, a new random sample will be selected for each query based on the seed generated by the system.

- TIMECAPSULE { TIMESTAMP | CSN } expression

Queries the table data of a specified CSN or at a specified time point.

Currently, the following tables do not support flashback query: system catalogs, MOTs, DFS tables, global temporary tables, local temporary tables, unlogged tables, views, sequence tables, hash bucket tables, shared tables, inherited tables, and tables with the PARTIAL CLUSTER KEY constraint.

- TIMECAPSULE TIMESTAMP

Searches for the result set of a specified time point based on the date as the flashback query flag. *date* must be a valid past timestamp.



- TIMECAPSULE CSN

Searches for the result set of a specified CSN based on the CSN flashback of the table as the flashback query flag. The CSN can be obtained from **snpcsn** recorded in **gs\_txn\_snapshot**.

- expression

Constants, functions, or SQL expressions.

 NOTE

- A flashback query cannot span statements that affect the table structure or physical storage. Otherwise, an error is reported. Between the flashback point and the current point, if a statement (TRUNCATE, DDL, DCL, or VACUUM FULL) has been executed to modify the table structure or affect physical storage, the flashback fails.
- Flashback query does not support index query. Flashback query supports only seq scan for full table scanning.
- When the flashback point is too old, the source version cannot be obtained because the flashback version is recycled. As a result, the flashback fails and the error message "Restore point too old" is displayed.
- The flashback point is specified by time. The maximum difference between the flashback point and the actual time is 3 seconds.
- After truncating a table, perform a flashback query or flashback on the table. The error message "Snapshot too old" is displayed when a flashback is performed at a specified time point. Data cannot be found or the error message "Snapshot too old" is reported during the CSN-based flashback.

- column\_alias

Specifies the column alias.

- PARTITION

Queries data in the specified partition in a partitioned table.

- partition\_name

Specifies the name of a partition.

- partition\_value

Specifies the value of the specified partition key. If a partitioned table has multiple partition keys, you can use PARTITION FOR to uniquely identify a partition.

- SUBPARTITION

Queries data in the specified level-2 partition in a partitioned table.

- subpartition\_name

Specifies the name of a level-2 partition name.

- subpartition\_value

Specifies the key values of specified level-1 and level-2 partitions. The values of the two partition keys specified by the SUBPARTITION FOR clause can be used to uniquely identify a level-2 partition.

- subquery

Performs a subquery in the FROM clause. A temporary table is created to save subquery results.

- with\_query\_name  
References a WITH clause as the source of a FROM clause by the name of a WITH query.
  - function\_name  
Function name. Functions can be used in a FROM clause.
  - join\_type  
The options are as follows:
    - [ INNER ] JOIN  
A JOIN clause combines two FROM items. You can use parentheses to determine the order of nesting. In the absence of parentheses, JOIN nests left-to-right.
    - LEFT [ OUTER ] JOIN  
Returns all rows that meet the join conditions in the Cartesian product, as well as the rows in the left table that do not match the right table rows according to join conditions. This left-hand row is extended to the full width of the joined table by inserting **NULL** values for the right-hand columns. Note that the condition of the JOIN clause is used only when matches are being calculated. The condition of the outer layer is applied after the calculation.
    - RIGHT [ OUTER ] JOIN  
Returns all result rows of INNER JOIN, as well as all rows on the right that do not have matches (rows on the left are filled with **NULL**).  
This is just a notational convenience because you can always convert it to a LEFT OUTER JOIN by switching the left and right inputs.
    - FULL [ OUTER ] JOIN  
Returns all the joined rows, pluses one row for each unmatched left-hand row (extended with **NULL** on the right), and pluses one row for each unmatched right-hand row (extended with **NULL** on the left).
    - CROSS JOIN  
Is equivalent to **INNER JOIN ON (TRUE)**, which means no rows are removed by qualification. These join types are just a notational convenience, since they do nothing you could not do with plain FROM and WHERE.
-  **NOTE**
- For the INNER and OUTER join types, a join condition must be specified, namely exactly one of **NATURAL ON**, **join\_condition**, or **USING (join\_column [, ...])**. For CROSS JOIN, none of these clauses can appear.
- CROSS JOIN and INNER JOIN produce a simple Cartesian product, the same result as you get from listing the two items at the top level of FROM.
- ON join\_condition  
Defines which rows have matches in joins. Example: **ON left\_table.a = right\_table.a**. You are advised not to use numeric types such as int for

**join\_condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and **0** is implicitly converted to **false**), which may cause unexpected results.

- USING(join\_column[, ...])

Abbreviation of **ON left\_table.a = right\_table.a AND left\_table.b = right\_table.b** .... The names of the corresponding columns must be the same.

- NATURAL

Shorthand for a USING list that contains all columns in the two tables that have the same names.

- from item

Specifies the name of the query source object connected.

- WHERE clause

The WHERE clause forms an expression for row selection to narrow down the query range of SELECT. **condition** indicates any expression that returns a value of Boolean type. Rows that do not meet this condition will not be retrieved. You are advised not to use numeric types such as int for **condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and **0** is implicitly converted to **false**), which may cause unexpected results.

In a WHERE clause, you can use the operator (+) to convert a table join to an outer join. However, this method is not recommended because it is not the standard SQL syntax and may raise syntax compatibility issues during platform migration. There are many restrictions on using the operator (+):

- a. It can appear only in the WHERE clause.
- b. If a table join has been specified in the FROM clause, the operator (+) cannot be used in the WHERE clause.
- c. The operator (+) can work only on columns of tables or views, instead of on expressions.
- d. If table A and table B have multiple join conditions, the operator (+) must be specified in all the conditions. Otherwise, the operator (+) will not take effect, and the table join will be converted into an inner join without any prompt information.
- e. Tables specified in a join condition where the operator (+) works cannot cross queries or subqueries. If tables where the operator (+) works are not in the FROM clause of the current query or subquery, an error will be reported. If a peer table for the operator (+) does not exist, no error will be reported and the table join will be converted into an inner join. Expressions where the operator (+) is used cannot be directly connected through OR.
- f. If a column where the operator (+) works is compared with a constant, the expression becomes a part of the join condition.
- g. A table cannot have multiple foreign tables.
- h. The operator (+) can appear only in the following expressions: comparison, NOT, ANY, ALL, IN, NULLIF, IS DISTINCT FROM, and IS OF. It is not allowed in other types of expressions. In addition, these expressions cannot be connected through AND or OR.

- i. The operator (+) can be used to convert a table join only to a left or right outer join, instead of a full join. That is, the operator (+) cannot be specified on both tables of an expression.

**NOTICE**

- For the WHERE clause, if special character "%", "\_", or "\" is queried in LIKE, add the slash "\" before each character.
- For a hierarchical query, the WHERE expression is processed as follows in a multi-table join query:
  - Decompose the WHERE expression based on the disjunctive and conjunctive actions and check whether each subexpression involves multiple tables (potential join conditions) queried at the current layer. If a subexpression is not a subconnection and involves only multiple tables queried at the current layer, push it down to the non-recursive START WITH clause or the recursive CONNECT BY clause of the hierarchical query for prior execution as the JOIN condition. Otherwise, place it after the recursive CONNECT BY clause for execution as a filter condition.
  - During internal implementation, the subexpressions that do not meet pushdown requirements in the parse tree of the original WHERE expression are deleted to obtain the final expression that can be pushed down. Then, the subexpressions that meet the pushdown condition are removed from the parse tree of the original WHERE expression to obtain the final expression that is not pushed down.


To sum up, pay attention to the following points:

- When the WHERE expression is split, the disjunctive action is also split. Therefore, the two expressions joined by OR are also split.
- Sublinks must not be pushed down.
- If an input column is not queried at the current layer, the column will not be pushed down.


Then, for the following example:

```
SELECT * FROM t1,t2
WHERE t1.a=1 or t1.a=t2.a and t1.b=t2.b or t1.c in (SELECT t3.c FROM t3)
START WITH t1.c=1
CONNECT BY PRIOR t2.c=t1.c;
```

The following figure shows the pushdown conditions after the WHERE expression is split disjunctively and conjunctively.



It is equivalent to "t1.a=t2.a AND t1.b=t2.b". The following figure shows the conditions that are not pushed down.



It is equivalent to "t1.a=1 OR t1.c in (SELECT t3.c FROM t3)".

For another example:

```
SELECT * FROM t3 WHERE EXISTS(
 SELECT * FROM t1,t2
 WHERE t1.a+t2.a=t3.a
 START WITH t1.c=1
 CONNECT BY PRIOR t2.c=t1.c;
)
```

The WHERE condition contains **t3.a** inputted by the outer query. Therefore, the WHERE condition is not pushed down.

- **START WITH clause**

The START WITH clause is usually used together with the CONNECT BY clause and indicates the initial condition of recursion: Data is traversed recursively and hierarchically. If this clause is omitted and the CONNECT BY clause is used alone, all rows in the table are used as the initial set. For details, see • [CONNECT BY clause](#).

- **CONNECT BY clause**

CONNECT BY indicates a recursive join condition. It is used together with START WITH to implement data traversal and recursion. Example:

```
openGauss=# CREATE TABLE test(name varchar, id int, fatherid int);
openGauss=# INSERT INTO test VALUES('A', 1, 0), ('B', 2, 1),('C',3,1),('D',4,1),('E',5,2);
openGauss=# SELECT * FROM test START WITH id = 1 CONNECT BY prior id = fatherid ORDER
SIBLINGS BY id desc;
name | id | fatherid
-----+-----+-----
A | 1 | 0
D | 4 | 1
C | 3 | 1
B | 2 | 1
E | 5 | 2
(5 rows)
```

In the CONNECT BY condition, the PRIOR keyword can be specified for a column to indicate that the column is recursive. If NOCYCLE is added before the recursive join condition, recursion stops when a loop record is encountered. (Note: A SELECT statement containing the START WITH .. CONNECT BY clause does not support the FOR SHARE or UPDATE lock.)

The process of executing the **START WITH** statement is as follows:

- The initial dataset is selected based on the condition in the START WITH clause. In the preceding example, **('A', 1, 0)** is selected first. Then, this initial dataset is set as the working set.
- If the working set is not empty, the data in the working set is used as the input for the next query. The filter criteria are specified in the CONNECT

BY clause. The keyword PRIOR indicates the current record. For example, **PRIOR id = fatherid** in the preceding example indicates that the ID of the current record is **fatherid** of the next record.

- c. Set the dataset filtered in step 2 as the working set and repeat the operation in step 2.

In addition, the database adds the following pseudocolumns to each selected data record so that users can learn about the location of the data in the recursive or tree structure.

- **LEVEL**: node level.
- **CONNECT\_BY\_ISLEAF**: specifies whether a node is a leaf node.

NOTICE

**connect\_by\_isleaf**: If a node does not have any subnode, the node is a leaf node and **connect\_by\_isleaf** is set to 1. Otherwise, **connect\_by\_isleaf** is set to 0.

Assume that table T1 exists and the data in table T1 is shown in [Figure 7-12](#).

Figure 7-12 Data structure


| T1 |   |
|----|---|
| A  | B |
| 1  | 2 |
| 2  | 3 |
| 2  | 4 |
| 3  | 5 |

Run the following statement:

```
SELECT * FROM T1 CONNECT BY PRIOR B=A AND LEVEL<=3;
```

The tree structure of the result is shown in [Figure 7-13](#), where the blue nodes are leaf nodes.

Figure 7-13 Logical structure of the execution





In addition to pseudocolumns, the following query functions are provided (for details, see [Hierarchical Recursion Query Functions](#)):

- **sys\_connect\_by\_path(col, separator)**: returns the connection path from the root node to the current row. The **col** parameter indicates the name of the column displayed in the path, and the **separator** parameter indicates the connector.
- **connect\_by\_root(col)**: displays the top-level node of the node. **col** indicates the name of the output column.

If a loop exists in the dataset, the database provides loop detection. By default, if a loop is detected, an error is reported and no data is returned. In addition, the NOCYCLE keyword is provided. With it, the query can be executed normally and when the first duplicate data record is found, the query exits directly instead of reporting an error.

Besides, in the hierarchical query, the search is performed strictly according to the depth-first order. If ROWNUM is used as the filter condition in START WITH or CONNECT BY, the value of **ROWNUM** is increased by 1 for each record to be returned. Then, the record is verified based on **ROWNUM** conditions. Records that do not satisfy the conditions are discarded and the value of **ROWNUM** is decreased by 1.

#### NOTICE

- The PRIOR keyword can be used only in the CONNECT BY clause instead of the START WITH clause.
  - The PRIOR keyword can be specified only for columns in the table instead of expressions, pseudocolumns, or type conversion. For example, PRIOR (a + 1) is not allowed.
  - In the CONNECT BY clause, the column using the keyword PRIOR cannot be in the same condition with pseudocolumns such as level and rownum. For example, (PRIOR a = level) is not allowed, but (PRIOR a = b) AND (level = 1) is allowed. Different conditions refer to the conditions connected by AND at the top of the CONNECT BY clause. For example, (PRIOR a = 1 OR level = 1) is considered as a condition and is not allowed.
  - In the START WITH and CONNECT BY clauses, pseudocolumns cannot be used for sublinks, for example, "rownum = (subquery)" or "rownum IN (subquery)".
  - You are advised not to use ROWNUM in the CONNECT BY clause. If ROWNUM is required, test it thoroughly to avoid unexpected results.
  - When START WITH or CONNECT BY is called on the CTE defined by WITH AS, if there are multiple CTEs, ensure that the definition of each CTE does not depend on other CTEs.
  - If no loop exists in the data but the error message "runs into cycle" is reported, increase the value of **max\_recursive\_times**.
  - **connect\_by\_isleaf**, **connect\_by\_iscycle**, and **level** are of the int type.
  - **connect\_by\_isleaf**, **connect\_by\_iscycle**, and **level** are parsed as pseudocolumns in hierarchical queries.
  - The alias defined in the projection column cannot be called for START WITH or CONNECT BY.
  - Optimization suggestions for START WITH:
    - Create indexes based on the conditions in the CONNECT BY clause to improve the performance of the START WITH clause.
    - Identify bottlenecks based on the plan collected by running EXPLAIN PERFORMANCE or in WDRs. If the recursive operator (inner plan) of RECURSIVE UNION is the HASH JOIN operator, but the hash table is created for the temporary table **tmp\_result** or the hash table in plan is materialized (that is, the batch size is greater than 1), the possible cause is that the value of **work\_mem** is too small. As a result, the hash table cannot be created for the outer data table. You can increase the value of **work\_mem** to improve performance.
- Note: GaussDB optimizes tables with a small volume of data and caches table results in hash tables to improve performance. In this case, indexes are not required. However, if the data volume exceeds the limit specified by **work\_mem**, the optimization becomes invalid. In this case, you can create indexes for optimization.

 **WARNING**

You are advised not to use recursive query when creating a view definition. Otherwise, the obtained view definition is incorrect or an error is reported. This is because the bottom layer of the recursive query statement is rewritten as the RECURSIVE UNION statement, and the displayed view definition is obtained by reversely parsing the RECURSIVE UNION query tree. After a view with recursive query is created, the view definition is the RECURSIVE UNION statement, instead of the START WITH statement, or even the statement may not be valid. The current database view functions are not affected. However, the restore operation will fail because some tools, such as `gs_dump` and `gs_restore`, migrate views by checking the view definition. In this case, you need to manually create the view again.

- **ORDER SIBLINGS BY clause**

The output of the START WITH statement is returned level by level. However, there is no sequence guarantee at each level because the database automatically selects the optimal execution path during each round of query. In the preceding example, A is output first, but the sequence of B, C, and D is not fixed. If you have requirements on the final output sequence, you can use ORDER SIBLINGS BY. The usage of ORDER SIBLINGS BY is the same as that of ORDER BY. ORDER SIBLINGS BY is used for sorting at each level during recursion.

**NOTICE**

The expression following ORDER SIBLINGS BY only supports sorting by calling non-aggregate and window functions for ordinary columns, column name offsets, and column names. It does not support calling system functions related to START WITH for column names or using pseudocolumns related to START WITH.

- **GROUP BY clause**

Condenses query results into a single row all selected rows that share the same values for the grouped expressions.

- **CUBE ( { expression | ( expression [, ...] ) } [, ...] )**

A CUBE grouping is an extension to the GROUP BY clause that creates subtotals for all of the possible combinations of the given list of grouping columns (or expressions). In terms of multidimensional analysis, CUBE generates all the subtotals that could be calculated for a data cube with the specified dimensions. For example, given three expressions ( $n=3$ ) in the CUBE clause, the operation results in  $2^n = 2^3 = 8$  groupings. Rows grouped on the values of  $n$  expressions are called regular rows, and the rest are called super-aggregate rows.

- **GROUPING SETS ( grouping\_element [, ...] )**

GROUPING SETS is another extension to the GROUP BY clause. It allows you to specify multiple GROUP BY clauses. This improves efficiency by trimming away unnecessary data. After you specify the required data group, the database does not need to compute a whole **ROLLUP** or **CUBE**.

**NOTICE**

- If a SELECT list expression references ungrouped columns and no aggregate function is used, an error is reported. This is because multiple values may be returned for ungrouped columns.
- If a SELECT list expression references a constant, the GROUP BY clause does not need to group the constant. Otherwise, an error is reported.

- **HAVING clause**

Used together with the GROUP BY clause to select special groups. The HAVING clause compares some attributes of groups with a constant. Only groups that match the logical expression in the HAVING clause are extracted.

- **WINDOW clause**

The general format is **WINDOW window\_name AS ( window\_definition ) [, ...]**. **window\_name** is a name can be referenced by **window\_definition**. **window\_definition** can be expressed in the following forms:

```
[existing_window_name]
[PARTITION BY expression [, ...]]
[ORDER BY expression [ASC | DESC | USING operator] [NULLS { FIRST | LAST }] [, ...]]
[frame_clause]
```

**frame\_clause** defines a **window frame** for the window function. The window function (not all window functions) depends on **window frame** and **window frame** is a set of relevant rows of the current query row. **frame\_clause** can be expressed in the following forms:

```
[RANGE | ROWS] frame_start
[RANGE | ROWS] BETWEEN frame_start AND frame_end
```

**frame\_start** and **frame\_end** can be expressed in the following forms:

```
UNBOUNDED PRECEDING
value PRECEDING
CURRENT ROW
value FOLLOWING
UNBOUNDED FOLLOWING
```

- **UNION clause**

Computes the set union of the rows returned by multiple SELECT statements. The UNION clause has the following constraints:

- The default result of UNION does not contain any duplicate rows unless the ALL clause is specified.
- Multiple UNION operators in the same SELECT statement are evaluated from left to right, unless otherwise specified by parentheses.
- **FOR UPDATE**, **FOR NO KEY UPDATE**, **FOR SHARE**, and **FOR KEY SHARE** cannot be specified in the result or input of **UNION**.

General expression:

```
select_statement UNION [ALL] select_statement
```

- **select\_statement** can be any SELECT statement without the ORDER BY, LIMIT, FOR UPDATE, FOR NO KEY UPDATE, FOR SHARE, or FOR KEY SHARE clause.
- ORDER BY and LIMIT can be attached to the subexpression if it is enclosed in parentheses.

- **INTERSECT clause**

Computes the set intersection of rows returned by the involved SELECT statements. The result of INTERSECT does not contain any duplicate rows.

The INTERSECT clause has the following constraints:

- Multiple INTERSECT operators in the same SELECT statement are evaluated from left to right, unless otherwise specified by parentheses.
- Processing INTERSECT preferentially when UNION and INTERSECT operations are executed for results of multiple SELECT statements.

General format:

```
select_statement INTERSECT select_statement
```

**select\_statement** can be any SELECT statement without the FOR UPDATE, FOR NO KEY UPDATE, FOR SHARE, or FOR KEY SHARE clause.

- EXCEPT clause

Has the following common form:

```
select_statement EXCEPT [ALL] select_statement
```

**select\_statement** can be any SELECT statement without the FOR UPDATE, FOR NO KEY UPDATE, FOR SHARE, or FOR KEY SHARE clause.

The EXCEPT operator computes the set of rows that are in the result of the left SELECT statement but not in the result of the right one.

The result of EXCEPT does not contain any duplicate rows unless the ALL clause is declared. To execute **ALL**, a row that has  $m$  duplicates in the left table and  $n$  duplicates in the right table will appear  $\text{MAX}(m-n, 0)$  times in the result set.

Multiple EXCEPT operators in the same SELECT statement are evaluated from left to right, unless parentheses dictate otherwise. EXCEPT binds at the same level as UNION.

Currently, the FOR UPDATE, FOR NO KEY UPDATE, FOR SHARE, and FOR KEY SHARE clauses cannot be specified for the result of EXCEPT or any input of EXCEPT.

- MINUS clause

Has the same function and syntax as EXCEPT clause.

- ORDER BY clause

Sorts data retrieved by SELECT in descending or ascending order. If the ORDER BY expression contains multiple columns:

- If two columns are equal according to the leftmost expression, they are compared according to the next expression.
- If they are equal according to all specified expressions, they are returned in an implementation-dependent order.
- When used with the DISTINCT keyword, the columns to be sorted in ORDER BY must be included in the columns of the result set retrieved by the SELECT statement.
- When used with the GROUP BY clause, the columns to be sorted in ORDER BY must be included in the columns of the result set retrieved by the SELECT statement.

**NOTICE**

To support Chinese pinyin order, specify the **UTF-8**, **GB18030**, or **GBK** encoding mode during database initiation. The statements are as follows:

```
initdb -E UTF8 -D ../data -locale=zh_CN.UTF-8, initdb -E GB18030 -D ../data -locale=zh_CN.GB18030,
or initdb -E GBK -D ../data -locale=zh_CN.GBK.
```

- **LIMIT clause**

Consists of two independent sub-clauses:

**LIMIT { count | ALL }** limits the number of rows to be returned. **count** specifies the number of rows to be returned. The effect of **LIMIT ALL** is the same as that of omitting the **LIMIT** clause.

**OFFSET start count** specifies the maximum number of rows to return, while **start** specifies the number of rows to skip before starting to return rows. When both are specified, **start** rows are skipped before starting to count the **count** rows to be returned.

- **OFFSET clause**

The SQL: 2008 standard has introduced a different clause:

**OFFSET start { ROW | ROWS }**

**start** specifies the number of rows to skip before starting to return rows.

- **FETCH { FIRST | NEXT } [ count ] { ROW | ROWS } ONLY**

If **count** is omitted in a **FETCH** clause, it defaults to **1**.

- **Locking clause**

The **FOR UPDATE** clause locks the rows retrieved by **SELECT**. This prevents these rows from being modified or deleted by other transactions before the current transaction ends. That is, other transactions that attempt to run **UPDATE**, **DELETE**, **SELECT FOR UPDATE**, **SELECT FOR NO KEY UPDATE**, **SELECT FOR SHARE**, or **SELECT FOR KEY SHARE** for these rows will be blocked until the current transaction ends. Any **DELETE** on the same row will also acquire the **FOR UPDATE** lock mode, as will **UPDATE** that modifies values on the primary key column. Conversely, **SELECT FOR UPDATE** will wait for concurrent transactions that are running the preceding commands on the same row, and will continue to lock and return the updated rows (there may be no row because the rows may be deleted).

**FOR NO KEY UPDATE** behaves similarly to **FOR UPDATE**, except that it acquires a weaker lock that will not block **SELECT FOR KEY SHARE** that attempts to acquire the lock on the same row. Any **UPDATE** that does not acquire the **FOR UPDATE** lock will also acquire this locking mode.

**FOR SHARE** behaves similarly, except that it requires a shared lock rather than an exclusive lock on each retrieved row. A shared lock blocks other transactions from executing **UPDATE**, **DELETE**, **SELECT FOR UPDATE**, or **SELECT FOR NO KEY UPDATE**, except **SELECT FOR SHARE** and **SELECT FOR KEY SHARE**.

**FOR KEY SHARE** is similar to **FOR SHARE** except that its lock is weak. **SELECT FOR UPDATE** is blocked but **SELECT FOR NO KEY UPDATE** is not blocked. A key-shared lock blocks other transactions from executing **DELETE** or **UPDATE** that modifies the key value, except **UPDATE**, **SELECT FOR NO KEY UPDATE**, **SELECT FOR SHARE**, and **SELECT FOR KEY SHARE**.

To prevent the operation from waiting for the commit of other transactions, you can use **NOWAIT**. If the selected row cannot be locked immediately, an error is reported immediately. If you use **WAIT** *n* and the selected row cannot be locked immediately, the operation needs to wait for *n* seconds (the value of *n* is of the int type with a range of  $0 \leq n \leq 2147483$ ). If the lock is obtained within *n* seconds, the operation is performed normally. Otherwise, an error is reported.

If table names are specified in a locking clause, then only those tables will be locked. Other tables used in SELECT will not be locked. Otherwise, locking all tables in the statement.

If a locking clause is applied to a view or sub-query, it affects all tables used in the view or sub-query.

Multiple locking clauses can be written if it is necessary to specify different locking behaviors for different tables.

If a table appears (or implicitly appears) in multiple clauses at the same time, the strongest lock is used. Similarly, a table is processed as **NOWAIT** if that is specified in any of the clauses affecting it.

---

**NOTICE**

Only FOR SHARE and FOR UPDATE can be used to query the Ustore table.

---

- **NLS\_SORT**

Specifies that a column is sorted in a special order. Currently, only Chinese Pinyin and case-insensitive sorting are supported. To support this sorting mode, you need to set the encoding format to UTF8, GB18030, or GBK when creating a database. If you set the encoding format to another format, for example, SQL\_ASCII, an error may be reported or the sorting mode may be invalid.

Value range:

- **SCHINESE\_PINYIN\_M**, sorted by Pinyin order.
- **generic\_m\_ci**: sorted in case-insensitive order (optional; only English characters are supported in the case-insensitive order.)

- **PARTITION** clause

Queries data in the specified partition in a partitioned table.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason values(3,'AAAAAAAAABAAAAAAA','reason 1'),
(10,'AAAAAAAAABAAAAAAA','reason 2'),(4,'AAAAAAAAABAAAAAAA','reason 3'),
(10,'AAAAAAAAABAAAAAAA','reason 4'),(10,'AAAAAAAAABAAAAAAA','reason 5'),
```

```
(20,'AAAAAAACAAAAAA','N%reason 6'),(30,'AAAAAAACAAAAAA','W%reason 7');

-- Obtain the temp_t temporary table by a subquery and query all records in this table.
openGauss=# WITH temp_t(name,isdba) AS (SELECT username,usesuper FROM pg_user) SELECT * FROM
temp_t;

-- Query all r_reason_sk records in the tpcds.reason table and delete duplicate records.
openGauss=# SELECT DISTINCT(r_reason_sk) FROM tpcds.reason;

-- Example of a LIMIT clause: Obtain a record from the table.
openGauss=# SELECT * FROM tpcds.reason LIMIT 1;

-- Query all records and sort them in alphabetic order.
openGauss=# SELECT r_reason_desc FROM tpcds.reason ORDER BY r_reason_desc;

-- Use table aliases to obtain data from pg_user and pg_user_status:
openGauss=# SELECT a.username,b.locktime FROM pg_user a,pg_user_status b WHERE a.usesysid=b.roloid;

-- Example of the FULL JOIN clause: Join data in pg_user and pg_user_status.
openGauss=# SELECT a.username,b.locktime,a.usesuper FROM pg_user a FULL JOIN pg_user_status b on
a.usesysid=b.roloid;

-- Example of the GROUP BY clause: Filter data based on query conditions, and group the results.
openGauss=# SELECT r_reason_id, AVG(r_reason_sk) FROM tpcds.reason GROUP BY r_reason_id HAVING
AVG(r_reason_sk) > 25;

-- Example of the GROUP BY CUBE clause: Filter data based on query conditions, and group the results.
openGauss=# SELECT r_reason_id,AVG(r_reason_sk) FROM tpcds.reason GROUP BY
CUBE(r_reason_id,r_reason_sk);

-- Example of the GROUP BY GROUPING SETS clause: Filter data based on query conditions, and group the
results.
openGauss=# SELECT r_reason_id,AVG(r_reason_sk) FROM tpcds.reason GROUP BY GROUPING
SETS((r_reason_id,r_reason_sk),r_reason_sk);

-- Example of the UNION clause: Merge the names started with W and N in the r_reason_desc column in
the tpcds.reason table.
openGauss=# SELECT r_reason_sk, tpcds.reason.r_reason_desc
FROM tpcds.reason
WHERE tpcds.reason.r_reason_desc LIKE 'W%'
UNION
SELECT r_reason_sk, tpcds.reason.r_reason_desc
FROM tpcds.reason
WHERE tpcds.reason.r_reason_desc LIKE 'N%';

-- Example of the NLS_SORT clause: Sort by Chinese Pinyin.
openGauss=# SELECT * FROM tpcds.reason ORDER BY NLSSORT(r_reason_desc, 'NLS_SORT =
SCHINESE_PINYIN_M');

-- sorting in case-insensitive order (optional; only English characters are supported in the case-insensitive
order.)
openGauss=# SELECT * FROM tpcds.reason ORDER BY NLSSORT(r_reason_desc, 'NLS_SORT =
generic_m_ci');

-- Create a range-partitioned table tpcds.reason_p.
openGauss=# CREATE TABLE tpcds.reason_p
(
r_reason_sk integer,
r_reason_id character(16),
r_reason_desc character(100)
)
PARTITION BY RANGE (r_reason_sk)
(
partition P_05_BEFORE values less than (05),
partition P_15 values less than (15),
partition P_25 values less than (25),
partition P_35 values less than (35),
partition P_45_AFTER values less than (MAXVALUE)
```



```
)
;

-- Insert data.
openGauss=# INSERT INTO tpcds.reason_p values(3,'AAAAAAAABAAAAAAA','reason 1'),
(10,'AAAAAAAABAAAAAAA','reason 2'),(4,'AAAAAAAABAAAAAAA','reason 3'),
(10,'AAAAAAAABAAAAAAA','reason 4'),(10,'AAAAAAAABAAAAAAA','reason 5'),
(20,'AAAAAAACAAAAAAA','reason 6'),(30,'AAAAAAACAAAAAAA','reason 7');

-- Example of the PARTITION clause: Obtain data from the P_05_BEFORE partition in the tpcds.reason_p
table.
openGauss=# SELECT * FROM tpcds.reason_p PARTITION (P_05_BEFORE);
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----
 4 | AAAAAAABAAAAAAA | reason 3
 3 | AAAAAAABAAAAAAA | reason 1
(2 rows)

-- Example of the GROUP BY clause: Group records in the tpcds.reason_p table by r_reason_id, and count
the number of records in each group.
openGauss=# SELECT COUNT(*),r_reason_id FROM tpcds.reason_p GROUP BY r_reason_id;
 count | r_reason_id
-----+-----
 2 | AAAAAAACAAAAAAA
 5 | AAAAAAABAAAAAAA
(2 rows)

-- Example of the GROUP BY CUBE clause: Filter data based on query conditions, and group the results.
openGauss=# SELECT * FROM tpcds.reason GROUP BY CUBE (r_reason_id,r_reason_sk,r_reason_desc);

-- Example of the GROUP BY GROUPING SETS clause: Filter data based on query conditions, and group the
results.
openGauss=# SELECT * FROM tpcds.reason GROUP BY GROUPING SETS
((r_reason_id,r_reason_sk),r_reason_desc);

-- Example of the HAVING clause: Group records in the tpcds.reason_p table by r_reason_id, count the
number of records in each group, and display only values whose number of r_reason_id is greater than 2.
openGauss=# SELECT COUNT(*) c,r_reason_id FROM tpcds.reason_p GROUP BY r_reason_id HAVING c>2;
 c | r_reason_id
---+-----
 5 | AAAAAAABAAAAAAA
(1 row)

-- Example of the IN clause: Group records in the tpcds.reason_p table by r_reason_id, count the number
of records in each group, and display only the numbers of records whose r_reason_id is
AAAAAAABAAAAAAA or AAAAAADAAAAAAA.
openGauss=# SELECT COUNT(*),r_reason_id FROM tpcds.reason_p GROUP BY r_reason_id HAVING
r_reason_id IN('AAAAAABAAAAAAA','AAAAAADAAAAAAA');
 count | r_reason_id
-----+-----
 5 | AAAAAAABAAAAAAA
(1 row)

-- Example of the INTERSECT clause: Query records whose r_reason_id is AAAAAAABAAAAAAA and
whose r_reason_sk is smaller than 5.
openGauss=# SELECT * FROM tpcds.reason_p WHERE r_reason_id='AAAAAABAAAAAAA' INTERSECT
SELECT * FROM tpcds.reason_p WHERE r_reason_sk<5;
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----
 4 | AAAAAAABAAAAAAA | reason 3
 3 | AAAAAAABAAAAAAA | reason 1
(2 rows)

-- Example of the EXCEPT clause: Query records whose r_reason_id is AAAAAAABAAAAAAA and whose
r_reason_sk is greater than or equal to 4.
openGauss=# SELECT * FROM tpcds.reason_p WHERE r_reason_id='AAAAAABAAAAAAA' EXCEPT SELECT *
FROM tpcds.reason_p WHERE r_reason_sk<4;
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----
```

```
10 | AAAAAAAAAABAAAAAAA | reason 2
10 | AAAAAAAAAABAAAAAAA | reason 5
10 | AAAAAAAAAABAAAAAAA | reason 4
4 | AAAAAAAAAABAAAAAAA | reason 3
(4 rows)

-- Create the store_returns and customer tables.
openGauss=# CREATE TABLE tpcds.store_returns (sr_item_sk int, sr_customer_id varchar(50),sr_customer_sk
int);
openGauss=# CREATE TABLE tpcds.customer (c_item_sk int, c_customer_id varchar(50),c_customer_sk int);
openGauss=# INSERT INTO tpcds.store_returns VALUES(18000);
openGauss=# INSERT INTO tpcds.customer (c_customer_id) VALUES('AAAAAAAJNGEBAAA');
-- Specify the operator (+) in the WHERE clause to indicate a left join.
openGauss=#
SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2 WHERE
t1.sr_customer_sk = t2.c_customer_sk(+) ORDER BY 1 DESC LIMIT 1;
 sr_item_sk | c_customer_id
-----+-----
 18000 |
(1 row)

-- Specify the operator (+) in the WHERE clause to indicate a right join.
openGauss=#
SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2 WHERE
t1.sr_customer_sk(+) = t2.c_customer_sk
ORDER BY 1 DESC LIMIT 1;
 sr_item_sk | c_customer_id
-----+-----
 | AAAAAAAJNGEBAAA
(1 row)

-- Specify the operator (+) in the WHERE clause to indicate a left join and add a join condition.
openGauss=#
SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2 WHERE
t1.sr_customer_sk = t2.c_customer_sk(+) AND t2.c_customer_sk(+) < 1 ORDER BY 1 LIMIT 1;
 sr_item_sk | c_customer_id
-----+-----
 18000 |
(1 row)

-- If the operator (+) is specified in the WHERE clause, do not use expressions connected through AND or
OR.
openGauss=# SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2
WHERE NOT(t1.sr_customer_sk = t2.c_customer_sk(+) AND t2.c_customer_sk(+) < 1);
ERROR: Operator "(+)" can not be used in nesting expression.
LINE 1: ...tomer_id from store_returns t1, customer t2 where not(t1.sr_...
 ^

-- If the operator (+) is specified in the WHERE clause which does not support expression macros, an error
will be reported.
openGauss=# SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2
WHERE (t1.sr_customer_sk = t2.c_customer_sk(+))::bool;
ERROR: Operator "(+)" can only be used in common expression.

-- If the operator (+) is specified on both sides of an expression in the WHERE clause, an error will be
reported.
openGauss=# SELECT t1.sr_item_sk ,t2.c_customer_id FROM tpcds.store_returns t1, tpcds.customer t2
WHERE t1.sr_customer_sk(+) = t2.c_customer_sk(+);
ERROR: Operator "(+)" can't be specified on more than one relation in one join condition
HINT: "t1", "t2"...are specified Operator "(+)" in one condition.

-- Drop the table.
openGauss=# DROP TABLE tpcds.reason_p;

-- Flashback query example. To use the flashback function, you need to set the undo_retention_time
parameter.
-- Create the tpcds.time_table table.
openGauss=# CREATE TABLE tpcds.time_table(idx integer, snaptime timestamp, snapcsn bigint, timeDesc
character(100));
-- Insert records into the tpcds.time_table table.
```

```
openGauss=# INSERT INTO tpcds.time_table SELECT 1, now(),int8in(xidout(next_csn)), 'time1' FROM
gs_get_next_xid_csn();
openGauss=# INSERT INTO tpcds.time_table SELECT 2, now(),int8in(xidout(next_csn)), 'time2' FROM
gs_get_next_xid_csn();
openGauss=# INSERT INTO tpcds.time_table SELECT 3, now(),int8in(xidout(next_csn)), 'time3' FROM
gs_get_next_xid_csn();
openGauss=# INSERT INTO tpcds.time_table SELECT 4, now(),int8in(xidout(next_csn)), 'time4' FROM
gs_get_next_xid_csn();
openGauss=# SELECT * FROM tpcds.time_table;
```

| idx | snaptime                   | snapsn | timedesc |
|-----|----------------------------|--------|----------|
| 1   | 2021-04-25 17:50:05.360326 | 107322 | time1    |
| 2   | 2021-04-25 17:50:10.886848 | 107324 | time2    |
| 3   | 2021-04-25 17:50:16.12921  | 107327 | time3    |
| 4   | 2021-04-25 17:50:22.311176 | 107330 | time4    |

(4 rows)

```
openGauss=# DELETE tpcds.time_table;
DELETE 4
-- 2021-04-25 17:50:22.311176: The value of the fourth snaptime column in tpcds.time_table is
recommended.
openGauss=# SELECT * FROM tpcds.time_table TIMECAPSULE TIMESTAMP to_timestamp('2021-04-25
17:50:22.311176','YYYY-MM-DD HH24:MI:SS.FF');
```

| idx | snaptime                   | snapsn | timedesc |
|-----|----------------------------|--------|----------|
| 1   | 2021-04-25 17:50:05.360326 | 107322 | time1    |
| 2   | 2021-04-25 17:50:10.886848 | 107324 | time2    |
| 3   | 2021-04-25 17:50:16.12921  | 107327 | time3    |

(3 rows)

```
-- 107330 csn: The value of the fourth snapsn column in tpcds.time_table is recommended.
openGauss=# SELECT * FROM tpcds.time_table TIMECAPSULE CSN 107330;
```

| idx | snaptime                   | snapsn | timedesc |
|-----|----------------------------|--------|----------|
| 1   | 2021-04-25 17:50:05.360326 | 107322 | time1    |
| 2   | 2021-04-25 17:50:10.886848 | 107324 | time2    |
| 3   | 2021-04-25 17:50:16.12921  | 107327 | time3    |

(3 rows)

```
-- Example of a WITH RECURSIVE query: Calculate the accumulated value from 1 to 100.
openGauss=# WITH RECURSIVE t1(a) AS (
 SELECT 100
),
t(n) AS (
 VALUES (1)
 UNION ALL
 SELECT n+1 FROM t WHERE n < (SELECT max(a) FROM t1)
)
SELECT sum(n) FROM t;
sum

5050
(1 row)
```

```
openGauss=# DROP TABLE t;
-- Drop the table.
openGauss=# DROP TABLE tpcds.reason;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.143 SELECT INTO

### Description

SELECT INTO creates a table based on the query result and inserts data retrieved by the query into the new table.

Different from SELECT, data is not returned to the client. The table columns have the same names and data types as the output columns of SELECT.

### Precautions

CREATE TABLE AS provides functions similar to SELECT INTO in functions and provides a superset of functions provided by SELECT INTO. You are advised to use CREATE TABLE AS, because SELECT INTO cannot be used in a stored procedure.

### Syntax

```
[WITH [RECURSIVE] with_query [, ...]]
SELECT [ALL | DISTINCT [ON (expression [, ...])]]
 { * | {expression [[AS] output_name]} [, ...] }
 INTO [[GLOBAL | LOCAL] [TEMPORARY | TEMP] | UNLOGGED] [TABLE] new_table
 [FROM from_item [, ...]]
 [WHERE condition]
 [GROUP BY expression [, ...]]
 [HAVING condition [, ...]]
 [WINDOW {window_name AS (window_definition)} [, ...]]
 [{ UNION | INTERSECT | EXCEPT | MINUS } [ALL | DISTINCT] select]
 [ORDER BY {expression [[ASC | DESC | USING operator] | nlssort_expression_clause] [NULLS { FIRST |
LAST }]} [, ...]]
 [LIMIT { count | ALL }]
 [OFFSET start [ROW | ROWS]]
 [FETCH { FIRST | NEXT } [count] { ROW | ROWS } ONLY]
 [{FOR {UPDATE | SHARE } [OF table_name [, ...]] [NOWAIT |WAIT N]} [, ...]];
```

### Parameters

- **new\_table**

**new\_table** specifies the name of the new table.

- **UNLOGGED**

Specifies that the table is created as an unlogged table. Data written to unlogged tables is not written to the WALs, which makes them considerably faster than ordinary tables. However, it is also insecure. After the database is restarted due to a conflict or abnormal shutdown, the data in the unlogged table is cleared. Contents of an unlogged table are also not replicated to standby nodes. Any indexes created on an unlogged table are automatically unlogged as well.

- Usage scenario: Unlogged tables do not ensure data security. Users can back up data before using unlogged tables; for example, users should back up the data before a system upgrade.
- Troubleshooting: If data is missing in the indexes of unlogged tables due to an abnormal shutdown or other unexpected operations, users should rebuild indexes with errors.

- **GLOBAL | LOCAL**

When creating a temporary table, you can specify the **GLOBAL** or **LOCAL** keyword before **TEMP** or **TEMPORARY**. If the keyword **GLOBAL** is specified,

GaussDB creates a global temporary table. Otherwise, GaussDB creates a local temporary table.

- **TEMPORARY | TEMP**

If **TEMP** or **TEMPORARY** is specified, the created table is a temporary table. Temporary tables are classified into global temporary tables and local temporary tables. If the keyword **GLOBAL** is specified when a temporary table is created, the table is a global temporary table. Otherwise, the table is a local temporary table.

The metadata of the global temporary table is visible to all sessions. After the sessions end, the metadata still exists. The user data, indexes, and statistics of a session are isolated from those of another session. Each session can only view and modify the data committed by itself. Global temporary tables can be created using **ON COMMIT PRESERVE ROWS** or **ON COMMIT DELETE ROWS**. The former creates session-level tables, where user data is automatically cleared when a session ends; the latter creates transaction-level tables, where user data is automatically cleared when a commit or rollback operation is performed. If the **ON COMMIT** option is not specified during table creation, the session level is used by default. Different from local temporary tables, you can specify a schema that does not start with **pg\_temp\_** when creating a global temporary table.

A local temporary table is automatically dropped at the end of the current session. Therefore, you can create and use temporary tables in the current session as long as the connected database node in the session is normal. Temporary tables are created only in the current session. If a DDL statement involves operations on temporary tables, a DDL error will be generated. Therefore, you are advised not to perform operations on temporary tables in DDL statements. **TEMP** is equivalent to **TEMPORARY**.

---

**NOTICE**

- Local temporary tables are visible to the current session through the schema starting with **pg\_temp**. Users should not delete schemas starting with **pg\_temp** or **pg\_toast\_temp**.
- If **TEMPORARY** or **TEMP** is not specified when you create a table but its schema is set to that starting with **pg\_temp\_** in the current session, the table will be created as a temporary table.
- If another session is using a global temporary table or index, the **ALTER** or **DROP** operation cannot be performed on it.
- The DDL of a global temporary table affects only the user data and indexes of the current session. For example, **TRUNCATE**, **REINDEX**, and **ANALYZE** are valid only for the current session.

---

 **NOTE**

For details about other parameters of **SELECT INTO**, see [Parameters](#) in "SELECT."

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpceds;
```

```
-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason values(1,'AAAAAAAABAAAAAAA','reason 1'),
(2,'AAAAAAAABAAAAAAA','reason 2'),(3,'AAAAAAAABAAAAAAA','reason 3'),
(4,'AAAAAAAABAAAAAAA','reason 4'),(4,'AAAAAAAABAAAAAAA','reason 5'),
(4,'AAAAAAAACAAAAAAA','reason 6'),(5,'AAAAAAAACAAAAAAA','reason 7');

-- Add the values that are less than 5 in the r_reason_sk column of the tpcds.reason table to the new table.
openGauss=# SELECT * INTO tpcds.reason_t1 FROM tpcds.reason WHERE r_reason_sk < 5;
INSERT 0 6

-- Drop the tpcds.reason_t1 table.
openGauss=# DROP TABLE tpcds.reason_t1;

-- Drop the tpcds.reason table.
openGauss=# DROP TABLE tpcds.reason;

-- Drop the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Helpful Links

[SELECT](#)

## 7.13.144 SET

### Description

Modifies GUC parameters.

### Precautions

Most GUC parameters can be modified at runtime by using SET, but some parameters cannot be modified after a service or session starts.

### Syntax

- Set the system time zone.  
SET [ SESSION | LOCAL ] TIME ZONE { timezone | LOCAL | DEFAULT };
- Set the schema of the table.  
SET [ SESSION | LOCAL ]  
    {CURRENT\_SCHEMA { TO | = } { schema | DEFAULT }  
    | SCHEMA 'schema'};
- Set client-side encoding.  
SET [ SESSION | LOCAL ] NAMES encoding\_name;
- Set XML parsing mode.  
SET [ SESSION | LOCAL ] XML OPTION { DOCUMENT | CONTENT };
- Set other GUC parameters.  
SET [ LOCAL | SESSION ]  
    {config\_parameter { { TO | = } { value | DEFAULT }  
    | FROM CURRENT }};

## Parameters

- **SESSION**

Specifies that the specified parameters take effect for the current session. If neither **SESSION** nor **LOCAL** appears, **SESSION** is the default value.

If this command is executed in a transaction, the effect of the command disappears after the transaction is rolled back. Once the surrounding transaction is committed, the effect will persist until the end of the session, unless the parameters are reset by another SET statement.

- **LOCAL**

Specifies that the specified parameters take effect for the current transaction. After COMMIT or ROLLBACK, session-level settings will take effect again.

The effect of this command lasts only until the end of the current transaction, whether the transaction is committed or not. A special case is **SET** followed by **SET LOCAL** within a single transaction: the **SET LOCAL** value will be seen until the end of the transaction, but afterward (if the transaction is committed) the **SET** value will take effect.

- **TIME\_ZONE timezone**

Specifies the local time zone for the current session.

Value range: a valid local time zone. The corresponding GUC parameter is **TimeZone**. The default value is **PRC**.

- **CURRENT\_SCHEMA schema**

Specifies the current schema.

Value range: an existing schema name. If the schema name does not exist, the value of **CURRENT\_SCHEMA** will be empty.

- **SCHEMA schema**

Specifies the current schema. Here the schema is a string.

Example: set schema 'public';

- **NAMES encoding\_name**

Specifies the client character encoding. This statement is equivalent to **set client\_encoding to encoding\_name**.

Value range: a valid character encoding name. The GUC parameter corresponding to this option is **client\_encoding**. The default encoding is **UTF8**.

- **XML OPTION option**

Specifies the XML parsing mode.

Value range: **CONTENT** (default) and **DOCUMENT**.

- **config\_parameter**

Specifies the name of a configurable GUC parameter. You can use SHOW ALL to view the available GUC parameters.

### NOTE

Some parameters viewed by SHOW ALL cannot be set by using SET. For example, **max\_datanodes**.

- **value**

Specifies the new value of **config\_parameter**. This parameter can be specified as string constants, identifiers, numbers, or comma-separated lists of these.

**DEFAULT** can be written to indicate resetting the parameter to its default value.

## Examples

```
-- Set the search path of a schema.
openGauss=# SET search_path TO tpeds, public;

-- Set the date style to the traditional POSTGRES style (date placed before month).
openGauss=# SET datestyle TO postgres,dmy;
```

## Helpful Links

[RESET](#) and [SHOW](#)

## 7.13.145 SET CONSTRAINTS

### Function

**SET CONSTRAINTS** sets the behavior of constraint checking within the current transaction.

**IMMEDIATE** constraints are checked at the end of each statement. **DEFERRED** constraints are not checked until transaction commit. Each constraint has its own **IMMEDIATE** or **DEFERRED** mode.

Upon creation, a constraint is given one of three characteristics **DEFERRABLE INITIALLY DEFERRED**, **DEFERRABLE INITIALLY IMMEDIATE**, or **NOT DEFERRABLE**. The third class is always **IMMEDIATE** and is not affected by the **SET CONSTRAINTS** statement. The first two classes start every transaction in specified modes, but its behaviors can be changed within a transaction by **SET CONSTRAINTS**.

**SET CONSTRAINTS** with a list of constraint names changes the mode of just those constraints (which must all be deferrable). If multiple constraints match a name, the name is affected by all of these constraints. **SET CONSTRAINTS ALL** changes the modes of all deferrable constraints.

When **SET CONSTRAINTS** changes the mode of a constraint from **DEFERRED** to **IMMEDIATE**, the new mode takes effect retroactively: any outstanding data modifications that would have been checked at the end of the transaction are instead checked during the execution of the **SET CONSTRAINTS** statement. If any such constraint is violated, the **SET CONSTRAINTS** fails (and does not change the constraint mode). Therefore, **SET CONSTRAINTS** can be used to force checking of constraints to occur at a specific point in a transaction.

Constraints are always checked immediately when a row is inserted or modified.

### Precautions

**SET CONSTRAINTS** sets the behavior of constraint checking only within the current transaction. Therefore, if you execute this statement outside of a transaction block (**START TRANSACTION/COMMIT** pair), it will not appear to have any effect.



## Syntax

```
SET CONSTRAINTS { ALL | { name } [, ...] } { DEFERRED | IMMEDIATE } ;
```

## Parameter Description

- **name**  
Specifies the constraint name.  
Value range: an existing table name, which can be found in the system catalog **pg\_constraint**.
- **ALL**  
Specifies all constraints.
- **DEFERRED**  
Specifies that constraints are not checked until transaction commit.
- **IMMEDIATE**  
Specifies that constraints are checked at the end of each statement.

## Examples

```
-- Set that constraints are checked when a transaction is committed.
openGauss=# SET CONSTRAINTS ALL DEFERRED;
```

## 7.13.146 SET ROLE

### Description

Sets the current user identifier of the current session.

### Precautions

- Users of the current session must be members of specified **rolename**, but system administrators can choose any roles.
- This statement may add permissions to or restrict permissions of a user. If the role of a session user has the **INHERITS** attribute, it automatically has all permissions of roles that SET ROLE enables the role to be. In this case, SET ROLE physically deletes all permissions directly granted to session users and permissions of its belonging roles and only leaves permissions of the specified roles. If the role of the session user has the **NOINHERITS** attribute, SET ROLE deletes permissions directly granted to the session user and obtains permissions of the specified role.

## Syntax

- Set the current user identifier of the current session.  

```
SET [SESSION | LOCAL] ROLE role_name PASSWORD 'password';
```
- Reset the current user identifier to that of the current session.  

```
RESET ROLE;
```

## Parameters

- **SESSION**  
Specifies that the statement takes effect only for the current session. This parameter is used by default.

- **LOCAL**  
Specifies that the specified statement takes effect only for the current transaction.
- **role\_name**  
Specifies the role name.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **password**  
Specifies the password of a role. It must comply with the password convention. Encrypted passwords are not supported.
- **RESET ROLE**  
Resets the current user identifier.

## Examples

```
-- Create a role paul.
openGauss=# CREATE ROLE paul IDENTIFIED BY '*****';

-- Set the current user to paul.
openGauss=# SET ROLE paul PASSWORD '*****';

-- View the current session user and the current user.
openGauss=> SELECT SESSION_USER, CURRENT_USER;

-- Reset the current user.
openGauss=> RESET ROLE;

-- Drop the user.
openGauss=# DROP USER paul;
```

## 7.13.147 SET SESSION AUTHORIZATION

### Description

Sets the session user identifier and the current user identifier of the current session to a specified user.

### Precautions

The session identifier can be changed only when the initial session user has the SYSADMIN permission. Otherwise, the system supports the statement only when the authenticated username is specified.

### Syntax

- Set the session user identifier and the current user identifier of the current session.  
`SET [ SESSION | LOCAL ] SESSION AUTHORIZATION role_name PASSWORD 'password';`
- Reset the identifiers of the session and current users to the initially authenticated usernames.  
`{SET [ SESSION | LOCAL ] SESSION AUTHORIZATION DEFAULT  
| RESET SESSION AUTHORIZATION};`

## Parameters

- **SESSION**  
Specifies that the specified parameters take effect for the current session.
- **LOCAL**  
Specifies that the specified statement takes effect only for the current transaction.
- **role\_name**  
Username.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **password**  
Specifies the password of a role. It must comply with the password convention. Encrypted passwords are not supported.
- **DEFAULT**  
Resets the identifiers of the session and current users to the initially authenticated usernames.

## Examples

```
-- Create a role paul.
openGauss=# CREATE ROLE paul IDENTIFIED BY '*****!';

-- Set the current user to paul.
openGauss=# SET SESSION AUTHORIZATION paul password '*****!';

-- View the current session user and the current user.
openGauss=# SELECT SESSION_USER, CURRENT_USER;

-- Reset the current user.
openGauss=# RESET SESSION AUTHORIZATION;

-- Drop the user.
openGauss=# DROP USER paul;
```

## Reference

### [SET ROLE](#)

## 7.13.148 SET TRANSACTION

### Description

Sets the transaction characteristics. Transaction characteristics include transaction isolation level and transaction access mode (read/write or read-only). You can set the current transaction characteristics using **LOCAL** or the default transaction characteristics of a session using **SESSION**.

### Precautions

The current transaction characteristics must be set in a transaction, that is, **START TRANSACTION** or **BEGIN** must be executed before **SET TRANSACTION** is executed. Otherwise, the setting does not take effect.

## Syntax

Set the isolation level and access mode of the transaction.

```
{ SET [LOCAL] TRANSACTION|SET SESSION CHARACTERISTICS AS TRANSACTION }
{ ISOLATION LEVEL { READ COMMITTED | SERIALIZABLE | REPEATABLE READ }
| { READ WRITE | READ ONLY } } [, ...]
```

## Parameters

- **LOCAL**

Specifies that the specified statement takes effect only for the current transaction.

- **SESSION**

Specifies that the specified parameters take effect for the current session.

- **ISOLATION\_LEVEL**

Specifies the transaction isolation level that determines the data that a transaction can view if other concurrent transactions exist.

 **NOTE**

The isolation level cannot be changed after data is modified using SELECT, INSERT, DELETE, UPDATE, FETCH, or COPY in the current transaction.

Value range:

- **READ COMMITTED**: Only committed data can be read. It is the default value.
- **REPEATABLE READ**: Only the data committed before transaction start is read. Uncommitted data or data committed in other concurrent transactions cannot be read.
- **SERIALIZABLE**: Currently, this isolation level is not supported in GaussDB. It is equivalent to **REPEATABLE READ**.

- **READ WRITE | READ ONLY**

Specifies the transaction access mode (read/write or read only).

## Examples

```
-- Start a transaction and set its isolation level to READ COMMITTED and access mode to READ ONLY.
openGauss=# START TRANSACTION;
openGauss=# SET LOCAL TRANSACTION ISOLATION LEVEL READ COMMITTED READ ONLY;
openGauss=# COMMIT;
```

## 7.13.149 SHOW

### Description

SHOW shows the current value of a GUC parameter.

### Syntax

```
SHOW
{
 [VARIABLES LIKE] configuration_parameter |
 CURRENT_SCHEMA |
 TIME_ZONE |
 TRANSACTION ISOLATION LEVEL |
```

```
SESSION AUTHORIZATION |
ALL
};
```

## Parameters

See [Parameters](#) in "RESET."

## Examples

```
-- Show the value of timezone.
openGauss=# SHOW timezone;

-- Show all parameters.
openGauss=# SHOW ALL;

-- Show all parameters whose names contain var.
openGauss=# SHOW VARIABLES LIKE var;
```

## Helpful Links

[SET](#) and [RESET](#)

## 7.13.150 SHUTDOWN

### Description

SHUTDOWN shuts down the currently connected database node.

### Precautions

Only administrators can run this command.

### Syntax

```
SHUTDOWN [FAST | IMMEDIATE];
```

### Parameters

- **""**  
If the shutdown mode is not specified, the default mode **fast** is used.
- **fast**  
Rolls back all active transactions, forcibly disconnects the client, and shuts down the database node without waiting for the client to disconnect.
- **immediate**  
Shuts down the server forcibly. Fault recovery will occur on the next startup.

### Examples

```
-- Shut down the current database node.
openGauss=# SHUTDOWN;

-- Shut down the current database node in fast mode.
openGauss=# SHUTDOWN FAST;
```

## 7.13.151 START TRANSACTION

### Description

START TRANSACTION starts a transaction. If the isolation level or read/write mode is specified, a new transaction will have those characteristics. You can also specify them using [SET TRANSACTION](#).

### Syntax

#### Format 1: START TRANSACTION

```
START TRANSACTION
[
 {
 ISOLATION LEVEL { READ COMMITTED | SERIALIZABLE | REPEATABLE READ }
 | { READ WRITE | READ ONLY }
 } [, ...]
];
```

#### Format 2: BEGIN

```
BEGIN [WORK | TRANSACTION]
[
 {
 ISOLATION LEVEL { READ COMMITTED | SERIALIZABLE | REPEATABLE READ }
 | { READ WRITE | READ ONLY }
 } [, ...]
];
```

### Parameters

- **WORK | TRANSACTION**

Specifies the optional keyword in BEGIN format without functions.

- **ISOLATION LEVEL**

Specifies the transaction isolation level that determines the data that a transaction can view if other concurrent transactions exist.

 **NOTE**

The isolation level cannot be changed after data is modified using SELECT, INSERT, DELETE, UPDATE, FETCH, or COPY in the transaction.

Value range:

- **READ COMMITTED**: Only committed data can be read. It is the default value.
- **REPEATABLE READ**: Only the data committed before transaction start is read. Uncommitted data or data committed in other concurrent transactions cannot be read.
- **SERIALIZABLE**: Currently, this isolation level is not supported in GaussDB. It is equivalent to **REPEATABLE READ**.

- **READ WRITE | READ ONLY**

Specifies the transaction access mode (read/write or read only).

### Examples

```
-- Start a transaction in default mode.
openGauss=# START TRANSACTION;
```

```
openGauss=# SELECT * FROM tpceds.reason;
openGauss=# END;

-- Start a transaction in default mode.
openGauss=# BEGIN;
openGauss=# SELECT * FROM tpceds.reason;
openGauss=# END;

-- Start a transaction with the isolation level being READ COMMITTED and the access mode being READ WRITE:
openGauss=# START TRANSACTION ISOLATION LEVEL READ COMMITTED READ WRITE;
openGauss=# SELECT * FROM tpceds.reason;
openGauss=# COMMIT;
```

## Helpful Links

[COMMIT | END, ROLLBACK](#), and [SET TRANSACTION](#)

## 7.13.152 TIMECAPSULE TABLE

### Description

The TIMECAPSULE TABLE statement restores a table to an earlier state in the event of human or application errors.

The table can flash back to a past point in time, depending on the old data stored in the system. In addition, GaussDB cannot restore a table to an earlier state through DDL operations that has changed the structure of the table.

### Precautions

- The TIMECAPSULE TABLE statement can be used to flash back the old data or the data from the recycle bin.
  - **TO TIMECAPSULE** and **TO CSN** can flash back a table to an earlier version.
  - The recycle bin records the objects dropped or truncated by running **DROP** and **TRUNCATE**. **TO BEFORE DROP** and **TO BEFORE TRUNCATE** flash back from the recycle bin.
- The following object types do not support flashback: system catalogs, MOTs, DFS tables, global temporary tables, local temporary tables, unlogged tables, sequence tables, and hash bucket tables.
- Between the flashback point and the current point, a statement (DDL, DCL, or VACUUM FULL) that modifies the table structure or affects physical storage has been executed. Therefore, the flashback fails.
- To run **DROP**, you must have the CREATE or USAGE permission on the schema to which the junk object belongs, and you must be the owner of the schema or the owner of the junk object.  
To run **TRUNCATE**, you must have the CREATE or USAGE permission on the schema to which the junk object belongs, and you must be the owner of the schema or the junk object. In addition, you must have the TRUNCATE permission on the junk object.
- Scenarios or tables that do not support DROP or TRUNCATE FLASHBACK
  - Scenario where the recycle bin is disabled (**enable\_recyclebin** is set to **off**)

- Scenario where the system is being maintained (**xc\_maintenance\_mode** is set to **on**) or is being upgraded
- Scenario where multiple objects are deleted (The **DROP** or **TRUNCATE TABLE** command is executed to drop multiple objects at the same time.)
- System catalogs, MOTs, DFS tables, global temporary tables, local temporary tables, UNLOGGED tables, sequence tables, and hash bucket tables.
- After TRUNCATE FLASHBACK is performed, the statistics remain unchanged and are still displayed as **0**. You can modify the statistics by manual ANALYZE during off-peak hours (to reduce the impact on performance).

## Syntax

```
TIMECAPSULE TABLE [schema.]table_name TO { CSN expr | TIMESTAMP expr | BEFORE { DROP [RENAME
TO table_name] | TRUNCATE } }
```

## Parameters

- **schema**  
Specifies a schema containing the table to be flashed back. If this parameter is not specified, the current schema is used.
- **table\_name**  
Specifies a table name.
- **TO CSN**  
Specifies the CSN corresponding to the time point when the table is to be flashed back. *expr* must be a number representing a valid CSN.
- **TO TIMESTAMP**  
Specifies a timestamp value corresponding to the point in time to which you want to flash back the table. The result of *expr* must be a valid past timestamp (convert a string to a time type using the **TO\_TIMESTAMP** function). The table will be flashed back to a time within approximately 3 seconds of the specified timestamp.  
  
Note: When the flashback point is too old, the source version cannot be obtained because it is recycled. As a result, the flashback fails and the error message "Restore point too old" is displayed.
- **TO BEFORE DROP**  
Retrieves dropped tables and their subobjects from the recycle bin.  
You can specify either the original user-specified name of the table or the system-generated name assigned to the object when it was deleted.
  - System-generated recycle bin object names are unique. Therefore, if you specify the system-generated name, the database retrieves that specified object. To see the content in your recycle bin, run **select \* from gs\_recyclebin;**
  - If you specify the user-specified name and the recycle bin contains more than one object of that name, the database retrieves the object that was moved to the recycle bin most recently. If you want to retrieve an older version of the table, then do one of these things:
    - Specify the system-generated recycle bin name of the table you want to retrieve.



- Run the **TIMECAPSULE TABLE ... TO BEFORE DROP** statement until you retrieve the table you want.
  - When a dropped table is restored, only the base table name is restored, and the names of other subobjects remain the same as those in the recycle bin. You can run the DDL command to manually change the names of subobjects as required.
  - If a table has default values that reference sequences and user-defined functions, flashing back the dropped table will succeed, but the default values will not be restored.
  - Similarly, if a table is referenced by views, dropping the table will cascadingly drop the views. Therefore, flashing back the dropped table will succeed, but the views will not be restored.
  - The recycle bin does not support write operations such as DML, DCL, and DDL, and does not support DQL query operations (supported in later versions).
  - The **recyclebin\_retention\_time** parameter has been set for specifying the retention period of objects in the recycle bin. The objects will be automatically deleted after the retention period expires.
- **RENAME TO**  
Specifies a new name for the table retrieved from the recycle bin.
  - **TRUNCATE**  
Flashes back to the point in time before the TRUNCATE operation.

## Examples

```
-- Drop the tpcds.reason_t2 table.
DROP TABLE IF EXISTS tpcds.reason_t2;
-- Create the tpcds.reason_t2 table.
openGauss=# CREATE TABLE tpcds.reason_t2
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
)with(storage_type = ustore);
-- Insert records into the tpcds.reason_t2 table.
openGauss=# INSERT INTO tpcds.reason_t2 VALUES (1, 'AA', 'reason1'),(2, 'AB', 'reason2'),(3, 'AC',
'reason3');
INSERT 0 3
-- Truncate data from the tpcds.reason_t2 table.
openGauss=# TRUNCATE TABLE tpcds.reason_t2;
-- Query data in the tpcds.reason_t2 table.
openGauss=# select * from tpcds.reason_t2;
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----+-----
(0 rows)
-- Perform the TRUNCATE FLASHBACK operation.
openGauss=# TIMECAPSULE TABLE tpcds.reason_t2 to BEFORE TRUNCATE;
openGauss=# select * from tpcds.reason_t2;
 r_reason_sk | r_reason_id | r_reason_desc
-----+-----+-----
+-----+-----+-----
 1 | AA | reason1
 2 | AB | reason2
 3 | AC | reason3
(3 rows)
-- Drop the tpcds.reason_t2 table.
openGauss=# DROP TABLE tpcds.reason_t2;
-- Perform the DROP FLASHBACK operation.
```

```
openGauss=# TIMECAPSULE TABLE tpcds.reason_t2 to BEFORE DROP;
TimeCapsule Table
```

## 7.13.153 TRUNCATE

### Function

**TRUNCATE** quickly removes all rows from a database table.

It has the same effect as an unqualified **DELETE** on each table, but **TRUNCATE** is faster because it does not actually scan the tables. This is most useful on large tables.

### Precautions

- **TRUNCATE TABLE** has the same function as a **DELETE** statement with no **WHERE** clause, emptying a table.
- **TRUNCATE TABLE** uses less system and transaction log resources as compared with **DELETE**.
  - **DELETE** deletes a row each time, and records the deletion of each row in the transaction log.
  - **TRUNCATE TABLE** deletes all rows in a table by releasing the data page storing the table data, and records the releasing of the data page only in the transaction log.
- The differences between **TRUNCATE**, **DELETE**, and **DROP** are as follows:
  - **TRUNCATE TABLE** deletes content, releases space, but does not delete definitions.
  - **DELETE TABLE** deletes content, but does not delete definitions nor release space.
  - **DROP TABLE** deletes content and definitions, and releases space.

### Syntax

- Delete data from a table.

```
TRUNCATE [TABLE] [ONLY] { table_name [*] } [, ...]
[CONTINUE IDENTITY] [CASCADE | RESTRICT] [PURGE];
```

- Truncate the data in a partition.

```
ALTER TABLE [IF EXISTS] { [ONLY] table_name
| table_name *
| ONLY (table_name) }
TRUNCATE PARTITION { partition_name
| FOR (partition_value [, ...]) } [UPDATE GLOBAL INDEX];
```

### Parameter Description

- **ONLY**  
If **ONLY** is specified, only the specified table is cleared. Otherwise, the table and all its subtables (if any) are cleared.
- **table\_name**  
Specifies the name (optionally schema-qualified) of the target table.  
Value range: an existing table name

- **CONTINUE IDENTITY**  
Does not change the values of sequences. This is the default action.
- **CASCADE | RESTRICT**
  - **CASCADE**: Clears all tables that are added to a group.
  - **RESTRICT** (default value): Clears all data.
- **PURGE**: Purges table data in the recycle bin by default.
- **partition\_name**  
Specifies the partition in the target partitioned table.  
Value range: an existing table name
- **partition\_value**  
Specifies the value of the specified partition key.  
The value specified by **PARTITION FOR** can uniquely identify a partition.  
Value range: value range of the partition key for the partition to be renamed

**NOTICE**

When the **PARTITION FOR** clause is used, the entire partition where **partition\_value** is located is cleared.

- **UPDATE GLOBAL INDEX**  
If this parameter is used, all global indexes in the partitioned table are updated to ensure that data can be queried correctly using global indexes. If this parameter is not used, all global indexes in the partitioned table will become invalid.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason values(1,'AAAAAAAABAAAAAAA','reason 1'),
(5,'AAAAAAAABAAAAAAA','reason 2'),(15,'AAAAAAAABAAAAAAA','reason 3'),
(25,'AAAAAAAABAAAAAAA','reason 4'),(35,'AAAAAAAABAAAAAAA','reason 5'),
(45,'AAAAAAAACAAAAAAA','reason 6'),(55,'AAAAAAAACAAAAAAA','reason 7');

-- Create a table.
openGauss=# CREATE TABLE tpcds.reason_t1 AS TABLE tpcds.reason;

-- Truncate the tpcds.reason_t1 table.
openGauss=# TRUNCATE TABLE tpcds.reason_t1;

-- Delete the table.
openGauss=# DROP TABLE tpcds.reason_t1;
-- Create a partitioned table.
openGauss=# CREATE TABLE tpcds.reason_p
(
```

```
r_reason_sk integer,
r_reason_id character(16),
r_reason_desc character(100)
)PARTITION BY RANGE (r_reason_sk)
(
 partition p_05_before values less than (05),
 partition p_15 values less than (15),
 partition p_25 values less than (25),
 partition p_35 values less than (35),
 partition p_45_after values less than (MAXVALUE)
);

-- Insert data.
openGauss=# INSERT INTO tpcds.reason_p SELECT * FROM tpcds.reason;

-- Clear the p_05_before partition.
openGauss=# ALTER TABLE tpcds.reason_p TRUNCATE PARTITION p_05_before;

-- Clear the p_15 partition.
openGauss=# ALTER TABLE tpcds.reason_p TRUNCATE PARTITION for (15);

-- Clear the partitioned table.
openGauss=# TRUNCATE TABLE tpcds.reason_p;

-- Delete the table.
openGauss=# DROP TABLE tpcds.reason_p;

-- Delete the table.
openGauss=# DROP TABLE tpcds.reason;

-- Delete the schema.
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## 7.13.154 UPDATE

### Description

Updates data in a table. UPDATE changes the values of the specified columns in all rows that satisfy the condition. The WHERE clause clarifies conditions. The columns to be modified need to be mentioned in the SET clause; columns not explicitly modified retain their previous values.

### Precautions

- The owner of a table, users granted with the UPDATE permission on the table, or users granted with the UPDATE ANY TABLE permission can update data in the table. System administrators have the permission to update data in the table by default.
- You must have the SELECT permission on all tables involved in the expressions or conditions.
- The generated column cannot be directly written. In the **UPDATE** statement, values cannot be specified for generated columns, but the keyword **DEFAULT** can be specified.

### Syntax

```
[WITH [RECURSIVE] with_query [, ...]]
UPDATE [/*+ plan_hint */] [ONLY] table_name [partition_clause] [*] [[AS] alias]
SET {column_name = { expression | DEFAULT }
 | (column_name [, ...]) = {({ expression | DEFAULT } [, ...]) | sub_query }}, ...]
[FROM from_list] [WHERE condition]
[RETURNING {*
```

```
| {output_expression [[AS] output_name]} [, ...] }];
```

where sub\_query can be:

```
SELECT [ALL | DISTINCT [ON (expression [, ...])]]
{ * | {expression [[AS] output_name]} [, ...] }
[FROM from_item [, ...]]
[WHERE condition]
[GROUP BY grouping_element [, ...]]
[HAVING condition [, ...]]
[ORDER BY {expression [[ASC | DESC | USING operator] | nlssort_expression_clause } [NULLS { FIRST |
LAST }]} [, ...]]
[LIMIT { [offset,] count | ALL }]
```

The subquery **with\_query** is as follows:

```
with_query_name [(column_name [, ...])]
AS [[NOT] MATERIALIZED] ({select | values | insert | update | delete})
```

## Parameters

- **WITH [ RECURSIVE ] with\_query [, ...]**

Specifies one or more subqueries that can be referenced by name in the main query, which is equivalent to a temporary table. This subquery statement structure is called the common table expression (CTE) structure. When this structure is used, the execution plan contains the CTE SCAN content.

If **RECURSIVE** is specified, it allows a SELECT subquery to reference itself by name.

Format of **with\_query**:

```
with_query_name [(column_name [, ...])] AS [[NOT] MATERIALIZED] ({select | values | insert |
update | delete})
```

- **with\_query\_name** specifies the name of the result set generated by a subquery. Such names can be used to access the result sets of subqueries in a query.
- **column\_name** specifies the column name displayed in the subquery result set.
- Each subquery can be a SELECT, VALUES, INSERT, UPDATE, or DELETE statement.
- You can use **MATERIALIZED** or **NOT MATERIALIZED** to modify the CTE.
  - If **MATERIALIZED** is specified, the WITH query will be materialized, and a copy of the subquery result set is generated. The copy is directly queried at the reference point. Therefore, the WITH subquery cannot be jointly optimized with the SELECT statement trunk (for example, predicate pushdown and equivalence class transfer). In this scenario, you can use **NOT MATERIALIZED** for modification. If the WITH query can be executed as a subquery inline, the preceding optimization can be performed.
  - If the user does not explicitly declare the materialized attribute, comply with the following rules: If the CTE is referenced only once in the SELECT statement trunk to which it belongs and semantically supports inline execution, it will be rewritten as subquery inline execution. Otherwise, the materialized execution will be performed in CTE Scan mode.

- **plan\_hint** clause

Follows the UPDATE keyword in the **/\*+ \*/** format. It is used to optimize the plan of an UPDATE statement block. For details, see [Hint-based Tuning](#). In

each statement, only the first */\*+ plan\_hint\*/* comment block takes effect as a hint. Multiple hints can be written.

- **table\_name**  
Specifies the name (optionally schema-qualified) of the table to be updated.  
Value range: an existing table name.
- **partition\_clause**  
Updates a specified partition.  
PARTITION { ( partition\_name ) | FOR ( partition\_value [, ...] ) } |  
SUBPARTITION { ( subpartition\_name ) | FOR ( subpartition\_value [, ...] ) }  
For details about the keywords, see [SELECT](#).  
For details, see [CREATE TABLE SUBPARTITION](#).
- **alias**  
Specifies the alias of the target table.  
Value range: a string that complies with the [Identifier Naming Conventions](#).
- **column\_name**  
Specifies the name of the column to be modified.  
You can refer to this column by specifying the target table alias and the column name. For example:  
UPDATE foo AS f SET f.col\_name = 'namecol';  
Value range: name of an existing column.
- **expression**  
Specifies a value assigned to a column or an expression that assigns the value.
- **DEFAULT**  
Specifies the default value of a column.  
The value is **NULL** if no default value has been assigned to it.
- **sub\_query**  
Specifies a subquery.  
This statement can be executed to update a table with information for other tables in the same database. For details about clauses in the **SELECT** statement, see [SELECT](#).  
When a single column is updated, the ORDER BY and LIMIT clauses can be used. When multiple columns are updated, the ORDER BY and LIMIT clauses cannot be used.
- **from\_list**  
Specifies a list of table expressions. You can use columns of other tables in the WHERE condition. It is similar to specifying a table list in a FROM clause of a SELECT statement.

---

#### NOTICE

Note that the target table cannot appear in the **from\_list**, unless you intend a self-join (in which case it must appear with an alias in the **from\_list**).

---

- **condition**  
Specifies an expression that returns a value of type Boolean. Only rows for which this expression returns **true** are updated. You are advised not to use numeric types such as int for **condition**, because such types can be implicitly converted to bool values (non-zero values are implicitly converted to **true** and 0 is implicitly converted to **false**), which may cause unexpected results.
- **output\_expression**  
Specifies an expression to be computed and returned by the **UPDATE** command after each row is updated.  
Value range: any table and table columns listed in FROM. Write \* to return all columns.
- **output\_name**  
Specifies a name to use for a returned column.

## Examples

```
-- Create the student1 table.
openGauss=# CREATE TABLE student1
(
 stuno int,
 classno int
);

-- Insert data.
openGauss=# INSERT INTO student1 VALUES(1,1);
openGauss=# INSERT INTO student1 VALUES(2,2);
openGauss=# INSERT INTO student1 VALUES(3,3);

-- View data.
openGauss=# SELECT * FROM student1;

-- Update the values of all records.
openGauss=# UPDATE student1 SET classno = classno*2;

-- View data.
openGauss=# SELECT * FROM student1;

-- Drop the table.
openGauss=# DROP TABLE student1;
```

## 7.13.155 VACUUM

### Description

VACUUM recycles storage space occupied by deleted rows in tables or B-Tree indexes. In normal database operation, rows that have been deleted are not physically removed from their table; they remain present until a **VACUUM** is done. Therefore, it is necessary to execute **VACUUM** periodically, especially on frequently-updated tables.

### Precautions

- If no table is specified, VACUUM processes the tables on which the user has the corresponding permission in the current database. If a table is specified, VACUUM processes only the specified table.
- To perform VACUUM operation to a table, you must be a table owner or a user granted the VACUUM permission on the table. By default, system

administrators have this permission. However, database owners are allowed to VACUUM all tables in their databases, except shared catalogs. (The restriction for shared catalogs means that a true database-wide VACUUM can only be executed by a system administrator). VACUUM skips over any tables that the calling user does not have permission to vacuum.

- **VACUUM** cannot be executed inside a transaction block.
- It is recommended that active production databases be vacuumed frequently (at least nightly), in order to remove dead rows. After adding or deleting a large number of rows, it might be a good idea to run **VACUUM ANALYZE** for the affected table. This will update the system catalogs with the results of all recent changes, and allow the query planner to make better choices in planning queries.
- **FULL** is recommended only in special scenarios. For example, you wish to physically narrow the table to decrease the occupied disk space after deleting most rows of a table. **VACUUM FULL** usually shrinks a table more than **VACUUM** does. The **FULL** option does not clear indexes. You are advised to periodically run the **REINDEX** statement. Deleting all indexes, running **VACUUM FULL**, and rebuilding indexes is usually a faster choice. If the physical space usage does not decrease after you run the statement, check whether there are other active transactions (that have started before you delete data transactions and not ended before you run **VACUUM FULL**). If there are such transactions, run this statement again when the transactions quit.
- VACUUM causes a significant increase in I/O traffic, which may affect the performance of other active sessions. Therefore, it is sometimes advisable to use the cost-based VACUUM delay feature.
- When **VERBOSE** is specified, VACUUM prints progress messages to indicate which table is currently being processed. Various statistics about the tables are printed as well.
- When the option list is surrounded by parentheses, the options can be written in any order. If there are no brackets, the options must be given in the order displayed in the syntax.
- **VACUUM** and **VACUUM FULL** clear deleted tuples after the delay specified by **vacuum\_defer\_cleanup\_age**.
- When VACUUM ANALYZE is used, a VACUUM operation and then an ANALYZE operation are performed for each selected table. This is a handy combination form for routine maintenance scripts.
- A plain VACUUM statement (without the **FULL** option) simply recycles space and makes it available for reuse. This form of statement can operate in parallel with normal reading and writing of the table, as an exclusive lock is not obtained. **VACUUM FULL** executes wider processing, including moving rows across blocks to compress tables so they occupy the minimum number of disk blocks. This form is much slower and requires an exclusive lock on each table while it is being processed.
- A deadlock may occur when multiple VACUUM FULL statements are executed simultaneously.
- If the **xc\_maintenance\_mode** parameter is not enabled, VACUUM FULL will skip all system catalogs.
- If you run **VACUUM FULL** immediately after running **DELETE**, the space will not be recycled. After executing **DELETE**, execute 1000 non-SELECT



transactions, or wait for 1s and then execute one transaction. Then, run **VACUUM FULL** to the space.

## Syntax

- Recycle space and update statistics information, without requirements for keyword orders.

```
VACUUM [({ FULL | FREEZE | VERBOSE | {ANALYZE | ANALYSE } } [,...])]
[table_name [(column_name [, ...])] [PARTITION (partition_name) | SUBPARTITION
(subpartition_name)]];
```

- Recycle space, without updating statistics information.

```
VACUUM [FULL [COMPACT]] [FREEZE] [VERBOSE] [table_name
[PARTITION (partition_name) | SUBPARTITION (subpartition_name)]];
```

- Recycle space and update statistics information, and require keywords in order.

```
VACUUM [FULL] [FREEZE] [VERBOSE] { ANALYZE | ANALYSE } [VERBOSE]
[table_name [(column_name [, ...])]] [PARTITION (partition_name)];
```

## Parameters

- FULL**

Selects "FULL" vacuum, which can recycle more space, but takes much longer and exclusively locks the table.

### NOTE

- Using FULL will cause statistics missing. To collect statistics, add the keyword ANALYZE to VACUUM FULL.
- Ustore does not support the DDL statement VACUUM FULL. After VACUUM FULL is executed, "INFO: skipping "test" --- Ustore table does not support vacuum full" is printed.

- FREEZE**

Is equivalent to running **VACUUM** with the **vacuum\_freeze\_min\_age** parameter set to **zero**.

- VERBOSE**

Prints a detailed vacuum activity report for each table.

- ANALYZE | ANALYSE**

Updates statistics used by the planner to determine the most efficient way to execute a query.

### NOTE

VACUUM is also triggered when **autovacuum** is set to **analyze** for a Ustore partitioned table.

- table\_name**

Specifies the name (optionally schema-qualified) of a specific table to vacuum.

Value range: name of a specific table to vacuum. Defaults are all tables in the current database.

- column\_name**

Specifies the name of the column to be analyzed. This parameter must be used together with **ANALYZE**.

Value range: name of the column to be analyzed Defaults are all columns.

 NOTE

The mechanism of the VACUUM ANALYZE statement is to execute **VACUUM** and **ANALYZE** in sequence. Therefore, if **column\_name** is incorrect, **VACUUM** may be successfully executed but **ANALYZE** may fail to be executed. For a partitioned table, **ANALYZE** may fail to be executed after **VACUUM** is successfully executed on a partition.

- **PARTITION**  
**COMPACT** and **PARTITION** cannot be used at the same time.
- **partition\_name**  
Specifies the level-1 partition name of the table to be cleared. If it is left empty, all level-1 partitions are cleared.
- **subpartition\_name**  
Specifies the level-2 partition name of the table to be cleared. If it is left empty, all level-2 partitions are cleared.

## Examples

```
-- Create a schema.
openGauss=# CREATE SCHEMA tpcds;

-- Create the tpcds.reason table.
openGauss=# CREATE TABLE tpcds.reason
(
 r_reason_sk integer,
 r_reason_id character(16),
 r_reason_desc character(100)
);

-- Insert multiple records into the table.
openGauss=# INSERT INTO tpcds.reason values(1,'AAAAAAAABAAAAAAA','reason 1'),
(2,'AAAAAAAABAAAAAAA','reason 2');

-- Create an index in the tpcds.reason table.
openGauss=# CREATE UNIQUE INDEX ds_reason_index1 ON tpcds.reason(r_reason_sk);

-- Vacuum the tpcds.reason table that has indexes.
openGauss=# VACUUM (VERBOSE, ANALYZE) tpcds.reason;

-- Drop the index.
openGauss=# DROP INDEX tpcds.ds_reason_index1 CASCADE;
openGauss=# DROP TABLE tpcds.reason;
openGauss=# DROP SCHEMA tpcds CASCADE;
```

## Suggestions

- vacuum
  - **VACUUM** cannot be executed inside a transaction block.
  - It is recommended that active production databases be vacuumed frequently (at least nightly), in order to remove dead rows. It is strongly recommended that you run **VACUUM ANALYZE** after adding or deleting a large number of records.
  - **FULL** is recommended only in special scenarios. For example, you wish to physically narrow the table to decrease the occupied disk space after deleting most rows of a table.
  - Before performing the VACUUM FULL operation, you are advised to delete all indexes in related tables, run **VACUUM FULL**, and then rebuild the index.

## 7.13.156 VALUES

### Description

Computes a row or a set of rows based on given values. It is most commonly used to generate a constant table within a large statement.

### Precautions

- VALUES lists with large numbers of rows should be avoided, as you might encounter out-of-memory failures or poor performance. VALUES appearing within INSERT is a special case, because the desired column types are known from the INSERT's target table, and need not be inferred by scanning the VALUES list. In this case, a large number of result rows can be processed.
- If more than one row is specified, all the rows must have the same number of elements.

### Syntax

```
VALUES {(expression [, ...])} [, ...]
 [ORDER BY { sort_expression [ASC | DESC | USING operator] } [, ...]]
 [LIMIT { count | ALL }]
 [OFFSET start [ROW | ROWS]]
 [FETCH { FIRST | NEXT } [count] { ROW | ROWS } ONLY];
```

### Parameters

- expression**  
Specifies a constant or expression to compute and insert at the indicated place in the resulting table or set of rows.  
In a VALUES list appearing at the top level of an INSERT, an expression can be replaced by DEFAULT to indicate that the destination column's default value should be inserted. DEFAULT cannot be used when VALUES appears in other contexts.
- sort\_expression**  
Specifies an expression or integer constant indicating how to sort the result rows.
- ASC**  
Specifies an ascending sort order.
- DESC**  
Specifies a descending sort order.
- operator**  
Specifies a sorting operator.
- count**  
Specifies the maximum number of rows to return.
- OFFSET start [ ROW | ROWS ]**  
Specifies the maximum number of returned rows, whereas **start** specifies the number of rows to skip before starting to return rows.
- FETCH { FIRST | NEXT } [ count ] { ROW | ROWS } ONLY**

The FETCH clause restricts the total number of rows starting from the first row of the return query result, and the default value of **count** is **1**.

## Examples

See [Examples](#) in "INSERT."

## 7.14 Appendix

### 7.14.1 Extended Functions

The following table lists the extended functions supported by GaussDB. These functions are for reference only.

| Category                          | Name                                              | Description                                                     |
|-----------------------------------|---------------------------------------------------|-----------------------------------------------------------------|
| Access privilege inquiry function | has_sequence_privilege(user, sequence, privilege) | Queries whether a specified user has privilege for sequences.   |
|                                   | has_sequence_privilege(sequence, privilege)       | Queries whether the current user has privilege for sequence.    |
| Trigger function                  | pg_get_triggerdef(oid)                            | Gets <b>CREATE [ CONSTRAINT ] TRIGGER</b> command for triggers. |
|                                   | pg_get_triggerdef(oid, boolean)                   | Gets <b>CREATE [ CONSTRAINT ] TRIGGER</b> command for triggers. |

### 7.14.2 Extended Syntax

GaussDB provides the following extended syntax:

**Table 7-110** Extended SQL syntax

| Category         | Keyword                                   | Description                                                                 |
|------------------|-------------------------------------------|-----------------------------------------------------------------------------|
| Creating a table | <b>INHERITS ( parent_table [, ... ] )</b> | Currently, the syntax of inherited tables is retained but is not supported. |

| Category            | Keyword                                                                                                                                                                                                      | Description                                                                                                                                                                        |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                     | column_constraint:<br><b>REFERENCES reftable</b><br><b>[ ( refcolumn ) ] [ MATCH</b><br><b>FULL   MATCH PARTIAL   MATCH</b><br><b>SIMPLE ] [ ON</b><br><b>DELETE action ] [ ON</b><br><b>UPDATE action ]</b> | You can run <b>REFERENCES reftable[(refcolumn)] [MATCH FULL  MATCH PARTIAL   MATCH SIMPLE] [ON DELETE action] [ON UPDATE action]</b> to create foreign key constraints for tables. |
| Aggregate functions | <b>CREATE AGGREGATE</b>                                                                                                                                                                                      | Defines a new aggregate function.                                                                                                                                                  |
|                     | <b>ALTER AGGREGATE</b>                                                                                                                                                                                       | Alters the definition of an aggregate function.                                                                                                                                    |
|                     | <b>DROP AGGREGATE</b>                                                                                                                                                                                        | Drops an existing aggregate function.                                                                                                                                              |

# 8 Best Practices

## 8.1 Best Practices for Table Design

### 8.1.1 Selecting a Storage Model

During database design, some key factors about table design will greatly affect the subsequent query performance of the database. Table design affects data storage as well. A good table design reduces I/O operations and minimizes memory usage, improving the query performance.

Selecting a model for table storage is the first step of table definition. Select a proper storage model for your service based on the following table:

| Storage Model | Applicable Scenario                                                                                                                                       |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Row store     | Point queries (simple index-based queries that only return a few records)<br>Scenarios requiring frequent addition, deletion, and modification operations |

### 8.1.2 Using Partitioned Tables

A partitioned table is a logical table that is divided into several physical partitions based on a specific plan. The table based on the logic is called a partitioned table, and each physical block is called a partition. A partitioned table is a logical table and does not store data. Data is actually stored in partitions. A partitioned table has the following advantages over an ordinary table:

1. High query performance: You can specify partitions when querying partitioned tables, improving query efficiency.
2. High availability: If a certain partition in a partitioned table is faulty, data in the other partitions is still available.

3. **Easy maintenance:** To fix a partitioned table having a faulty partition, you only need to fix the partition.

Partitioned tables supported by the GaussDB database are level-1 and level-2 partitioned tables. Level-1 partitioned tables include range partitioned tables, interval partitioned tables, list partitioned tables, and hash partitioned tables. Level-2 partitioned tables include nine combinations of any two of range partitioned tables, list partitioned tables, and hash partitioned tables.

- **Range partitioned table:** Data within a certain range is mapped to each partition. The range is determined by the partition key specified when the partitioned table is created. This partitioning method is most commonly used. The partition key is usually a date. For example, sales data is partitioned by month.
- **Interval partitioned table:** a special type of range partitioned tables. Compared with range partitioned tables, interval value definition is added. When no matching partition can be found for an inserted record, a partition can be automatically created based on the interval value.
- **List partitioned table:** Key values contained in the data are stored in different partitions, and the data is mapped to each partition in sequence. The key values contained in the partitions are specified when the partitioned table is created.
- **Hash partitioned table:** Data is mapped to each partition based on the internal hash algorithm. The number of partitions is specified when the partitioned table is created.
- **Level-2 partitioned table:** a partitioned table obtained by randomly combining range partitioning, list partitioning, and hash partitioning. Both level-1 and level-2 partitions can be defined in the preceding three ways.

### 8.1.3 Selecting a Data type

Efficient data types include the following:

1. **Select data types that facilitate data calculation.**

Generally, the calculation of integers (including common comparison calculations, such as =, >, <, ≥, ≤, and ≠, and GROUP BY) is more efficient than that of strings and floating-point numbers.

2. **Select data types with a short length.**

Data types with short length reduce both the data file size and the memory used for computing, improving the I/O and computing performance. For example, use **SMALLINT** instead of **INT**, and **INT** instead of **BIGINT**.

3. **Use the same data type for a join.**

You are advised to use the same data type for a join. To join columns with different data types, the database needs to convert them to the same type, which leads to additional performance overheads.

## 8.2 Best Practices for SQL Queries

Based on the SQL execution mechanism and a large number of practices, SQL statements can be optimized by following certain rules to enable the database to execute SQL statements more quickly and obtain correct results.

- Replace **UNION** with **UNION ALL**.  
**UNION** eliminates duplicate rows while merging two result sets but **UNION ALL** merges the two result sets without deduplication. Therefore, replace **UNION** with **UNION ALL** if you are sure that the two result sets do not contain duplicate rows based on the service logic.
- Add **NOT NULL** to the join columns.  
If there are many **NULL** values in the **JOIN** columns, you can add the filter criterion **IS NOT NULL** to filter data in advance to improve the **JOIN** efficiency.
- Replace **NOT IN** with **NOT EXISTS**.  
**Nested Loop Anti Join** must be used to implement **NOT IN**, and **Hash Anti Join** is required for **NOT EXISTS**. If no **NULL** value exists in the **JOIN** columns, **NOT IN** is equivalent to **NOT EXISTS**. Therefore, if you are sure that no **NULL** value exists, you can convert **NOT IN** to **NOT EXISTS** to generate hash join and to improve the query performance.

The statements for creating a foreign table are as follows:

```
DROP SCHEMA IF EXISTS no_in_to_no_exists_test CASCADE;
CREATE SCHEMA no_in_to_no_exists_test;
SET CURRENT_SCHEMA=no_in_to_no_exists_test;
CREATE TABLE t1(c1 int, c2 int, c3 int);
CREATE TABLE t2(d1 int, d2 int NOT NULL, d3 int);
```

The statement for implementing the query using **NOT IN** is as follows:

```
SELECT * FROM t1 WHERE c1 NOT IN (SELECT d2 FROM t2);
```

The plan is as follows:

```
openGauss=# EXPLAIN SELECT * FROM t1 WHERE c1 NOT IN (SELECT d2 FROM t2);
 QUERY PLAN

Nested Loop Anti Join (cost=0.00..29749.02 rows=968 width=12)
 Join Filter: ((t1.c1 = t2.d2) OR (t1.c1 IS NULL))
 -> Seq Scan on t1 (cost=0.00..29.45 rows=1945 width=12)
 -> Materialize (cost=0.00..39.17 rows=1945 width=4)
 -> Seq Scan on t2 (cost=0.00..29.45 rows=1945 width=4)
(5 rows)
```

Because there is no null value in the **t2.d2** column (the **t2.d2** column is **NOT NULL** in the table definition), the query can be equivalently modified as follows:

```
SELECT * FROM t1 WHERE NOT EXISTS (SELECT * FROM t2 WHERE t1.c1=t2.d2);
```

The generated plan is as follows:

```
openGauss=# EXPLAIN SELECT * FROM t1 WHERE NOT EXISTS (SELECT * FROM t2 WHERE
t1.c1=t2.d2);
 QUERY PLAN

Hash Anti Join (cost=53.76..99.14 rows=972 width=12)
 Hash Cond: (t1.c1 = t2.d2)
 -> Seq Scan on t1 (cost=0.00..29.45 rows=1945 width=12)
 -> Hash (cost=29.45..29.45 rows=1945 width=4)
 -> Seq Scan on t2 (cost=0.00..29.45 rows=1945 width=4)
(5 rows)
```

- Use **hashagg**.  
If the **GROUP BY** condition exists in the query statement, the generated plan may contain sorting operations, that is, the plan contains the **GroupAgg+Sort** operator. As a result, the performance is poor. You can set the GUC parameter **work\_mem** to increase the available memory and generate a plan with **HashAgg** to avoid sorting operations and improve performance. For details about how to set **work\_mem**, contact the administrator.



- Replace functions with **CASE** statements.  
The database performance greatly deteriorates if a large number of functions are called. In this case, you can change the pushdown functions to CASE statements.
- Do not use functions or expressions for indexes.  
Using functions or expressions for indexes will stop indexing and enable scanning on the full table.
- Do not use operators **!=**, **<**, **>**, **NULL**, **OR**, or implicit parameter conversion in **WHERE** clauses.
- Split complex SQL statements.  
You can split an SQL statement into several ones and save the execution result to a temporary table if the SQL statement is too complex to be tuned using the solutions above, including but not limited to the following scenarios:
  - The same subquery is involved in multiple SQL statements of a job and the subquery contains a large amount of data.
  - Incorrect plan cost causes a small hash bucket of subquery. For example, the actual number of rows is 10 million, but only 1000 rows are in hash bucket.
  - Functions such as substr and to\_number cause incorrect measures for subqueries containing a large amount of data.
  - BROADCAST subqueries are performed on large tables in multi-DN environment.

For details about optimization, see [Typical SQL Optimization Methods](#).

# 9 User-defined Functions

---

When the database instance is started, the UDF master process is started in addition to the GaussDB main process. To execute a UDF in fenced mode, the UDF master process forks itself to UDF worker processes, and UDF worker processes execute the UDF in fenced mode.

## 9.1 PL/SQL Functions

PL/SQL is a loadable procedural language.

Functions created using PL/SQL can be used in any place where you can use built-in functions. For example, you can create calculation functions with complex conditions and use them to define operators or use them for index expressions.

SQL is used by most databases as a query language. It is portable and easy to learn. Each SQL statement must be executed independently by a database server.

This means that the client application performs the following processes for each query: send a query to the database server, wait for the query to be received, receive and process the result, perform related calculation, and then send more queries to the server. If the client and the database server are not on the same machine, this process also causes inter-process communication and network load.

PL/SQL enables a whole computing part and a series of queries to be grouped inside a database server. This makes procedural language available and SQL easier to use. In addition, the client/server communication cost is reduced.

- Extra round-trip communication between clients and servers is eliminated.
- Intermediate results that are not required by clients do not need to be sorted or transmitted between the clients and servers.
- Parsing can be skipped in multiple rounds of queries.

PL/SQL can use all data types, operators, and functions in SQL. For details about the PL/SQL syntax for creating functions, see [CREATE FUNCTION](#).

PL/SQL is a loadable procedural language. Its application method is similar to that of [Stored Procedure](#). However, stored procedures have no return values but the PL/SQL functions have.

# 10 Stored Procedure

---

## 10.1 Stored Procedure

In GaussDB, business rules and logics are saved as stored procedures.

A stored procedure is a combination of SQL and PL/SQL. Stored procedures can move the code that executes business rules from applications to databases. Therefore, the code storage can be used by multiple programs at a time.

For details about how to create and call a stored procedure, see [CREATE PROCEDURE](#).

The application methods for PL/SQL functions mentioned in [PL/SQL Functions](#) are the same as those for stored procedures. Unless otherwise specified, the following sections apply to stored procedures and PL/SQL functions.

## 10.2 Data Types

A data type refers to a value set and an operation set defined on the value set. The GaussDB database consists of tables, each of which is defined by its own columns. Each column corresponds to a data type. GaussDB uses corresponding functions to perform operations on data based on data types. For example, GaussDB can perform addition, subtraction, multiplication, and division operations on data of numeric values.

## 10.3 Data Type Conversion

Certain data types in the database support implicit data type conversions, such as assignments and parameters called by functions. For other data types, you can use the type conversion functions provided by GaussDB, such as the **CAST** function, to forcibly convert them.

[Table 10-1](#) lists common implicit data type conversions in GaussDB.

**NOTICE**

The valid value range of **DATE** supported by GaussDB is from 4713 BC to 294276 AD.

**Table 10-1** Implicit data type conversions

| Raw Data Type | Target Data Type | Remarks                                      |
|---------------|------------------|----------------------------------------------|
| CHAR          | VARCHAR2         | -                                            |
| CHAR          | NUMBER           | Raw data must consist of digits.             |
| CHAR          | DATE             | Raw data cannot exceed the valid date range. |
| CHAR          | RAW              | -                                            |
| CHAR          | CLOB             | -                                            |
| VARCHAR2      | CHAR             | -                                            |
| VARCHAR2      | NUMBER           | Raw data must consist of digits.             |
| VARCHAR2      | DATE             | Raw data cannot exceed the valid date range. |
| VARCHAR2      | CLOB             | -                                            |
| NUMBER        | CHAR             | -                                            |
| NUMBER        | VARCHAR2         | -                                            |
| DATE          | CHAR             | -                                            |
| DATE          | VARCHAR2         | -                                            |
| RAW           | CHAR             | -                                            |
| RAW           | VARCHAR2         | -                                            |
| CLOB          | CHAR             | -                                            |
| CLOB          | VARCHAR2         | -                                            |
| CLOB          | NUMBER           | Raw data must consist of digits.             |
| INT4          | CHAR             | -                                            |
| INT4          | BOOLEAN          | -                                            |
| BOOLEAN       | INT4             | -                                            |

## 10.4 Arrays, Sets, and Records

### 10.4.1 Arrays

#### Use of Array Types

Before the use of arrays, an array type needs to be defined.

Define an array type immediately after the **AS** keyword in a stored procedure. The definition method is as follows:

```
TYPE array_type IS VARRAY(size) OF data_type;
```

In the preceding information:

- **array\_type**: indicates the name of the array type to be defined.
- **VARRAY**: indicates the array type to be defined.
- **size**: indicates the maximum number of members in the array to be defined. The value is a positive integer.
- **data\_type**: indicates the types of members in the array to be created.

#### NOTE

- In GaussDB, an array automatically increases. If an access violation occurs, a null value is returned, and no error message is reported.
- The scope of an array type defined in a stored procedure takes effect only in this stored procedure.
- It is recommended that you use one of the preceding methods to define an array type. If both methods are used to define the same array type, GaussDB prefers the array type defined in a stored procedure to declare array variables.
- **data\_type** can also be the record type defined in a stored procedure (anonymous blocks are not supported), but cannot be the array or set type defined in the stored procedure.
- When **data\_type** is set to a type that can define the length and precision, such as varchar and numeric, if an array of varray type is created in a package and called outside the package, the restrictions for the length or precision become invalid.
- The constructors of the array type can be used only in O-compatible mode.

GaussDB supports access to array elements by using parentheses, and it also supports the **extend**, **count**, **first**, **last**, **prior**, **exists**, **trim**, **next**, and **delete** functions.

#### NOTE

- If a stored procedure contains a DML statement (such as **SELECT**, **UPDATE**, **INSERT**, or **DELETE**), you are advised to use square brackets to access array elements. Using parentheses will access arrays by default. If no array exists, function expressions will be identified.
- Exercise caution when using the **DELETE** statement to delete a single element. Otherwise, the element sequence may be incorrect.
- When the CLOB size is greater than 1 GB, the table of type, record type, and CLOB cannot be used in the input or output parameter, cursor, or raise info in a stored procedure.

## Examples

```
-- Perform operations on an array in the stored procedure.
openGauss=# CREATE OR REPLACE PROCEDURE array_proc AS
DECLARE
 TYPE ARRAY_INTEGER IS VARRAY(1024) OF INTEGER;-- Define the array type.
 ARRINT ARRAY_INTEGER:= ARRAY_INTEGER(); -- Declare the variable of the array type.
BEGIN
 ARRINT.extend(10);
 FOR I IN 1..10 LOOP
 ARRINT(I) := I;
 END LOOP;
 DBE_OUTPUT.PRINT_LINE(ARRINT.COUNT);
 DBE_OUTPUT.PRINT_LINE(ARRINT(1));
 DBE_OUTPUT.PRINT_LINE(ARRINT(10));
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.FIRST));
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.LAST));
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.NEXT(ARRINT.FIRST)));
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.PRIOR(ARRINT.LAST)));
 ARRINT.TRIM();

 IF ARRINT.EXISTS(10) THEN
 DBE_OUTPUT.PRINT_LINE('Exist 10th element');
 ELSE
 DBE_OUTPUT.PRINT_LINE('Not exist 10th element');
 END IF;
 DBE_OUTPUT.PRINT_LINE(ARRINT.COUNT);
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.FIRST));
 DBE_OUTPUT.PRINT_LINE(ARRINT(ARRINT.LAST));
 ARRINT.DELETE();
END;
/

-- Call the stored procedure.
openGauss=# CALL array_proc();

-- Delete the stored procedure.
openGauss=# DROP PROCEDURE array_proc;
```

## 10.4.2 Collections

### Use of Set Types

Before the use of sets, a set type needs to be defined.

Define a collection type immediately after the AS keyword in a stored procedure. The definition method is as follows:



In the preceding information:

- **table\_type**: indicates the name of the collection type to be defined.
- **TABLE**: indicates the collection type to be defined.
- **data\_type**: indicates the type of members in the collection to be created.
- **indexby\_type**: indicates the type of the set index to be created.

 NOTE

- In GaussDB, a set automatically increases. If an access violation occurs, a null value is returned, and no error message is reported.
- The scope of a set type defined in a stored procedure takes effect only in this stored procedure.
- The index can only be of the integer or varchar type. The length of the varchar type is not restricted.
- **NOT NULL** has no function but only takes effect in the syntax.
- **data\_type** can also be the record type or set type defined in a stored procedure (anonymous blocks are not supported), but cannot be the array type.
- Variables of the nested set type cannot be used across packages.
- Variables of the TABLE OF index by type cannot be nested in a record as the input and output parameters of a stored procedure.
- Variables of the TABLE OF index by type cannot be used as input and output parameters of functions.
- The **RAISE INFO** command cannot be used to print the entire nested TABLE OF variable.
- The TABLE OF variable cannot be passed across autonomous transactions.
- The input and output parameters of a stored procedure cannot be defined as the nested TABLE OF type.

GaussDB supports access to collection elements by using parentheses, and it also supports the extend, count, first, last, prior, next, and delete functions.

The set functions support multiset union, intersect, except all, and distinct.

 NOTE

- An expression can contain only one variable of the TABLE OF index by type.
- Exercise caution when using the DELETE statement to delete a single element. Otherwise, the element sequence may be incorrect.

## Example

```
-- Perform operations on a set in the stored procedure.
openGauss=# CREATE OR REPLACE PROCEDURE table_proc AS
DECLARE
 TYPE TABLE_INTEGER IS TABLE OF INTEGER;-- Define a set type.
 TABLE_INTEGER := TABLE_INTEGER(); -- Declare the variable of the set type.
BEGIN
 TABLE_INTEGER.extend(10);
 FOR I IN 1..10 LOOP
 TABLE_INTEGER(I) := I;
 END LOOP;
 DBE_OUTPUT.PRINT_LINE(TABLE_INTEGER.COUNT);
 DBE_OUTPUT.PRINT_LINE(TABLE_INTEGER(1));
 DBE_OUTPUT.PRINT_LINE(TABLE_INTEGER(10));
END;
/

-- Call the stored procedure.
openGauss=# CALL table_proc();

-- Delete the stored procedure.
openGauss=# DROP PROCEDURE table_proc;

-- Perform operations on a nested table in the stored procedure.
openGauss=# CREATE OR REPLACE PROCEDURE nest_table_proc AS
DECLARE
 TYPE TABLE_INTEGER IS TABLE OF INTEGER;-- Define a set type.
```

```
TYPE NEST_TABLE_INTEGER IS TABLE OF TABLE_INTEGER;-- Define a set type.
NEST_TABLE_VAR NEST_TABLE_INTEGER; -- Declare a variable of the nested table type.
BEGIN
 FOR I IN 1..10 LOOP
 NEST_TABLE_VAR(I)(I) := I;
 END LOOP;
 DBE_OUTPUT.PRINT_LINE(NEST_TABLE_VAR.COUNT);
 DBE_OUTPUT.PRINT_LINE(NEST_TABLE_VAR(1)(1));
 DBE_OUTPUT.PRINT_LINE(NEST_TABLE_VAR(10)(10));
END;
/

-- Call the stored procedure.
openGauss=# CALL nest_table_proc();

-- Delete the stored procedure.
openGauss=# DROP PROCEDURE nest_table_proc;
```

## 10.4.3 Records

### Record Variables


Perform the following operations to create a record variable:

Define a record type and use this type to declare a variable.

### Syntax

For the syntax of the record type, see [Figure 10-1](#).

**Figure 10-1** Syntax of the record type



The above syntax diagram is explained as follows:

- **record\_type**: record name
- **field**: record columns
- *datatype*: specifies data types for the member variables.
- *expression*: specifies the expression for setting a default value.



 NOTE

In GaussDB:

- When assigning values to record variables, you can:
  - Declare a record type and define member variables of this type when you declare a function or stored procedure.
  - Assign the value of a record variable to another record variable.
  - Use SELECT INTO or FETCH to assign values to a record type.
  - Assign the **NULL** value to a record variable.
- The INSERT and UPDATE statements cannot use a record variable to insert or update data.
- If a member has a record type, the default value of the inner record type cannot be not transferred to the outer record type.
- When **package\_name.record\_type** is used to declare a record variable, the default value of the record type becomes invalid. Therefore, you are advised not to use **package\_name.record\_type** to declare a record variable when the record type has a default value.
- Just like a variable, a record column of the compound type does not have a default value in the declaration.
- The data type can be the record type, array type, or collection type defined in a stored procedure (anonymous blocks are not supported).

## Examples

The table definition used in the following is defined as follows:

```
openGauss=# \d emp_rec
 Table "public.emp_rec"
 Column | Type | Modifiers
-----+-----+-----
 empno | numeric(4,0) | not null
 ename | character varying(10) |
 job | character varying(9) |
 mgr | numeric(4,0) |
 hiredate | timestamp(0) without time zone |
 sal | numeric(7,2) |
 comm | numeric(7,2) |
 deptno | numeric(2,0) |

-- Perform array operations in the function.
openGauss=# CREATE OR REPLACE FUNCTION regress_record(p_w VARCHAR2)
RETURNS
VARCHAR2 AS $$
DECLARE

 -- Declare a record type.
 type rec_type is record (name varchar2(100), epno int);
 employer rec_type;

 -- Use %type to declare the record type.
 type rec_type1 is record (name emp_rec.ename%type, epno int not null :=10);
 employer1 rec_type1;

 -- Declare a record type with a default value.
 type rec_type2 is record (
 name varchar2 not null := 'SCOTT',
 epno int not null :=10);
 employer2 rec_type2;
 CURSOR C1 IS select ename,epno from emp_rec order by 1 limit 1;

BEGIN
 -- Assign a value to a member record variable.
 employer.name := 'WARD';
```

```
employer.epno = 18;
raise info 'employer name: % , epno:%', employer.name, employer.epno;

-- Assign the value of a record variable to another variable.
employer1 := employer;
raise info 'employer1 name: % , epno: %', employer1.name, employer1.epno;

-- Assign the NULL value to a record variable.
employer1 := NULL;
raise info 'employer1 name: % , epno: %', employer1.name, employer1.epno;

-- Obtain the default value of a record variable.
raise info 'employer2 name: % , epno: %', employer2.name, employer2.epno;

-- Use a record variable in the FOR loop.
for employer in select ename,empno from emp_rec order by 1 limit 1
loop
 raise info 'employer name: % , epno: %', employer.name, employer.epno;
end loop;

-- Use a record variable in the SELECT INTO statement.
select ename,empno into employer2 from emp_rec order by 1 limit 1;
raise info 'employer name: % , epno: %', employer2.name, employer2.epno;

-- Use a record variable in a cursor.
OPEN C1;
FETCH C1 INTO employer2;
raise info 'employer name: % , epno: %', employer2.name, employer2.epno;
CLOSE C1;
RETURN employer.name;
END;
$$
LANGUAGE plpgsql;

-- Call this function.
openGauss=# CALL regress_record('abc');

-- Drop the function.
openGauss=# DROP FUNCTION regress_record;
```

## 10.5 DECLARE Syntax

### 10.5.1 Basic Structure

#### Structure

A PL/SQL block can contain a sub-block which can be placed in any section. The following describes the architecture of a PL/SQL block:

- **DECLARE:** declares variables, types, cursors, and regional stored procedures and functions used in the PL/SQL block.

DECLARE

#### NOTE

This part is optional if no variable needs to be declared.

- An anonymous block may omit the **DECLARE** keyword if no variable needs to be declared.
- For a stored procedure, **AS** is used, which is equivalent to **DECLARE**. The **AS** keyword must be reserved even if there is no variable declaration part.
- **EXECUTION:** specifies procedure and SQL statements. It is the main part of a program. It is required.

```
BEGIN
```

- Exception part: processes errors. It is optional.

```
EXCEPTION
```

- End. It is required.

```
END;
```

```
/
```

#### NOTICE

You are not allowed to use consecutive tabs in the PL/SQL block because they may result in an exception when the **gsql** tool is executed with the **-r** parameter specified.

## Category

PL/SQL blocks are classified into the following types:

- Anonymous block: a dynamic block that can be executed only for once. For details about the syntax, see [Figure 10-2](#).
- Subprogram: a stored procedure, function, operator, or packages stored in a database. A subprogram created in a database can be called by other programs.


## 10.5.2 Anonymous Blocks

An anonymous block applies to a script infrequently executed or a one-off activity. An anonymous block is executed in a session and is not stored.

## Syntax

[Figure 10-2](#) shows the syntax diagrams for an anonymous block.

**Figure 10-2** anonymous\_block::=



Details about the syntax diagram are as follows:

- The execute part of an anonymous block starts with a **BEGIN** statement, has a break with an **END** statement, and ends with a semicolon (;). Type a slash (/) and press **Enter** to execute the statement.

**NOTICE**

The terminator "/" must be written in an independent row.

- The declaration section includes the variable definition, type, and cursor definition.
- A simplest anonymous block does not execute any commands. At least one statement, even a **NULL** statement, must be presented in any implementation blocks.

### 10.5.3 Subprogram

A subprogram stores stored procedures, functions, operators, and advanced packages. A subprogram created in a database can be called by other programs.

## 10.6 Basic Statements

During PL/SQL programming, you may define some variables, assign values to variables, and call other stored procedures. This chapter describes basic PL/SQL statements, including variable definition statements, value assignment statements, call statements, and return statements.

**NOTE**

You are advised not to call the SQL statements containing passwords in the stored procedures because authorized users may view the stored procedure file in the database and password information is leaked. If a stored procedure contains other sensitive information, permission to access this procedure must be configured, preventing information leakage.


### 10.6.1 Variable Definition Statements

This section describes the declaration of variables in the PL/SQL and the scope of this variable in codes.

#### Variable Declaration

For details about the variable declaration syntax, see [Figure 10-3](#).

**Figure 10-3** declare\_variable::=



The above syntax diagram is explained as follows:

- *variable\_name*: specifies the name of a variable.
- **type** indicates the type of a variable.
- **value** indicates the initial value of the variable. (If the initial value is not given, NULL is taken as the initial value.) **value** can also be an expression.

### Example

```
openGauss=# DECLARE
emp_id INTEGER := 7788; -- Define a variable and assign a value to it.
BEGIN
emp_id := 5*7784; -- Assign a value to the variable.
END;
/
```

In addition to the declaration of basic variable types, **%TYPE** and **%ROWTYPE** can be used to declare variables related to table columns or table structures.

#### **%TYPE** Attribute

**%TYPE** declares a variable to be of the same data type as a previously declared variable (for example, a column in a table). For example, if you want to define the *my\_name* variable whose data type is the same as the data type of **firstname** in **employee**, you can define the variable as follows:

```
my_name employee.firstname%TYPE
```

In this way, you can declare *my\_name* without the need of knowing the data type of **firstname** in **employee**, and the data type of *my\_name* can be automatically updated when the data type of **firstname** changes.

```
TYPE employee_record is record (id INTEGER, firstname VARCHAR2(20));
my_employee employee_record;
my_id my_employee.id%TYPE;
my_id_copy my_id%TYPE;
```

#### **%ROWTYPE** Attribute

**%ROWTYPE** declares data types of a set of data. It stores a row of table data or results fetched from a cursor. For example, if you want to define a set of data with the same column names and column data types as the **employee** table, you can define the data as follows:

```
my_employee employee%ROWTYPE
```

The attribute can also be used on the cursor. The column names and column data types of this set of data are the same as those of the **employee** table. For the cursor in a package, **%ROWTYPE** can be omitted. **%TYPE** can also reference the type of a column in the cursor. You can define the data as follows:

```
cursor cur is select * from employee;
my_employee cur%ROWTYPE
my_name cur.firstname%TYPE
my_employee2 cur -- For the cursor defined in a package, %ROWTYPE can be omitted.
```

#### NOTICE

- **%TYPE** cannot reference the type of a composite variable or a record variable, a column type of the record type, a column type of a variable of the cross-package composite type, or a column type of a cursor variable of the cross-package type.
- **%ROWTYPE** cannot reference the type of a composite variable or a record variable and the type of a cross-package cursor.

## Scope of a Variable

The scope of a variable indicates the accessibility and availability of the variable in code block. In other words, a variable takes effect only within its scope.


- To define a function scope, a variable must declare and create a **BEGIN-END** block in the declaration section. The necessity of such declaration is also determined by block structure, which requires that a variable has different scopes and lifetime during a process.
- A variable can be defined multiple times in different scopes, and inner definition can cover outer one.
- A variable defined in an outer block can also be used in a nested block. However, the outer block cannot access variables in the nested block.

## 10.6.2 Assignment Statements

### Syntax

**Figure 10-4** shows the syntax diagram for assigning a value to a variable.

**Figure 10-4** assignment\_value::=



The above syntax diagram is explained as follows:

- **variable\_name** indicates the name of a variable.
- **value** can be a value or an expression. The type of *value* must be compatible with the type of *variable\_name*.


Example:

```
openGauss=# DECLARE
emp_id INTEGER := 7788; -- Assignment
BEGIN
emp_id := 5; -- Assignment
emp_id := 5*7784;
END;
/
```

## Nested Value Assignment

**Figure 10-5** shows the syntax diagram for assigning a nested value to a variable.

**Figure 10-5** nested\_assignment\_value::=



The syntax in **Figure 10-5** is described as follows:

- **variable\_name**: variable name
- **col\_name**: column name
- **subscript**: index, which is used for an array variable. The value can be a value or an expression and must be of the int type.
- **value**: value or expression. The type of *value* must be compatible with the type of *variable\_name*.

Example:

```
openGauss=#CREATE TYPE o1 as (a int, b int);
openGauss=# DECLARE
 TYPE r1 is VARRAY(10) of o1;
 emp_id r1;
BEGIN
 emp_id(1).a := 5;-- Assign a value.
 emp_id(1).b := 5*7784;
END;
/
```

#### NOTICE

- In INTO mode, values can be assigned only to the columns at the first layer. Two-dimensional or above arrays are not supported.
- When a nested column value is referenced, if an array index exists, only one parenthesis can exist in the first three layers of columns. You are advised to use square brackets to reference the index.

## INTO/BULK COLLECT INTO

Values returned by statements in a stored procedure are stored in variables. BULK COLLECT INTO allows some or all returned values to be temporarily stored in an array.

Example:

```
openGauss=# DROP TABLE IF EXISTS customers;
openGauss=# CREATE TABLE customers(id int,name varchar);
openGauss=# INSERT INTO customers VALUES(1,'ab');
openGauss=# DECLARE
 my_id integer;
BEGIN
 select id into my_id from customers limit 1; -- Assign a value.
END;
/

openGauss=# DECLARE
 type id_list is varray(6) of customers.id%type;
 id_arr id_list;
BEGIN
 select id bulk collect into id_arr from customers order by id DESC limit 20; -- Assign values in batches.
END;
/
```

#### NOTICE


BULK COLLECT INTO can only assign values to arrays or collections in batches. Use LIMIT properly to prevent performance deterioration caused by excessive operations on data.

## 10.6.3 Call Statements

### Syntax

Figure 10-6 shows the syntax diagram for calling a clause.

Figure 10-6 call\_clause::=



The above syntax diagram is explained as follows:

- **procedure\_name** specifies the name of a stored procedure.
- **parameter** specifies the parameters for the stored procedure. You can set no parameter or multiple parameters.

### Examples

```
-- Create a table.
openGauss=# CREATE SCHEMA hr;
openGauss=# SET CURRENT_SCHEMA = hr;
openGauss=# CREATE TABLE staffs
(
 section_id INTEGER,
 salary INTEGER
);
openGauss=# INSERT INTO staffs VALUES (30, 10);
openGauss=# INSERT INTO staffs VALUES (30, 20);

-- Create the stored procedure proc_staffs.
openGauss=# CREATE OR REPLACE PROCEDURE proc_staffs
(
 section NUMBER(6),
 salary_sum out NUMBER(8,2),
 staffs_count out INTEGER
)
IS
BEGIN
SELECT sum(salary), count(*) INTO salary_sum, staffs_count FROM hr.staffs where section_id = section;
END;
/

-- Create the stored procedure proc_return.
openGauss=# CREATE OR REPLACE PROCEDURE proc_return
AS
v_num NUMBER(8,2);
v_sum INTEGER;
BEGIN
proc_staffs(30, v_sum, v_num); -- Call a statement.
dbe_output.print_line(v_sum||'#'||v_num);
RETURN; -- Return a statement.
END;
/

-- Call the stored procedure proc_return.
openGauss=# CALL proc_staffs(2,8,6);

-- Drop the stored procedure.
```



```
openGauss=# DROP PROCEDURE proc_staffs;
openGauss=# DROP PROCEDURE proc_return;

-- Create the function func_return.
openGauss=# CREATE OR REPLACE FUNCTION func_return returns void
language plpgsql
AS $$
DECLARE
v_num INTEGER := 1;
BEGIN
dbs_output.print_line(v_num);
RETURN; -- Return a statement.
END $$;

-- Call the function func_return.
openGauss=# CALL func_return();

-- Drop the function.
openGauss=# DROP FUNCTION func_return;

-- Drop the current database schema.
openGauss=# DROP SCHEMA hr CASCADE;
```

## 10.7 Dynamic Statements

### 10.7.1 Executing Dynamic Query Statements

You can perform dynamic queries GaussDB provides two modes: EXECUTE IMMEDIATE and OPEN FOR. The former dynamically executes SELECT statements while the latter combines the use of cursors. If you need to store query results in a dataset, use **OPEN FOR**.

#### EXECUTE IMMEDIATE

Figure 10-7 shows the syntax diagram.

Figure 10-7 EXECUTE IMMEDIATE dynamic\_select\_clause::=





Figure 10-8 shows the syntax diagram for **using\_clause**.

**Figure 10-8** using\_clause::=

The above syntax diagram is explained as follows:

- *define\_variable*: specifies a variable to store single-row query results.
- **USING IN** *bind\_argument*: specifies the variable whose value is passed to the dynamic SQL statement. The variable is used when a dynamic placeholder exists in *dynamic\_select\_string*.
- **USING OUT** *bind\_argument*: specifies the variable that stores a value returned by the dynamic SQL statement.


#### NOTICE

- In query statements, **INTO** and **OUT** cannot coexist.
- A placeholder name starts with a colon (:) followed by digits, characters, or strings, corresponding to *bind\_argument* in the **USING** clause.
- *bind\_argument* can only be a value, variable, or expression. It cannot be a database object such as a table name, column name, and data type. That is, *bind\_argument* cannot be used to transfer schema objects for dynamic SQL statements. If a stored procedure needs to transfer database objects by declaring the parameter to construct dynamic SQL statements (generally, DDL statements), you are advised to use double vertical bars (||) to concatenate *dynamic\_select\_clause* with a database object.
- A dynamic PL/SQL block allows duplicate placeholders. That is, a placeholder can correspond to only one *bind\_argument* in the USING clause.

## OPEN FOR

Dynamic queries can be executed by using **OPEN FOR** to open dynamic cursors.

**Figure 10-9** shows the syntax diagram.

**Figure 10-9** open\_for::=

Parameter description:

- **cursor\_name**: specifies the name of the cursor to be opened.
- **dynamic\_string**: specifies the dynamic query statement.
- **USING value**: applies when a placeholder exists in `dynamic_string`.

For the use of cursors, see [Cursors](#).

## 10.7.2 Executing Dynamic Non-Query Statements

### Syntax

Figure 10-10 shows the syntax diagram.

Figure 10-10 `noselect::=`





Figure 10-11 shows the syntax diagram for `using_clause`.

Figure 10-11 `using_clause::=`



The above syntax diagram is explained as follows:

**USING IN** *bind\_argument* is used to specify the variable whose value is passed to the dynamic SQL statement. The variable is used when a placeholder exists in *dynamic\_noselect\_string*. That is, a placeholder is replaced by the corresponding *bind\_argument* when a dynamic SQL statement is executed. Note that *bind\_argument* can only be a value, variable, or expression, and cannot be a database object such as a table name, column name, and data type. If a stored procedure needs to transfer database objects by declaring the parameter to construct dynamic SQL statements (generally, DDL statements), you are advised to use double vertical bars (||) to concatenate *dynamic\_select\_clause* with a database object. In addition, duplicate placeholders are allowed in dynamic statements. A placeholder can correspond to only one *bind\_argument* by position.

### Examples

```
-- Create a table.
openGauss=# CREATE TABLE sections_t1
(
 section NUMBER(4) ,
 section_name VARCHAR2(30),
 manager_id NUMBER(6),
 place_id NUMBER(4)
);
```

```
-- Declare a variable.
openGauss=# DECLARE
section NUMBER(4) := 280;
section_name VARCHAR2(30) := 'Info support';
manager_id NUMBER(6) := 103;
place_id NUMBER(4) := 1400;
new_colname VARCHAR2(10) := 'sec_name';
BEGIN
-- Execute the query.
EXECUTE IMMEDIATE 'insert into sections_t1 values(:1, :2, :3, :4)'
 USING section, section_name, manager_id, place_id;
-- Execute the query (duplicate placeholders).
EXECUTE IMMEDIATE 'insert into sections_t1 values(:1, :2, :3, :1)'
 USING section, section_name, manager_id;
-- Run the ALTER statement. (You are advised to use double vertical bars (||) to concatenate the dynamic
DDL statement with a database object.)
EXECUTE IMMEDIATE 'alter table sections_t1 rename section_name to ' || new_colname;
END;
/

-- Query data.
openGauss=# SELECT * FROM sections_t1;

-- Drop the table.
openGauss=# DROP TABLE sections_t1;
```

### 10.7.3 Dynamically Calling Stored Procedures

This section describes how to dynamically call store procedures. You must use anonymous statement blocks to package stored procedures or statement blocks and append **IN** and **OUT** behind the **EXECUTE IMMEDIATE...USING** statement to input and output parameters.

#### Syntax

Figure 10-12 shows the syntax diagram.

Figure 10-12 call\_procedure::=





Figure 10-13 shows the syntax diagram for **using\_clause**.

**Figure 10-13** using\_clause::=

The above syntax diagram is explained as follows:

- **CALL procedure\_name:** calls the stored procedure.
- **[;placeholder1;;placeholder2,...]:** specifies the placeholder list of the stored procedure parameters. The numbers of the placeholders and parameters are the same.
- **USING [IN|OUT|IN OUT]bind\_argument:** specifies where the variable passed to the stored procedure parameter value is stored. The modifiers in front of **bind\_argument** and of the corresponding parameter are the same.

## Example

```
-- Create the stored procedure proc_add.
openGauss=# CREATE OR REPLACE PROCEDURE proc_add
(
 param1 in INTEGER,
 param2 out INTEGER,
 param3 in INTEGER
)
AS
BEGIN
 param2:= param1 + param3;
END;
/

openGauss=# DECLARE
 input1 INTEGER:=1;
 input2 INTEGER:=2;
 statement VARCHAR2(200);
 param2 INTEGER;
BEGIN
 -- Declare the call statement.
 statement := 'call proc_add(:col_1, :col_2, :col_3)';
 -- Execute the statement.
 EXECUTE IMMEDIATE statement
 USING IN input1, OUT param2, IN input2;
 db_output.print_line('result is: '||to_char(param2));
END;
/

-- Drop the stored procedure.
openGauss=# DROP PROCEDURE proc_add;
```

## 10.7.4 Dynamically Calling Anonymous Blocks

This section describes how to execute anonymous blocks in dynamic statements. Append **IN** and **OUT** behind the **EXECUTE IMMEDIATE...USING** statement to input and output parameters.

## Syntax

Figure 10-14 shows the syntax diagram.

Figure 10-14 call\_anonymous\_block::=





Figure 10-15 shows the syntax diagram for using\_clause.

Figure 10-15 using\_clause::=



The above syntax diagram is explained as follows:

- The execute part of an anonymous block starts with a **BEGIN** statement, has a break with an **END** statement, and ends with a semicolon (;).
- **USING [IN|OUT|IN OUT]bind\_argument**: specifies where the variable passed to the stored procedure parameter value is stored. The modifiers in front of **bind\_argument** and of the corresponding parameter are the same.
- The input and output parameters in the middle of an anonymous block are designated by placeholders. The numbers of the placeholders and parameters are the same. The sequences of the parameters corresponding to the placeholders and the USING parameters are the same.
- Currently in GaussDB, when dynamic statements call anonymous blocks, placeholders cannot be used to pass input and output parameters in an EXCEPTION statement.
- Overloaded functions with placeholders cannot be called.
- Variables declared in an anonymous block and binding parameters cannot be used in the same statement at the same time.

- Only SQL statements in anonymous blocks can be called to bind parameters. Parameters cannot be bound in other scenarios. For example, a stored procedure is called in an anonymous block, expressions and cursors are used in an anonymous block, and dynamic statements are called in an anonymous block.

## Examples

```
-- Create a table.
openGauss=# CREATE SCHEMA hr;
openGauss=# SET CURRENT_SCHEMA = hr;
-- Create the stored procedure dynamic_proc.
openGauss=# CREATE OR REPLACE PROCEDURE dynamic_proc
AS
 staff_id NUMBER(6) := 200;
 first_name VARCHAR2(20);
 salary NUMBER(8,2);
BEGIN
 -- Execute the anonymous block.
 EXECUTE IMMEDIATE 'begin select first_name, salary into :first_name, :salary from hr.staffs where
staff_id= :dno; end;'
 USING OUT first_name, OUT salary, IN staff_id;
 db_output.print_line(first_name|| ' ' || salary);
END;
/

-- Call the stored procedure.
openGauss=# CALL dynamic_proc();

-- Drop the stored procedure.
openGauss=# DROP PROCEDURE dynamic_proc;

-- Drop the current database schema.
openGauss=# SET CURRENT_SCHEMA = public;
openGauss=# DROP SCHEMA hr CASCADE;
```

## 10.8 Control Statements

### 10.8.1 RETURN Statements


In GaussDB, data can be returned in either of the following ways: **RETURN**, **RETURN NEXT**, or **RETURN QUERY**. **RETURN NEXT** and **RETURN QUERY** are used only for functions and cannot be used for stored procedures.

#### 10.8.1.1 RETURN

##### Syntax

Figure 10-16 shows the syntax diagram for a return statement.

Figure 10-16 return\_clause::=



The above syntax diagram is explained as follows:

This statement returns control from a stored procedure or function to a caller.

## Examples

See [Examples](#) for call statement examples.

### 10.8.1.2 RETURN NEXT and RETURN QUERY


#### Syntax

When creating a function, specify **SETOF** *datatype* for the return values.

return\_next\_clause::=



return\_query\_clause::=



The above syntax diagram is explained as follows:

If a function needs to return a result set, use **RETURN NEXT** or **RETURN QUERY** to add results to the result set, and then continue to execute the next statement of the function. As the **RETURN NEXT** or **RETURN QUERY** statement is executed repeatedly, more and more results will be added to the result set. After the function is executed, all results are returned.

**RETURN NEXT** can be used for scalar and compound data types.

**RETURN QUERY** has a variant **RETURN QUERY EXECUTE**. You can add dynamic queries and add parameters to the queries by **USING**.

## Examples

```
openGauss=# CREATE TABLE t1(a int);
openGauss=# INSERT INTO t1 VALUES(1),(10);

--RETURN NEXT
openGauss=# CREATE OR REPLACE FUNCTION fun_for_return_next() RETURNS SETOF t1 AS $$
DECLARE
 r t1%ROWTYPE;
BEGIN
 FOR r IN select * from t1
 LOOP
 RETURN NEXT r;
 END LOOP;
 RETURN;
END;
$$ LANGUAGE plpgsql;
openGauss=# call fun_for_return_next();
 a

 1
10
(2 rows)

-- RETURN QUERY
openGauss=# CREATE OR REPLACE FUNCTION fun_for_return_query() RETURNS SETOF t1 AS $$
DECLARE
 r t1%ROWTYPE;
```



```
BEGIN
 RETURN QUERY select * from t1;
END;
$$
language plpgsql;
openGauss=# call fun_for_return_query();
a

1
10
(2 rows)
```


## 10.8.2 Conditional Statements

Conditional statements are used to decide whether given conditions are met. Operations are executed based on the decisions made.

GaussDB supports five usages of **IF**:

- **IF\_THEN**

**Figure 10-17 IF\_THEN::=**




**IF\_THEN** is the simplest form of **IF**. If the condition is true, statements are executed. If it is false, they are skipped.

Example:

```
openGauss=#
DECLARE
 v_user_id integer default 1;
BEGIN
 IF v_user_id <> 0 THEN
 raise info 'v_user_id is NOT 0';
 END IF;
END;
/
INFO: v_user_id is NOT 0
```

- **IF\_THEN\_ELSE**

**Figure 10-18 IF\_THEN\_ELSE::=**



**IF\_THEN\_ELSE** statements add **ELSE** branches and can be executed if the condition is false.

Example:

```
openGauss=#
DECLARE
 v_user_id integer default 1;
BEGIN
```

```
IF v_user_id <> 0 THEN
 raise info 'v_user_id is NOT 0';
ELSE
 raise info 'v_user_id is 0';
END IF;
END;
/
INFO: v_user_id is NOT 0
```

- **IF\_THEN\_ELSE IF**


**IF** statements can be nested in the following way:

```
openGauss=#
DECLARE
 v_user_id integer default 1;
BEGIN
 IF v_user_id = 0 THEN
 raise info 'v_user_id is 0';
 ELSE
 IF v_user_id > 0 THEN
 raise info 'v_user_id > 0';
 END IF;
 END IF;
END;
/
INFO: v_user_id > 0
```

Actually, this is a way of an **IF** statement nesting in the **ELSE** part of another **IF** statement. Therefore, an **END IF** statement is required for each nesting **IF** statement and another **END IF** statement is required to end the parent **IF-ELSE** statement. To set multiple options, use the following form:

- **IF\_THEN\_ELSEIF\_ELSE**

**Figure 10-19** IF\_THEN\_ELSEIF\_ELSE::=



Example:

```
DECLARE
 v_user_id integer default NULL;
BEGIN
 IF v_user_id = 0 THEN
 raise info 'v_user_id is 0';
 ELSIF v_user_id > 0 THEN
 raise info 'v_user_id > 0';
 ELSIF v_user_id < 0 THEN
 raise info 'v_user_id < 0';
 ELSE
 raise info 'v_user_id is NULL';
 END IF;
END;
/
INFO: v_user_id is NULL
```

- **IF\_THEN\_ELSEIF\_ELSE**

**ELSEIF** is an alias of **ELSIF**.

Example:

```
CREATE OR REPLACE PROCEDURE proc_control_structure(i in integer)
AS
BEGIN
 IF i > 0 THEN
 raise info 'i:% is greater than 0. ',i;
 ELSIF i < 0 THEN
 raise info 'i:% is smaller than 0. ',i;
 ELSE
 raise info 'i:% is equal to 0. ',i;
 END IF;
 RETURN;
END;
/

CALL proc_control_structure(3);


-- Delete the stored procedure.
DROP PROCEDURE proc_control_structure;
```

## 10.8.3 Loop Statements

### Simple LOOP Statements

#### Syntax diagram

Figure 10-20 loop::=



#### Example

```
CREATE OR REPLACE PROCEDURE proc_loop(i in integer, count out integer)
AS
BEGIN
 count:=0;
 LOOP
 IF count > i THEN
 raise info 'count is %. ', count;
 EXIT;
 ELSE
 count:=count+1;
 END IF;
 END LOOP;
END;
/


CALL proc_loop(10,5);
```

#### NOTICE

The loop must be exploited together with **EXIT**; otherwise, a dead loop occurs.

### WHILE\_LOOP Statements

#### Syntax diagram

**Figure 10-21** while\_loop::=

If the conditional expression is true, a series of statements in the WHILE statement are repeatedly executed and the condition is decided each time the loop body is executed.

### Example


```
CREATE TABLE integertable(c1 integer) ;
CREATE OR REPLACE PROCEDURE proc_while_loop(maxval in integer)
AS
 DECLARE
 i int :=1;
 BEGIN
 WHILE i < maxval LOOP
 INSERT INTO integertable VALUES(i);
 i:=i+1;
 END LOOP;
 END;
/

-- Call a function.
CALL proc_while_loop(10);

-- Delete the stored procedure and table.
DROP PROCEDURE proc_while_loop;
DROP TABLE integertable;
```

## FOR\_LOOP (*Integer variable*) Statement

### Syntax diagram

**Figure 10-22** for\_loop::=


### NOTE

- The variable *name* is automatically defined as the integer type and exists only in this loop. The value of *name* ranges from **lower\_bound** to **upper\_bound**.
- When the keyword **REVERSE** is used, the value of **lower\_bound** must be greater than or equal to that of **upper\_bound**; otherwise, the loop body is not executed.

## FOR\_LOOP Query Statements

### Syntax diagram

Figure 10-23 for\_loop\_query::=




### NOTE

The variable *target* is automatically defined, its type is the same as that in the query result, and it is valid only in this loop. The **target** value is the query result.

## FORALL Batch Query Statements

### Syntax diagram

Figure 10-24 forall::=



### NOTE

- The variable *index* is automatically defined as the integer type and exists only in this loop. The value of *index* falls between the value of **low\_bound** and the value of **upper\_bound**.
- If **SAVE EXCEPTIONS** is specified, exceptions occurred during DML execution in the loop body are saved in **SQL&BULK\_EXCEPTIONS** and an exception is thrown after the execution is complete. If there is no abnormal execution result in the loop, the loop will not be rolled back in the current subtransaction.

### Example

```

CREATE TABLE hdfs_t1 (
 title NUMBER(6),
 did VARCHAR2(20),
 data_period VARCHAR2(25),
 kind VARCHAR2(25),
 interval VARCHAR2(20),
 time DATE,
 isModified VARCHAR2(10)
);

```

```
INSERT INTO hdfs_t1 VALUES(8, 'Donald', 'OConnell', 'DOCONNEL', '650.507.9833', to_date('21-06-1999',
'dd-mm-yyyy'), 'SH_CLERK');

CREATE OR REPLACE PROCEDURE proc_forall()
AS
BEGIN
 FORALL i IN 100..120
 update hdfs_t1 set title = title + 100*i;
END;
/

-- Call a function.
CALL proc_forall();

-- Query the calling result of the stored procedure.
SELECT * FROM hdfs_t1 WHERE title BETWEEN 100 AND 120;

-- Delete the stored procedure and table.
DROP PROCEDURE proc_forall;
DROP TABLE hdfs_t1;
```

## 10.8.4 Branch Statements

### Syntax

Figure 10-25 shows the syntax diagram for a branch statement.

Figure 10-25 case\_when::=





Figure 10-26 shows the syntax diagram for when\_clause.

Figure 10-26 when\_clause::=



Parameters:

- *case\_expression*: specifies the variable or expression.
- *when\_expression*: specifies the constant or conditional expression.
- *statement*: specifies the statement to be executed.

## Examples

```
CREATE OR REPLACE PROCEDURE proc_case_branch(pi_result in integer, pi_return out integer)
AS
BEGIN
 CASE pi_result
 WHEN 1 THEN
 pi_return := 111;
 WHEN 2 THEN
 pi_return := 222;
 WHEN 3 THEN
 pi_return := 333;
 WHEN 6 THEN
 pi_return := 444;
 WHEN 7 THEN
 pi_return := 555;
 WHEN 8 THEN
 pi_return := 666;
 WHEN 9 THEN
 pi_return := 777;
 WHEN 10 THEN
 pi_return := 888;
 ELSE
 pi_return := 999;
 END CASE;
 raise info 'pi_return : %',pi_return ;
END;
/

CALL proc_case_branch(3,0);

-- Delete the stored procedure.
DROP PROCEDURE proc_case_branch;
```

## 10.8.5 NULL Statements

In PL/SQL programs, **NULL** statements are used to indicate "nothing should be done", equal to placeholders. They grant meanings to some statements and improve program readability.

### Syntax

The following shows example use of **NULL** statements.

```
DECLARE
...
BEGIN
...
 IF v_num IS NULL THEN
 NULL; -- No data needs to be processed.
 END IF;
END;
/
```

## 10.8.6 Error Trapping Statements

By default, any error occurring in a PL/SQL function aborts execution of the function, and indeed of the surrounding transaction as well. You can trap errors and restore from them by using a **BEGIN** block with an **EXCEPTION** clause. The syntax is an extension of the normal syntax for a **BEGIN** block:

```
[<<label>>]
[DECLARE
 declarations]
BEGIN
 statements
```

```
EXCEPTION
 WHEN condition [OR condition ...] THEN
 handler_statements
 [WHEN condition [OR condition ...] THEN
 handler_statements
 ...]
END;
```

If no error occurs, this form of block simply executes all the statements, and then control passes to the next statement after **END**. But if an error occurs within the statements, further processing of the statements is abandoned, and control passes to the **EXCEPTION** list. The list is searched for the first condition matching the error that occurred. If a match is found, the corresponding **handler\_statements** are executed, and then control passes to the next statement after **END**. If no match is found, the error propagates out as though the **EXCEPTION** clause were not there at all: Error codes can be used to catch other error codes of the same type.

The error can be caught by an enclosing block with **EXCEPTION**, or if there is none it aborts processing of the function.

The condition names can be any of those shown in SQL standard error codes. The special condition name **OTHERS** matches every error type except **QUERY\_CANCELED**.

If a new error occurs within the selected **handler\_statements**, it cannot be caught by this **EXCEPTION** clause, but is propagated out. A surrounding **EXCEPTION** clause could catch it.

When an error is caught by an **EXCEPTION** clause, the local variables of the PL/SQL function remain as they were when the error occurred, but all changes to persistent database state within the block are rolled back.

Example:

```
CREATE TABLE mytab(id INT,firstname VARCHAR(20),lastname VARCHAR(20)) ;

INSERT INTO mytab(firstname, lastname) VALUES('Tom', 'Jones');

CREATE FUNCTION fun_exp() RETURNS INT
AS $$
DECLARE
 x INT :=0;
 y INT;
BEGIN
 UPDATE mytab SET firstname = 'Joe' WHERE lastname = 'Jones';
 x := x + 1;
 y := x / 0;
EXCEPTION
 WHEN division_by_zero THEN
 RAISE NOTICE 'caught division_by_zero';
 RETURN x;
END;$$
LANGUAGE plpgsql;

call fun_exp();
NOTICE: caught division_by_zero
fun_exp

 1
(1 row)

select * from mytab;
 id | firstname | lastname
-----+-----+-----
```



```
| Tom | Jones
(1 row)
```

```
DROP FUNCTION fun_exp();
DROP TABLE mytab;
```

When control reaches the assignment to **y**, it will fail with a **division\_by\_zero** error. This will be caught by the **EXCEPTION** clause. The value returned in the **RETURN** statement will be the incremented value of **x**.

#### NOTE

A block containing an **EXCEPTION** clause is more expensive to enter and exit than a block without one. Therefore, do not use **EXCEPTION** without need.

In the following scenario, an exception cannot be caught, and the entire transaction rolls back. The threads of the nodes participating the stored procedure exit abnormally due to node failure and network fault, or the source data is inconsistent with that of the table structure of the target table during the COPY FROM operation.

#### Example: Exceptions with **UPDATE/INSERT**

This example uses exception handling to perform either **UPDATE** or **INSERT**, as appropriate:

```
CREATE TABLE db (a INT, b TEXT);

CREATE FUNCTION merge_db(key INT, data TEXT) RETURNS VOID AS
$$
BEGIN
 LOOP

 -- First try to update the key
 UPDATE db SET b = data WHERE a = key;
 IF found THEN
 RETURN;
 END IF;

 -- The key does not exist. Therefore, try to insert one. If someone else inserts the same key concurrently, the
 unique key may fail to be obtained.
 BEGIN
 INSERT INTO db(a,b) VALUES (key, data);
 RETURN;
 EXCEPTION WHEN unique_violation THEN
 -- Do nothing, and loop to try the UPDATE again.
 END;
 END LOOP;
 END;
$$
LANGUAGE plpgsql;

SELECT merge_db(1, 'david');
SELECT merge_db(1, 'dennis');


-- Drop the function and table.
DROP FUNCTION merge_db;
DROP TABLE db ;
```

## 10.8.7 GOTO Statements

A GOTO statement unconditionally transfers the control from the current statement to a labeled statement. The GOTO statement changes the execution logic. Therefore, use this statement only when necessary. Alternatively, you can use the EXCEPTION statement to handle issues in special scenarios. To run a GOTO statement, the labeled statement must be unique.

## Syntax

label declaration ::=



goto statement ::=



## Examples

```
openGauss=# CREATE OR REPLACE PROCEDURE GOTO_test()
AS
DECLARE
 v1 int;
BEGIN
 v1 := 0;
 LOOP
 EXIT WHEN v1 > 100;
 v1 := v1 + 2;
 if v1 > 25 THEN
 GOTO pos1;
 END IF;
 END LOOP;
<<pos1>>
v1 := v1 + 10;
raise info 'v1 is %.', v1;
END;
/

call GOTO_test();
```

## Constraints

Using GOTO statements has the following constraints:

- A GOTO statement does not allow multiple labeled statements even if the statements are in different blocks.

```
BEGIN
 GOTO pos1;
<<pos1>>
 SELECT * FROM ...
<<pos1>>
 UPDATE t1 SET ...
END;
```

- A GOTO statement cannot transfer control to the IF, CASE, or LOOP statement.

```
BEGIN
 GOTO pos1;
 IF valid THEN
 <<pos1>>
 SELECT * FROM ...
 END IF;
END;
```

- A GOTO statement cannot transfer control from one IF clause to another, or from one WHEN clause in the CASE statement to another.

```
BEGIN
 IF valid THEN
 GOTO pos1;
 SELECT * FROM ...
 ELSE
```

```
<<pos1>>
UPDATE t1 SET ...
END IF;
END;
```

- Transferring control from an outer block to an inner BEGIN-END block is not supported.

```
BEGIN
GOTO pos1;
BEGIN
<<pos1>>
UPDATE t1 SET ...
END;
END;
```

- Transferring control from an exception handler to the current BEGIN-END block is not supported. However, transferring control to the upper-layer BEGIN-END block is supported.

```
BEGIN
<<pos1>>
UPDATE t1 SET ...
EXCEPTION
WHEN condition THEN
GOTO pos1;
END;
```

- To branch to a position that does not have an executable statement, add the NULL statement.

```
DECLARE
done BOOLEAN;
BEGIN
FOR i IN 1..50 LOOP
IF done THEN
GOTO end_loop;
END IF;
<<end_loop>> -- not allowed unless an executable statement follows
NULL; -- add NULL statement to avoid error
END LOOP; -- raises an error without the previous NULL
END;
/
```

## 10.9 Transaction Management

Calling a stored procedure automatically starts a transaction. When the calling is complete, the transaction is automatically committed, or rolled back upon an exception. In addition to automatic transaction control, you can also use COMMIT/ROLLBACK to control transactions in stored procedures. Running the COMMIT/ROLLBACK commands in a stored procedure will commit or roll back the current transaction and automatically starts a new transaction. All subsequent operations will be performed in the new transaction.

A savepoint is a special mark inside a transaction. It allows all commands that are executed after it was established to be rolled back, restoring the transaction state to what it was at the time of the savepoint. In a stored procedure, you can use savepoints to manage transactions. Currently, you can create, roll back, and release savepoints. If a savepoint for rollback is used in a stored procedure, only the modification of the current transaction is rolled back. The execution process of the stored procedure is not changed, and the values of local variables in the stored procedure are not rolled back.

## Syntax

```
Define a savepoint.
SAVEPOINT savepoint_name;
Roll back a savepoint.
ROLLBACK TO [SAVEPOINT] savepoint_name;
Release a savepoint.
RELEASE [SAVEPOINT] savepoint_name;
```

## Usage Scenarios

The applicable contexts are as follows:

- COMMIT, ROLLBACK, and SAVEPOINT can be used in PL/SQL stored procedures.
- COMMIT, ROLLBACK, and SAVEPOINT can be used in stored procedures that contain EXCEPTION.
- COMMIT, ROLLBACK, and SAVEPOINT can be used in EXCEPTION statements of stored procedures.
- A stored procedure that contains COMMIT, ROLLBACK, or SAVEPOINT (which means the stored procedure is controlled by BEGIN, START, or END) can be called in a transaction block.
- A stored procedure that contains savepoints can be called in a subtransaction. That is, an externally defined savepoint is used in the stored procedure to roll back the transaction to the savepoint defined outside the stored procedure.
- A stored procedure is visible to a savepoint defined in the stored procedure. That is, the modification of the transaction can be rolled back to the savepoint defined in the stored procedure.
- COMMIT, ROLLBACK, and SAVEPOINT, as well as IF, FOR, CURSOR LOOP, and WHILE, can be called in most PL/SQL contexts and statements.
- A stored procedure or function that contains COMMIT, ROLLBACK, or SAVEPOINT can be called in the return values and simple expression calculation of stored procedures.

The following content can be committed or rolled back:

- DDL statements after COMMIT/ROLLBACK can be committed or rolled back.
- DML statements after COMMIT/ROLLBACK can be committed.
- GUC parameters in stored procedures can be committed or rolled back.

## Usage Restrictions

The unapplicable contexts are as follows:

- COMMIT, ROLLBACK, and SAVEPOINT cannot be called in stored procedures other than PL/SQL, such as PL/Python.
- COMMIT, ROLLBACK, SAVEPOINT and stored procedures that contain COMMIT, ROLLBACK, or SAVEPOINT cannot be called in functions.
- After SAVEPOINT is called in a transaction block, stored procedures that contain COMMIT/ROLLBACK cannot be called.
- Stored procedures that contain COMMIT, ROLLBACK, or SAVEPOINT cannot be called in TRIGGER.

- COMMIT, ROLLBACK, and SAVEPOINT cannot be called in EXECUTE statements.
- Stored procedures that contain COMMIT, ROLLBACK, or SAVEPOINT cannot be called in CURSOR statements.
- Stored procedures that contain IMMUTABLE or SHIPPABLE cannot call COMMIT, ROLLBACK, SAVEPOINT or another stored procedure that contain COMMIT, ROLLBACK, or SAVEPOINT.
- Stored procedures that contain COMMIT, ROLLBACK, or SAVEPOINT cannot be called in SQL statements other than SELECT PROC and CALL PROC.
- COMMIT, ROLLBACK, or SAVEPOINT cannot be called in a stored procedure whose header contains GUC parameters.
- COMMIT, ROLLBACK, or SAVEPOINT cannot be called in expressions or CURSOR and EXECUTE statements.
- Savepoints defined outside a stored procedure cannot be released in the stored procedure.
- An autonomous transaction and a stored procedure transaction are two independent transactions that cannot use the savepoints defined in each other.

The following content cannot be committed or rolled back:

- Variables declared or imported in stored procedures cannot be committed or rolled back.
- In stored procedures, GUC parameters that take effect only after a restart cannot be committed or rolled back.

## Examples

- Example 1: COMMIT and ROLLBACK can be used in PL/SQL stored procedures.

```
CREATE TABLE EXAMPLE1(COL1 INT);

CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE()
AS
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1(COL1) VALUES (i);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
 ROLLBACK;
 END IF;
 END LOOP;
END;
/

DROP PROCEDURE TRANSACTION_EXAMPLE;
DROP TABLE EXAMPLE1;
```

- Example 2:

COMMIT and ROLLBACK can be used in stored procedures that contain EXCEPTION.

COMMIT and ROLLBACK can be used in EXCEPTION statements of stored procedures.

DDL statements after COMMIT/ROLLBACK can be committed or rolled back.

```
CREATE OR REPLACE PROCEDURE TEST_COMMIT_INSERT_EXCEPTION_ROLLBACK()
AS
```

```
BEGIN
DROP TABLE IF EXISTS TEST_COMMIT;
CREATE TABLE TEST_COMMIT(A INT, B INT);
INSERT INTO TEST_COMMIT SELECT 1, 1;
COMMIT;
 CREATE TABLE TEST_ROLLBACK(A INT, B INT);
 RAISE EXCEPTION 'RAISE EXCEPTION AFTER COMMIT';
EXCEPTION
 WHEN OTHERS THEN
INSERT INTO TEST_COMMIT SELECT 2, 2;
ROLLBACK;
END;
/
```

- Example 3: A stored procedure that contains COMMIT or ROLLBACK (which means the stored procedure is controlled by BEGIN, START, or END) can be called in a transaction block.

```
BEGIN;
 CALL TEST_COMMIT_INSERT_EXCEPTION_ROLLBACK();
END;
DROP PROCEDURE TEST_COMMIT_INSERT_EXCEPTION_ROLLBACK;
```

- Example 4: COMMIT and ROLLBACK, as well as IF, FOR, CURSOR LOOP, and WHILE, can be called in most PL/SQL contexts and statements.

```
CREATE OR REPLACE PROCEDURE TEST_COMMIT2()
IS
BEGIN
 DROP TABLE IF EXISTS TEST_COMMIT;
 CREATE TABLE TEST_COMMIT(A INT);
 FOR I IN REVERSE 3..0 LOOP
INSERT INTO TEST_COMMIT SELECT I;
COMMIT;
 END LOOP;
 FOR I IN REVERSE 2..4 LOOP
UPDATE TEST_COMMIT SET A=I;
COMMIT;
 END LOOP;
EXCEPTION
 WHEN OTHERS THEN
INSERT INTO TEST_COMMIT SELECT 4;
COMMIT;
END;
/
DROP PROCEDURE TEST_COMMIT2;
```

- Example 5: Return values and simple expression calculation of stored procedures are supported.

```
CREATE OR REPLACE PROCEDURE exec_func3(RET_NUM OUT INT)
AS
BEGIN
 RET_NUM := 1+1;
COMMIT;
END;
/
CREATE OR REPLACE PROCEDURE exec_func4(ADD_NUM IN INT)
AS
SUM_NUM INT;
BEGIN
SUM_NUM := ADD_NUM + exec_func3();
COMMIT;
END;
/
DROP PROCEDURE exec_func3;
DROP PROCEDURE exec_func4;
```

- Example 6: GUC parameters in stored procedures can be rolled back to a commit.

Note: You are advised not to use the **enable\_force\_vector\_engine** GUC parameter.

```
SHOW explain_perf_mode;
SHOW enable_force_vector_engine;

CREATE OR REPLACE PROCEDURE GUC_ROLLBACK()
AS
BEGIN
 SET enable_force_vector_engine = on;
 COMMIT;
 SET explain_perf_mode TO pretty;
 ROLLBACK;
END;
/

call GUC_ROLLBACK();
SHOW explain_perf_mode;
SHOW enable_force_vector_engine;
SET enable_force_vector_engine = off;
DROP PROCEDURE GUC_ROLLBACK;
```

- Example 7: COMMIT, ROLLBACK, and stored procedures that contain COMMIT or ROLLBACK cannot be called in functions.

```
CREATE OR REPLACE FUNCTION FUNCTION_EXAMPLE1() RETURN INT
AS
EXP INT;
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1(col1) VALUES (i);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
 ROLLBACK;
 END IF;
 END LOOP;
 SELECT COUNT(*) FROM EXAMPLE1 INTO EXP;
 RETURN EXP;
END;
/
DROP FUNCTION FUNCTION_EXAMPLE1;
```

- Example 8: Stored procedures that contain COMMIT or ROLLBACK cannot be called in functions.

```
CREATE OR REPLACE FUNCTION FUNCTION_EXAMPLE2() RETURN INT
AS
EXP INT;
BEGIN
 -- transaction_example is a stored procedure and contains the COMMIT/ROLLBACK statement.
 CALL transaction_example();
 SELECT COUNT(*) FROM EXAMPLE1 INTO EXP;
 RETURN EXP;
END;
/
DROP FUNCTION FUNCTION_EXAMPLE2;
```

- Example 9: A TRIGGER stored procedure cannot contain COMMIT or ROLLBACK or call another stored procedure that contains COMMIT or ROLLBACK.

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE FUNCTION FUNCTION_TRI_EXAMPLE2() RETURN TRIGGER
AS
EXP INT;
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1(col1) VALUES (i);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
```

```
 ROLLBACK;
 END IF;
END LOOP;
SELECT COUNT(*) FROM EXAMPLE1 INTO EXP;
END;
/

CREATE TRIGGER TRIGGER_EXAMPLE AFTER DELETE ON EXAMPLE1
FOR EACH ROW EXECUTE PROCEDURE FUNCTION_TRI_EXAMPLE2();

DELETE FROM EXAMPLE1;
DROP TRIGGER TRIGGER_EXAMPLE ON EXAMPLE1;
DROP FUNCTION FUNCTION_TRI_EXAMPLE2;
DROP TABLE EXAMPLE1;
```

- Example 10: Stored procedures that contain IMMUTABLE or SHIPPABLE cannot call COMMIT, ROLLBACK, or another stored procedure that contains COMMIT or ROLLBACK.

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE1()
IMMUTABLE
AS
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1 (col1) VALUES (i);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
 ROLLBACK;
 END IF;
 END LOOP;
END;
/
DROP PROCEDURE TRANSACTION_EXAMPLE1;
DROP TABLE EXAMPLE1;
```

- Example 11: Variables declared or specified in stored procedures cannot be committed.

```
CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE2(EXP_OUT OUT INT)
AS
EXP INT;
BEGIN
 EXP_OUT := 0;
 COMMIT;
 DBE_OUTPUT.PRINT_LINE('EXP IS:'||EXP);
 EXP_OUT := 1;
 ROLLBACK;
 DBE_OUTPUT.PRINT_LINE('EXP IS:'||EXP);
END;
/
DROP PROCEDURE TRANSACTION_EXAMPLE2;
```

- Example 12: Calling in SQL statements (other than Select Procedure) is not supported.

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE3()
AS
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1 (col1) VALUES (i);
 IF i % 2 = 0 THEN
 EXECUTE IMMEDIATE 'COMMIT';
 ELSE
 EXECUTE IMMEDIATE 'ROLLBACK';
 END IF;
 END LOOP;
END;
/
```



```
DROP PROCEDURE TRANSACTION_EXAMPLE3;
DROP TABLE EXAMPLE1;
```

- **Example 13: COMMIT or ROLLBACK cannot be called in a stored procedure whose header contains GUC parameters.**

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE4()
SET ARRAY_NULLS TO "ON"
AS
BEGIN
 FOR i IN 0..20 LOOP
 INSERT INTO EXAMPLE1 (col1) VALUES (i);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
 ROLLBACK;
 END IF;
 END LOOP;
END;
/
DROP PROCEDURE TRANSACTION_EXAMPLE4;
DROP TABLE EXAMPLE1;
```

- **Example 14: A stored procedure whose cursor is open cannot contain COMMIT or ROLLBACK.**

```
CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE5(INTIN IN INT, INTOUT OUT INT)
AS
BEGIN
 INTOUT := INTIN + 1;
 COMMIT;
END;
/

CREATE OR REPLACE PROCEDURE TRANSACTION_EXAMPLE6()
AS
CURSOR CURSOR1(EXPIN INT)
IS SELECT TRANSACTION_EXAMPLE5(EXPIN);
INTEXP INT;
BEGIN
 FOR i IN 0..20 LOOP
 OPEN CURSOR1(i);
 FETCH CURSOR1 INTO INTEXP;
 INSERT INTO EXAMPLE1(COL1) VALUES (INTEXP);
 IF i % 2 = 0 THEN
 COMMIT;
 ELSE
 ROLLBACK;
 END IF;
 CLOSE CURSOR1;
 END LOOP;
END;
/

DROP PROCEDURE TRANSACTION_EXAMPLE5;
DROP PROCEDURE TRANSACTION_EXAMPLE6;
```

- **Example 15: COMMIT or ROLLBACK cannot be called in expressions or CURSOR and EXECUTE statements.**

```
CREATE OR REPLACE PROCEDURE exec_func1()
AS
BEGIN
 CREATE TABLE TEST_exec(A INT);
 COMMIT;
END;
/
CREATE OR REPLACE PROCEDURE exec_func2()
AS
BEGIN
 EXECUTE exec_func1();
 COMMIT;
```

```
END;
/
DROP PROCEDURE exec_func1;
DROP PROCEDURE exec_func2;
```

- **Example 16: Roll back some modifications of stored procedure on a transaction to a savepoint.**

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE STP_SAVEPOINT_EXAMPLE1()
AS
BEGIN
 INSERT INTO EXAMPLE1 VALUES(1);
 SAVEPOINT s1;
 INSERT INTO EXAMPLE1 VALUES(2);
 ROLLBACK TO s1; -- Roll back the insertion of record 2.
 INSERT INTO EXAMPLE1 VALUES(3);
END;
/
DROP PROCEDURE STP_SAVEPOINT_EXAMPLE1;
DROP TABLE EXAMPLE1;
```

- **Example 17: Roll back a stored procedure to a savepoint defined outside the stored procedure.**

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE STP_SAVEPOINT_EXAMPLE2()
AS
BEGIN
 INSERT INTO EXAMPLE1 VALUES(2);
 ROLLBACK TO s1; -- Roll back the insertion of record 2.
 INSERT INTO EXAMPLE1 VALUES(3);
END;
/

BEGIN;
INSERT INTO EXAMPLE1 VALUES(1);
SAVEPOINT s1;
CALL STP_SAVEPOINT_EXAMPLE2();
SELECT * FROM EXAMPLE1;
COMMIT;
DROP PROCEDURE STP_SAVEPOINT_EXAMPLE2;
DROP TABLE EXAMPLE1;
```

- **Example 18: Savepoints defined outside the stored procedure cannot be released in the stored procedure.**

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE STP_SAVEPOINT_EXAMPLE3()
AS
BEGIN
 INSERT INTO EXAMPLE1 VALUES(2);
 RELEASE SAVEPOINT s1; -- Release the savepoint defined outside the stored procedure.
 INSERT INTO EXAMPLE1 VALUES(3);
END;
/

BEGIN;
INSERT INTO EXAMPLE1 VALUES(1);
SAVEPOINT s1;
CALL STP_SAVEPOINT_EXAMPLE3();
COMMIT;
DROP PROCEDURE STP_SAVEPOINT_EXAMPLE3;
DROP TABLE EXAMPLE1;
```

- **Example 19: Roll back an external SQL or other stored procedure to a savepoint defined in the stored procedure.**

```
CREATE TABLE EXAMPLE1(COL1 INT);
CREATE OR REPLACE PROCEDURE STP_SAVEPOINT_EXAMPLE3()
AS
BEGIN
 INSERT INTO EXAMPLE1 VALUES(1);
```

```
SAVEPOINT s1;
INSERT INTO EXAMPLE1 VALUES(2);
END;
/

BEGIN;
INSERT INTO EXAMPLE1 VALUES(3);
CALL STP_SAVEPOINT_EXAMPLE3();
ROLLBACK TO SAVEPOINT s1; -- Roll back the insertion of record 2 to the stored procedure.
SELECT * FROM EXAMPLE1;
COMMIT;
DROP PROCEDURE STP_SAVEPOINT_EXAMPLE3;
DROP TABLE EXAMPLE1;
```

## 10.10 Other Statements

### 10.10.1 Lock Operations

GaussDB provides multiple lock modes to control concurrent accesses to table data. These modes are used when Multi-Version Concurrency Control (MVCC) cannot give expected behaviors. Alike, most GaussDB commands automatically apply appropriate locks to ensure that called tables are not deleted or modified in an incompatible manner during command execution. For example, when concurrent operations exist, ALTER TABLE cannot be executed on the same table.

For the complete lock operations, see [LOCK](#).

### 10.10.2 Cursor Operations

GaussDB provides cursors as a data buffer for users to store execution results of SQL statements. Each cursor region has a name. Users can use SQL statements to obtain records one by one from cursors and grant the records to master variables, then being processed further by host languages.

Cursor operations include cursor definition, open, fetch, and close operations.

For the complete example of cursor operations, see [Explicit Cursor](#).

## 10.11 Cursors

### 10.11.1 Overview

To process SQL statements, the stored procedure process assigns a memory segment to store context association. Cursors are handles or pointers pointing to context regions. With cursors, stored procedures can control alterations in context regions.

---

#### NOTICE

If JDBC is used to call a stored procedure whose returned value is a cursor, the returned cursor cannot be used.

---

Cursors are classified into explicit cursors and implicit cursors. [Table 10-2](#) shows the usage conditions of explicit and implicit cursors for different SQL statements.

**Table 10-2** Cursor usage conditions

| SQL Statement                             | Cursor               |
|-------------------------------------------|----------------------|
| Non-query statements                      | Implicit             |
| Query statements with single-line results | Implicit or explicit |
| Query statements with multi-line results  | Explicit             |

## 10.11.2 Explicit Cursor

An explicit cursor is used to process query statements, particularly when query results are multiple records.


### Procedure

An explicit cursor performs the following six PL/SQL steps to process query statements:

- Step 1** Define a static cursor: Define a cursor name and its corresponding **SELECT** statement.

[Figure 10-27](#) shows the syntax diagram for defining a static cursor.

**Figure 10-27** static\_cursor\_define::=



Parameter description:

- *cursor\_name*: defines a cursor name.
- *parameter*: specifies cursor parameters. Only input parameters are allowed. Its format is as follows:  
parameter\_name datatype
- *select\_statement*: specifies a query statement.

 NOTE

- The system automatically determines whether the cursor can be used for backward fetching based on the execution plan.
- In syntax, **parameter** can be an output parameter, but its behavior is the same as that of an input parameter.

Define a dynamic cursor: Define a **ref** cursor, which means that the cursor can be opened dynamically by a set of static SQL statements. Define the type of the **ref** cursor first, and then the cursor variable of this cursor type. Dynamically bind a **SELECT** statement through **OPEN FOR** when the cursor is opened.

**Figure 10-28** and **Figure 10-29** show the syntax diagrams for defining a dynamic cursor.

**Figure 10-28** cursor\_typename::=






**Figure 10-29** dynamic\_cursor\_define::=



**Step 2** Open the static cursor: Execute the **SELECT** statement corresponding to the cursor. The query result is placed in the workspace and the pointer directs to the head of the workspace to identify the cursor result set. If the cursor query statement carries the **FOR UPDATE** option, the **OPEN** statement locks the data rows corresponding to the cursor result set in the database table.

**Figure 10-30** shows the syntax diagram for opening a static cursor.


**Figure 10-30** open\_static\_cursor::=



Open the dynamic cursor: Use the **OPEN FOR** statement to open the dynamic cursor and the SQL statement is dynamically bound.

**Figure 10-31** shows the syntax diagrams for opening a dynamic cursor.

**Figure 10-31** open\_dynamic\_cursor::=




A PL/SQL program cannot use the **OPEN** statement to repeatedly open a cursor.

**Step 3** Fetch cursor data: Retrieve data rows in the result set and place them in specified output variables.

**Figure 10-32** shows the syntax diagrams for fetching cursor data.

**Figure 10-32** fetch\_cursor::=




**Step 4** Process the record.

**Step 5** Continue to process until the active set has no record.

**Step 6** Close the cursor: After you fetch and process the data in the cursor result set, close the cursor in time to release system resources used by the cursor and invalidate the workspace of the cursor so that the **FETCH** statement can no longer be used to fetch data. A closed cursor can be reopened by an **OPEN** statement.

**Figure 10-33** shows the syntax diagram for closing a cursor.

**Figure 10-33** close\_cursor::=



----End

## Attributes

Cursor attributes are used to control program procedures or know program status. When a DML statement is executed, the PL/SQL opens a built-in cursor and processes its result. A cursor is a memory segment for maintaining query results. It is opened when a DML statement is executed and closed when the execution is finished. An explicit cursor has the following attributes:

- **%FOUND**: Boolean attribute, which returns **TRUE** if the last fetch returns a row.
- **%NOTFOUND**: Boolean attribute, which works opposite to the **%FOUND** attribute.
- **%ISOPEN**: Boolean attribute, which returns **TRUE** if the cursor has been opened.
- **%ROWCOUNT**: numeric attribute, which returns the number of records fetched from the cursor.

## Examples

DDL and DML statements are prepared. Subsequent examples in this section depend on this case.

```
DROP SCHEMA IF EXISTS hr CASCADE;
CREATE SCHEMA hr;
SET current_schema = hr;
DROP TABLE IF EXISTS sections;
DROP TABLE IF EXISTS staffs;
DROP TABLE IF EXISTS department;
-- Create a department table.
CREATE TABLE sections(
 section_name varchar(100),
 place_id int,
 section_id int
);
INSERT INTO sections VALUES ('hr',1,1);

-- Create an employee table.
CREATE TABLE staffs(
 staff_id number(6),
 salary number(8,2),
 section_id int,
 first_name varchar(20)
);
INSERT INTO staffs VALUES (1,100,1,'Tom');

-- Create a department table.
CREATE TABLE department(
 section_id int
);
-- Specify the method for passing cursor parameters.
CREATE OR REPLACE PROCEDURE cursor_proc1()
AS
DECLARE
 DEPT_NAME VARCHAR(100);
 DEPT_LOC NUMBER(4);
 -- Define a cursor.
 CURSOR C1 IS
 SELECT section_name, place_id FROM hr.sections WHERE section_id <= 50;
 CURSOR C2(sect_id INTEGER) IS
 SELECT section_name, place_id FROM hr.sections WHERE section_id <= sect_id;
 TYPE CURSOR_TYPE IS REF CURSOR;
 C3 CURSOR_TYPE;
 SQL_STR VARCHAR(100);
BEGIN
 OPEN C1;-- Open the cursor.
 LOOP
 -- Fetch data from the cursor.
 FETCH C1 INTO DEPT_NAME, DEPT_LOC;
 EXIT WHEN C1%NOTFOUND;
 DBE_OUTPUT.PRINT_LINE(DEPT_NAME||'---'||DEPT_LOC);
 END LOOP;
 CLOSE C1;-- Close the cursor.

 OPEN C2(10);
 LOOP
 FETCH C2 INTO DEPT_NAME, DEPT_LOC;
 EXIT WHEN C2%NOTFOUND;
 DBE_OUTPUT.PRINT_LINE(DEPT_NAME||'---'||DEPT_LOC);
 END LOOP;
 CLOSE C2;

 SQL_STR := 'SELECT section_name, place_id FROM hr.sections WHERE section_id <= :DEPT_NO;';
 OPEN C3 FOR SQL_STR USING 50;
 LOOP
 FETCH C3 INTO DEPT_NAME, DEPT_LOC;
 EXIT WHEN C3%NOTFOUND;
 DBE_OUTPUT.PRINT_LINE(DEPT_NAME||'---'||DEPT_LOC);
 END LOOP;
 CLOSE C3;
END;
/
```

```
CREATE PROCEDURE
CALL cursor_proc1();

hr---1
hr---1
hr---1
cursor_proc1

(1 row)
DROP PROCEDURE cursor_proc1;
DROP PROCEDURE
-- Give a salary raise to employees whose salary is lower than 3000 by adding 500.
CREATE TABLE hr.staffs_t1 AS TABLE hr.staffs;
INSERT 0 1

CREATE OR REPLACE PROCEDURE cursor_proc2()
AS
DECLARE
 V_EMPNO NUMBER(6);
 V_SAL NUMBER(8,2);
 CURSOR C IS SELECT staff_id, salary FROM hr.staffs_t1;
BEGIN
 OPEN C;
 LOOP
 FETCH C INTO V_EMPNO, V_SAL;
 EXIT WHEN C%NOTFOUND;
 IF V_SAL<=3000 THEN
 UPDATE hr.staffs_t1 SET salary =salary + 500 WHERE staff_id = V_EMPNO;
 END IF;
 END LOOP;
 CLOSE C;
END;
/
CREATE PROCEDURE
CALL cursor_proc2();
cursor_proc2

(1 row)
-- Delete the stored procedure.
DROP PROCEDURE cursor_proc2;
DROP PROCEDURE
DROP TABLE hr.staffs_t1;
DROP TABLE
-- Use function parameters of the SYS_REFCURSOR type.
CREATE OR REPLACE PROCEDURE proc_sys_ref(O OUT SYS_REFCURSOR)
IS
C1 SYS_REFCURSOR;
BEGIN
 OPEN C1 FOR SELECT section_ID FROM HR.sections ORDER BY section_ID;
 O := C1;
END;
/

DECLARE
C1 SYS_REFCURSOR;
TEMP NUMBER(4);
BEGIN
proc_sys_ref(C1);
LOOP
 FETCH C1 INTO TEMP;
 DBE_OUTPUT.PRINT_LINE(C1%ROWCOUNT);
 EXIT WHEN C1%NOTFOUND;
END LOOP;
END;
/
```



```
-- Delete the stored procedure.
DROP PROCEDURE proc_sys_ref;
```

### 10.11.3 Implicit Cursor

Implicit cursors are automatically set by the system for non-query statements such as modify or delete operations, along with their workspace. Implicit cursors are named **SQL**, which is defined by the system.

#### Overview

Implicit cursor operations, such as definition, open, value-grant, and close operations, are automatically performed by the system and do not need users to process. Users can use only attributes related to implicit cursors to complete operations. In workspace of implicit cursors, the data of the latest SQL statement is stored and is not related to explicit cursors defined by users.

Format call: **SQL%**

#### NOTE

- INSERT, UPDATE, DELETE, and SELECT statements do not need defined cursors.
- In O-compatible mode, if the GUC parameter **behavior\_compat\_options** is set to **compat\_cursor**, implicit cursors are valid across stored procedures.

#### Attributes

An implicit cursor has the following attributes:

- **SQL%FOUND**: Boolean attribute, which returns **TRUE** if the last fetch returns a row.
- **SQL%NOTFOUND**: Boolean attribute, which works opposite to the **SQL%FOUND** attribute.
- **SQL%ROWCOUNT**: numeric attribute, which returns the number of records fetched from the cursor.
- **SQL%ISOPEN**: Boolean attribute, whose value is always **FALSE**. Close implicit cursors immediately after an SQL statement is run.

```
CREATE TABLE hr.staffs_t1 AS TABLE hr.staffs;
CREATE TABLE hr.sections_t1 AS TABLE hr.sections;
-- Delete all employees in a department from the hr.staffs table. If the department has no employees,
-- delete the department from the hr.sections table.
CREATE OR REPLACE PROCEDURE proc_cursor3()
AS
 DECLARE
 V_DEPTNO NUMBER(4) := 100;
 BEGIN
 DELETE FROM hr.staffs WHERE section_ID = V_DEPTNO;
 -- Proceed based on cursor status.
 IF SQL%NOTFOUND THEN
 DELETE FROM hr.sections_t1 WHERE section_ID = V_DEPTNO;
 END IF;
 END;
/

CALL proc_cursor3();

-- Delete the stored procedure and the temporary table.
DROP PROCEDURE proc_cursor3;
```

```
DROP TABLE hr.staffs_t1;
DROP TABLE hr.sections_t1;
```


## 10.11.4 Cursor Loop

Use of cursors in WHILE and LOOP statements is called a cursor loop. Generally, OPEN, FETCH, and CLOSE statements are called in this kind of loop. The following describes a loop that simplifies a cursor loop without the need for these operations. This mode is applicable to a static cursor loop, without executing four steps about a static cursor.

### Syntax

Figure 10-34 shows the syntax diagram of the FOR AS loop.

Figure 10-34 FOR\_AS\_loop::=



### Precautions

- The UPDATE operation for the queried table is not allowed in the loop statement.
- The variable *loop\_name* is automatically defined and is valid only in this loop. Its type is the same as that in the query result of *select\_statement*. The value of *loop\_name* is the query result of *select\_statement*.
- The **%FOUND**, **%NOTFOUND**, and **%ROWCOUNT** attributes access the same internal variable in GaussDB. Transactions and the anonymous block do not support multiple cursor accesses at the same time.
- The specific type of the **loop\_name** variable is not parsed during compilation. If the specific type needs to be parsed (for example, **loop\_name** is used as the input and output parameters of an overloaded function or stored procedure), a compilation error is reported. To parse the specific type of a variable, set **behavior\_compat\_options** to 'allow\_procedure\_compile\_check'.

### Examples

```
BEGIN
FOR ROW_TRANS IN
 SELECT first_name FROM hr.staffs
LOOP
 DBE_OUTPUT.PRINT_LINE (ROW_TRANS.first_name);
END LOOP;
END;
/

-- Create a table.
CREATE TABLE integerTable1(A INTEGER);
CREATE TABLE integerTable2(B INTEGER);
INSERT INTO integerTable2 VALUES(2);
```

```
-- Multiple cursors share the parameters of cursor attributes.
DECLARE
 CURSOR C1 IS SELECT A FROM integerTable1;-- Declare the cursor.
 CURSOR C2 IS SELECT B FROM integerTable2;
 PI_A INTEGER;
 PI_B INTEGER;
BEGIN
 OPEN C1;-- Open the cursor.
 OPEN C2;
 FETCH C1 INTO PI_A; ---- The values of C1%FOUND and C2%FOUND are FALSE.
 FETCH C2 INTO PI_B; ---- The values of C1%FOUND and C2%FOUND are TRUE.
 -- Determine the cursor status.
 IF C1%FOUND THEN
 IF C2%FOUND THEN
 DBE_OUTPUT.PRINT_LINE('Dual cursor share paremeter.');
```

## 10.12 Advanced Packages

Advanced packages have two sets of APIs. The first set is basic APIs, and the second set is secondary encapsulation APIs that are used improve usability. The second set is recommended.

### 10.12.1 Basic APIs

#### 10.12.1.1 PKG\_SERVICE

**Table 10-3** lists all APIs supported by the **PKG\_SERVICE** package.

**Table 10-3** PKG\_SERVICE

| API                                                | Description                             |
|----------------------------------------------------|-----------------------------------------|
| <a href="#">PKG_SERVICE.SQL_IS_CONTEXT_ACTIVE</a>  | Checks whether a context is registered. |
| <a href="#">PKG_SERVICE.SQL_CLEAN_ALL_CONTEXTS</a> | Deregisters all registered contexts.    |
| <a href="#">PKG_SERVICE.SQL_REGISTER_CONTEXT</a>   | Registers a context.                    |
| <a href="#">PKG_SERVICE.SQL_UNREGISTER_CONTEXT</a> | Deregisters a context.                  |

| API                                                 | Description                                                                                                                      |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">PKG_SERVICE.SQL_SET_SQL</a>             | Sets an SQL statement for a context. Currently, only the SELECT statement is supported.                                          |
| <a href="#">PKG_SERVICE.SQL_RUN</a>                 | Executes the configured SQL statement on a context.                                                                              |
| <a href="#">PKG_SERVICE.SQL_NEXT_ROW</a>            | Reads the next row of data in a context.                                                                                         |
| <a href="#">PKG_SERVICE.SQL_GET_VALUE</a>           | Reads a dynamically defined column value in a context.                                                                           |
| <a href="#">PKG_SERVICE.SQL_SET_RESULT_TYPE</a>     | Dynamically defines a column of a context based on the type OID.                                                                 |
| <a href="#">PKG_SERVICE.JOB_CANCEL</a>              | Removes a scheduled task by task ID.                                                                                             |
| <a href="#">PKG_SERVICE.JOB_FINISH</a>              | Disables or enables scheduled task execution.                                                                                    |
| <a href="#">PKG_SERVICE.JOB_SUBMIT</a>              | Submits a scheduled task. Job ID can be automatically generated by the system or specified manually.                             |
| <a href="#">PKG_SERVICE.JOB_UPDATE</a>              | Modifies user-definable attributes of a scheduled task, including the task content, next-execution time, and execution interval. |
| <a href="#">PKG_SERVICE.SUBMIT_ON_NODES</a>         | Submits a task to all nodes. The task ID is automatically generated by the system.                                               |
| <a href="#">PKG_SERVICE.ISUBMIT_ON_NODES</a>        | Submits a job to all nodes. The job ID is specified by the user.                                                                 |
| <a href="#">PKG_SERVICE.SQL_GET_ARRAY_RESULT</a>    | Obtains the array value returned in the context.                                                                                 |
| <a href="#">PKG_SERVICE.SQL_GET_VARIABLE_RESULT</a> | Obtains the column value returned in the context.                                                                                |

- [PKG\\_SERVICE.SQL\\_IS\\_CONTEXT\\_ACTIVE](#)

This function checks whether a context is registered. This function transfers the ID of the context to be queried. If the context exists, **TRUE** is returned. Otherwise, **FALSE** is returned.

The prototype of the [PKG\\_SERVICE.SQL\\_IS\\_CONTEXT\\_ACTIVE](#) function is as follows:

```
PKG_SERVICE.SQL_IS_CONTEXT_ACTIVE(
 context_id IN INTEGER
)
RETURN BOOLEAN;
```

**Table 10-4** PKG\_SERVICE.SQL\_IS\_CONTEXT\_ACTIVE API parameters

| Parameter  | Description                      |
|------------|----------------------------------|
| context_id | ID of the context to be queried. |

- PKG\_SERVICE.SQL\_CLEAN\_ALL\_CONTEXTS

This function cancels all contexts.

The prototype of the PKG\_SERVICE.SQL\_CLEAN\_ALL\_CONTEXTS function is as follows:

```
PKG_SERVICE.SQL_CLEAN_ALL_CONTEXTS(
)
RETURN VOID;
```

- PKG\_SERVICE.SQL\_REGISTER\_CONTEXT

This function opens a context, which is the prerequisite for the subsequent operations in the context. This function does not transfer any parameter. It automatically generates context IDs in an ascending order and returns values to integer variables.

The prototype of the PKG\_SERVICE.SQL\_REGISTER\_CONTEXT function is as follows:

```
DBE_SQL.REGISTER_CONTEXT(
)
RETURN INTEGER;
```

- PKG\_SERVICE.SQL\_UNREGISTER\_CONTEXT

This function closes a context, which is the end of each operation in the context. If this function is not called when the stored procedure ends, the memory is still occupied by the context. Therefore, remember to close a context when you do not need to use it. If an exception occurs, the stored procedure exits but the context is not closed. Therefore, you are advised to include this API in the exception handling of the stored procedure.

The prototype of the PKG\_SERVICE.SQL\_UNREGISTER\_CONTEXT function is as follows:

```
PKG_SERVICE.SQL_UNREGISTER_CONTEXT(
context_id IN INTEGER
)
RETURN INTEGER;
```

**Table 10-5** PKG\_SERVICE.SQL\_UNREGISTER\_CONTEXT API parameter

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the context to be closed. |

- PKG\_SERVICE.SQL\_SET\_SQL

This function parses the query statement of a given context. The input query statement is executed immediately. Currently, only the SELECT query statement can be parsed. The statement parameters can be specified only through the TEXT type. The length cannot exceed 1 GB.

The prototype of the PKG\_SERVICE.SQL\_SET\_SQL function is as follows:

```
PKG_SERVICE.SQL_SET_SQL(
context_id IN INTEGER,
```

```
query_string IN TEXT,
language_flag IN INTEGER
)
RETURN BOOLEAN;
```

**Table 10-6** PKG\_SERVICE.SQL\_SET\_SQL API parameters

| Parameter     | Description                                                     |
|---------------|-----------------------------------------------------------------|
| context_id    | ID of the context whose query statement is to be parsed.        |
| query_string  | Query statement to be parsed.                                   |
| language_flag | Version language number. Currently, only <b>1</b> is supported. |

- **PKG\_SERVICE.SQL\_RUN**

This function executes a given context. It receives a context ID first, and the data obtained after execution is used for subsequent operations. Currently, only the **SELECT** query statement can be executed.

The prototype of the PKG\_SERVICE.SQL\_RUN function is as follows:

```
PKG_SERVICE.SQL_RUN(
context_id IN INTEGER,
)
RETURN INTEGER;
```

**Table 10-7** PKG\_SERVICE.SQL\_RUN API parameters

| Parameter  | Description                                              |
|------------|----------------------------------------------------------|
| context_id | ID of the context whose query statement is to be parsed. |

- **PKG\_SERVICE.SQL\_NEXT\_ROW**

This function returns the number of data rows returned after the SQL statement is executed. Each time the API is executed, the system obtains a set of new rows until all data is read.

The prototype of the PKG\_SERVICE.SQL\_NEXT\_ROW function is as follows:

```
PKG_SERVICE.SQL_NEXT_ROW(
context_id IN INTEGER,
)
RETURN INTEGER;
```

**Table 10-8** PKG\_SERVICE.SQL\_NEXT\_ROW API parameters

| Parameter  | Description                       |
|------------|-----------------------------------|
| context_id | ID of the context to be executed. |

- **PKG\_SERVICE.SQL\_GET\_VALUE**

This function returns the context element value in a specified position of a context and accesses the data obtained by PKG\_SERVICE.SQL\_NEXT\_ROW.

The prototype of the PKG\_SERVICE.SQL\_GET\_VALUE function is as follows:

```
PKG_SERVICE.SQL_GET_VALUE(
context_id IN INTEGER,
pos IN INTEGER,
col_type IN ANYELEMENT
)
RETURN ANYELEMENT;
```

**Table 10-9** PKG\_SERVICE.SQL\_GET\_VALUE API parameters

| Parameter  | Description                                                           |
|------------|-----------------------------------------------------------------------|
| context_id | ID of the context to be executed.                                     |
| pos        | Position of a dynamically defined column in the query.                |
| col_type   | Variable of any type, which defines the return value type of columns. |

- PKG\_SERVICE.SQL\_SET\_RESULT\_TYPE

This function defines columns returned from a given context and can be used only for contexts defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The prototype of

PKG\_SERVICE.SQL\_SET\_RESULT\_TYPE is as follows:

```
PKG_SERVICE.SQL_SET_RESULT_TYPE(
context_id IN INTEGER,
pos IN INTEGER,
coltype_oid IN ANYELEMENT,
maxsize IN INTEGER
)
RETURN INTEGER;
```

**Table 10-10** PKG\_SERVICE.SQL\_SET\_RESULT\_TYPE API parameters

| Parameter   | Description                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------|
| context_id  | ID of the context to be executed                                                                    |
| pos         | Position of a dynamically defined column in the query                                               |
| coltype_oid | Variable of any type. The OID of the corresponding type can be obtained based on the variable type. |
| maxsize     | Length of a defined column                                                                          |

- PKG\_SERVICE.JOB\_CANCEL

The stored procedure CANCEL deletes a specified task.

The prototype of the PKG\_SERVICE.JOB\_CANCEL function is as follows:

```
PKG_SERVICE.JOB_CANCEL(
job IN INTEGER);
```

**Table 10-11** PKG\_SERVICE.JOB\_CANCEL API parameters

| Parameter | Type    | Input/Output Parameter | Empty or Not | Description           |
|-----------|---------|------------------------|--------------|-----------------------|
| id        | integer | IN                     | No           | Specifies the job ID. |

Example:

```
SELECT PKG_SERVICE.JOB_SUBMIT (101, 'insert_msg_statistic1;', sysdate, 'sysdate+3.0/24');
job_submit

 101
(1 row)
CALL PKG_SERVICE.JOB_CANCEL(101);
job_cancel

(1 row)
```

- **PKG\_SERVICE.JOB\_FINISH**

The stored procedure FINISH disables or enables a scheduled task.

The prototype of the PKG\_SERVICE.JOB\_FINISH function is as follows:

```
PKG_SERVICE.JOB_FINISH(
id IN INTEGER,
broken IN BOOLEAN,
next_time IN TIMESTAMP DEFAULT sysdate);
```

**Table 10-12** PKG\_SERVICE.JOB\_FINISH API parameters

| Parameter | Type    | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                                                |
|-----------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id        | integer | IN                     | No           | Specifies the job ID.                                                                                                                                                                                                                      |
| broken    | Boolean | IN                     | No           | Specifies the status flag, <b>true</b> for broken and <b>false</b> for not broken. The current job is updated based on the parameter value <b>true</b> or <b>false</b> . If the parameter is left empty, the job status remains unchanged. |



| Parameter | Type      | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|-----------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| next_time | timestamp | IN                     | Yes          | Specifies the next execution time. The default value is the current system time. If <b>broken</b> is set to <b>true</b> , <b>next_time</b> is updated to '4000-1-1'. If <b>broken</b> is set to <b>false</b> and <b>next_time</b> is not empty, <b>next_time</b> is updated for the job. If <b>next_time</b> is empty, it will not be updated. This parameter can be omitted, and its default value will be used in this case. |

- PKG\_SERVICE.JOB\_SUBMIT

The stored procedure JOB\_SUBMIT commits a scheduled task provided by the system.

The prototype of the PKG\_SERVICE.JOB\_SUBMIT function is as follows:

```
PKG_SERVICE.JOB_SUBMIT(
id IN BIGINT,
content IN TEXT,
next_date IN TIMESTAMP DEFAULT sysdate,
interval_time IN TEXT DEFAULT 'null',
job OUT INTEGER);
```

 **NOTE**

When a scheduled task (using **JOB**) is created, the system binds the current database and the username to the task by default. This function can be called by using **call** or **select**. If you call this function by using **select**, there is no need to specify output parameters. To call this function within a stored procedure, use **perform**. If the committed SQL statement task uses a non-public schema, specify the schema to a table schema or a function schema, or add **set current\_schema = xxx** before the SQL statement.

**Table 10-13** PKG\_SERVICE.JOB\_SUBMIT API parameters

| Parameter | Type   | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                      |
|-----------|--------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id        | bigint | IN                     | No           | Specifies the job ID. If the input ID is <b>NULL</b> , a job ID is generated internally.                                                                                         |
| context   | text   | IN                     | No           | Specifies the SQL statement to be executed. One or multiple DML statements, anonymous blocks, and statements for calling stored procedures, or all three combined are supported. |

| Parameter     | Type      | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                                                                                                                     |
|---------------|-----------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| next_time     | timestamp | IN                     | No           | Specifies the next time the job will be executed. The default value is the current system time ( <b>sysdate</b> ). If the specified time has past, the job is executed at the time it is submitted.                                                                                                             |
| interval_time | text      | IN                     | Yes          | Calculates the next time to execute the job. It can be an interval expression, or <b>sysdate</b> followed by a numeric value, for example, <b>sysdate+1.0/24</b> . If this parameter is left empty or set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. |
| job           | integer   | OUT                    | No           | Specifies the job ID. The value ranges from 1 to 32767. When <b>pkg_service.job_submit</b> is called using <b>select</b> , this parameter can be omitted.                                                                                                                                                       |

## Example:

```
CREATE TABLE test_table(a int);
CREATE TABLE

CREATE OR REPLACE PROCEDURE test_job(a in int) IS
BEGIN
INSERT INTO test_table VALUES(a);
COMMIT;
END;
/
CREATE PROCEDURE
SELECT PKG_SERVICE.JOB_SUBMIT(NULL, 'call pro_xxx();', to_date('20180101','yyyymmdd'),'sysdate
+1');
job_submit

28269
(1 row)
SELECT PKG_SERVICE.JOB_SUBMIT(NULL, 'call pro_xxx();', to_date('20180101','yyyymmdd'),'sysdate
+1.0/24');
job_submit

1506
(1 row)
CALL PKG_SERVICE.JOB_SUBMIT(NULL, 'INSERT INTO T_JOB VALUES(1); call pro_1(); call pro_2();',
add_months(to_date('201701','yyyymm'),1), 'date_trunc("day",SYSDATE) + 1 +(8*60+30.0)/
(24*60)' ,:jobid);
job

14131
(1 row)
SELECT PKG_SERVICE.JOB_SUBMIT (101, 'insert_msg_statistic1;', sysdate, 'sysdate+3.0/24');
job_submit

```

101  
(1 row)

- **PKG\_SERVICE.JOB\_UPDATE**

The stored procedure UPDATE modifies user-definable attributes of a task, including the task content, next-execution time, and execution interval.

The prototype of the PKG\_SERVICE.JOB\_UPDATE function is as follows:

```
PKG_SERVICE.JOB_UPDATE(
id IN BIGINT,
next_time IN TIMESTAMP,
interval_time IN TEXT,
content IN TEXT);
```

**Table 10-14** PKG\_SERVICE.JOB\_UPDATE API parameters

| Parameter     | Type      | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------|-----------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id            | integer   | IN                     | No           | Specifies the job ID.                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| next_time     | timestamp | IN                     | Yes          | Specifies the next execution time. If this parameter is left empty, the system does not update the <b>next_time</b> parameter for the specified job. Otherwise, the system updates the <b>next_time</b> parameter for the specified job.                                                                                                                                                                                                                          |
| interval_time | text      | IN                     | Yes          | Specifies the time expression for calculating the next time the job will be executed. If this parameter is left empty, the system does not update the <b>interval_time</b> parameter for the specified job. Otherwise, the system updates the <b>interval_time</b> parameter for the specified job after necessary validity check. If this parameter is set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. |
| content       | text      | IN                     | Yes          | Specifies the name of the stored procedure or SQL statement block that is executed. If this parameter is left empty, the system does not update the <b>content</b> parameter for the specified job. Otherwise, the system updates the <b>content</b> parameter for the specified job.                                                                                                                                                                             |

Example:

```
CALL PKG_SERVICE.JOB_UPDATE(101, 'call userproc();', sysdate, 'sysdate + 1.0/1440');
job_update

(1 row)
CALL PKG_SERVICE.JOB_UPDATE(101, 'insert into tbl_a values(sysdate);', sysdate, 'sysdate + 1.0/1440');
job_update

(1 row)
```

- **PKG\_SERVICE.SUBMIT\_ON\_NODES**

Creates a scheduled job on a node. Only SYSADMIN and MONADMIN have this permission.

The prototype of the PKG\_SERVICE.SUBMIT\_ON\_NODES function is as follows:

```
PKG_SERVICE.SUBMIT_ON_NODES(
node_name IN NAME,
database IN NAME,
what IN TEXT,
next_date IN TIMESTAMP WITHOUT TIME ZONE,
job_interval IN TEXT,
job OUT INTEGER);
```

**Table 10-15** PKG\_SERVICE.SUBMIT\_ON\_NODES API parameters

| Parameter | Type      | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                       |
|-----------|-----------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name | text      | IN                     | No           | Specifies the node where a job is executed. Currently, the value can be ' <b>ALL_NODE</b> ' (the job is executed on all nodes) or ' <b>CCN</b> '. (Note: CCN is invalid in a centralized deployment environment.) |
| database  | text      | IN                     | No           | Database used by a database instance job. When the node type is ' <b>ALL_NODE</b> ', the value can only be ' <b>postgres</b> '.                                                                                   |
| what      | text      | IN                     | No           | Specifies the SQL statement to be executed. One or multiple DML statements, anonymous blocks, and statements for calling stored procedures, or all three combined are supported.                                  |
| nextdate  | timestamp | IN                     | No           | Specifies the next time the job will be executed. The default value is the current system time ( <b>sysdate</b> ). If the specified time has past, the job is executed at the time it is submitted.               |

| Parameter    | Type    | Input/Output Parameter | Empty or Not | Description                                                                                                                                                                                                                                                                                                     |
|--------------|---------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_interval | text    | IN                     | No           | Calculates the next time to execute the job. It can be an interval expression, or <b>sysdate</b> followed by a numeric value, for example, <b>sysdate+1.0/24</b> . If this parameter is left empty or set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. |
| job          | integer | OUT                    | No           | Specifies the job ID. The value ranges from 1 to 32767. When dbms.submit_on_nodes is called using SELECT, this parameter can be omitted.                                                                                                                                                                        |

Example:

```
SELECT pkg_service.submit_on_nodes('ALL_NODE', 'postgres', 'select
capture_view_to_json("dbe_perf.statement", 0);', sysdate, 'interval "60 second"');
submit_on_nodes
```

```

25376
(1 row)
```

- **PKG\_SERVICE.ISUBMIT\_ON\_NODES**

ISUBMIT\_ON\_NODES has the same syntax function as SUBMIT\_ON\_NODES, but the first parameter of ISUBMIT\_ON\_NODES is an input parameter, that is, a specified task ID. In contrast, that last parameter of ISUBMIT\_ON\_NODES is an output parameter, indicating the task ID automatically generated by the system. Only SYSADMIN and MONADMIN have this permission.

- **PKG\_SERVICE.SQL\_GET\_ARRAY\_RESULT**

This function is used to return the value of the bound OUT parameter of the array type and obtain the OUT parameter in a stored procedure.

The prototype of the PKG\_SERVICE.SQL\_GET\_ARRAY\_RESULT function is as follows:

```
PKG_SERVICE.SQL_GET_ARRAY_RESULT(
 context_id in int,
 pos in VARCHAR2,
 column_value inout anyarray,
 result_type in anyelement
);
```

**Table 10-16** PKG\_SERVICE.SQL\_GET\_ARRAY\_RESULT parameters

| Parameter  | Description                      |
|------------|----------------------------------|
| context_id | ID of the context to be queried. |
| pos        | Name of the bound parameter      |

| Parameter    | Description  |
|--------------|--------------|
| column_value | Return value |
| result_type  | Return type  |

- **PKG\_SERVICE.SQL\_GET\_VARIABLE\_RESULT**

This function is used to return the value of the bound OUT parameter of the non-array type and obtain the OUT parameter in a stored procedure.

The prototype of the PKG\_SERVICE.SQL\_GET\_VARIABLE\_RESULT function is as follows:

```
PKG_SERVICE.SQL_GET_VARIABLE_RESULT(
 context_id in int,
 pos in VARCHAR2,
 result_type in anyelement
)
RETURNS anyelement;
```

**Table 10-17** PKG\_SERVICE.SQL\_GET\_VARIABLE\_RESULT parameters

| Parameter   | Description                      |
|-------------|----------------------------------|
| context_id  | ID of the context to be queried. |
| pos         | Name of the bound parameter.     |
| result_type | Return type.                     |

### 10.12.1.2 PKG\_UTIL

**Table 10-18** lists all APIs supported by the PKG\_UTIL package.

**Table 10-18** PKG\_UTIL

| API                            | Description                                                                   |
|--------------------------------|-------------------------------------------------------------------------------|
| <b>PKG_UTIL.LOB_GET_LENGTH</b> | Obtains the length of a LOB.                                                  |
| <b>PKG_UTIL.LOB_READ</b>       | Reads a part of a LOB.                                                        |
| <b>PKG_UTIL.LOB_WRITE</b>      | Writes the source object to the target object in the specified format.        |
| <b>PKG_UTIL.LOB_APPEND</b>     | Appends a specified number of characters of the source LOB to the target LOB. |
| <b>PKG_UTIL.LOB_COMPARE</b>    | Compares two LOBs based on the specified length.                              |
| <b>PKG_UTIL.LOB_MATCH</b>      | Returns the position of the Mth occurrence of a character string in a LOB.    |

| API                                                   | Description                                                                   |
|-------------------------------------------------------|-------------------------------------------------------------------------------|
| <a href="#">PKG_UTIL.LOB_RESET</a>                    | Resets the character in specified position of a LOB to a specified character. |
| <a href="#">PKG_UTIL.IO_PRINT</a>                     | Displays character strings.                                                   |
| <a href="#">PKG_UTIL.RAW_GET_LENGTH</a>               | Obtains the length of RAW data.                                               |
| <a href="#">PKG_UTIL.RAW_CAST_FROM_VARCHAR2</a>       | Converts VARCHAR2 data to RAW data.                                           |
| <a href="#">PKG_UTIL.RAW_CAST_FROM_BINARY_INTEGER</a> | Converts binary integers to RAW data.                                         |
| <a href="#">PKG_UTIL.RAW_CAST_TO_BINARY_INTEGER</a>   | Converts RAW data to binary integers.                                         |
| <a href="#">PKG_UTIL.RANDOM_SET_SEED</a>              | Sets a random seed.                                                           |
| <a href="#">PKG_UTIL.RANDOM_GET_VALUE</a>             | Returns a random value.                                                       |
| <a href="#">PKG_UTIL.FILE_SET_DIRNAME</a>             | Sets the directory to be operated.                                            |
| <a href="#">PKG_UTIL.FILE_OPEN</a>                    | Opens a file based on the specified file name and directory.                  |
| <a href="#">PKG_UTIL.FILE_SET_MAX_LINE_SIZE</a>       | Sets the maximum length of a line to be written to a file.                    |
| <a href="#">PKG_UTIL.FILE_IS_CLOSE</a>                | Checks whether a file handle is closed.                                       |
| <a href="#">PKG_UTIL.FILE_READ</a>                    | Reads data of a specified length from an open file handle.                    |
| <a href="#">PKG_UTIL.FILE_READLINE</a>                | Reads a line of data from an open file handle.                                |
| <a href="#">PKG_UTIL.FILE_WRITE</a>                   | Writes the data specified in the buffer to a file.                            |
| <a href="#">PKG_UTIL.FILE_WRITELINE</a>               | Writes the buffer to a file and adds newline characters.                      |
| <a href="#">PKG_UTIL.FILE_NEWLINE</a>                 | Adds a line.                                                                  |
| <a href="#">PKG_UTIL.FILE_READ_RAW</a>                | Reads binary data of a specified length from an open file handle.             |
| <a href="#">PKG_UTIL.FILE_WRITE_RAW</a>               | Writes binary data to a file.                                                 |
| <a href="#">PKG_UTIL.FILE_FLUSH</a>                   | Writes data from a file handle to a physical file.                            |
| <a href="#">PKG_UTIL.FILE_CLOSE</a>                   | Closes an open file handle.                                                   |

| API                                                     | Description                                                                    |
|---------------------------------------------------------|--------------------------------------------------------------------------------|
| <a href="#">PKG_UTIL.FILE_REMOVE</a>                    | Deletes a physical file. To do so, you must have the corresponding permission. |
| <a href="#">PKG_UTIL.FILE_RENAME</a>                    | Renames files on the disk, similar to <b>mv</b> in Unix.                       |
| <a href="#">PKG_UTIL.FILE_SIZE</a>                      | Returns the size of a file.                                                    |
| <a href="#">PKG_UTIL.FILE_BLOCK_SIZE</a>                | Returns the number of blocks contained in a file.                              |
| <a href="#">PKG_UTIL.FILE_EXISTS</a>                    | Checks whether a file exists.                                                  |
| <a href="#">PKG_UTIL.FILE_GETPOS</a>                    | Specifies the offset of a returned file, in bytes.                             |
| <a href="#">PKG_UTIL.FILE_SEEK</a>                      | Sets the offset for file position.                                             |
| <a href="#">PKG_UTIL.FILE_CLOSE_ALL</a>                 | Closes all file handles opened in a session.                                   |
| <a href="#">PKG_UTIL.EXCEPTION_REPORT_ERROR</a>         | Throws an exception.                                                           |
| <a href="#">PKG_UTIL.APP_READ_CLIENT_INFO</a>           | Reads the client information.                                                  |
| <a href="#">PKG_UTIL.APP_SET_CLIENT_INFO</a>            | Sets the client information.                                                   |
| <a href="#">PKG_UTIL.LOB_CONVERTTOBLOB</a>              | Converts the CLOB type to the BLOB type.                                       |
| <a href="#">PKG_UTIL.LOB_CONVERTTOCLOB</a>              | Converts the BLOB type to the CLOB type.                                       |
| <a href="#">PKG_UTIL.LOB_RAWTOTEXT</a>                  | Converts the raw type to the text type.                                        |
| <a href="#">PKG_UTIL.LOB_TEXTTORAW</a>                  | Converts the text type to the raw type.                                        |
| <a href="#">PKG_UTIL.MATCH_EDIT_DISTANCE_SIMILARITY</a> | Calculates the difference between two character strings.                       |
| <a href="#">PKG_UTIL.RAW_CAST_TO_VARCHAR2</a>           | Converts the raw type to the varchar2 type.                                    |
| <a href="#">PKG_UTIL.SESSION_CLEAR_CONTEXT</a>          | Clears the attribute values in <b>session_context</b> .                        |
| <a href="#">PKG_UTIL.SESSION_SEARCH_CONTEXT</a>         | Searches for an attribute value.                                               |
| <a href="#">PKG_UTIL.SESSION_SET_CONTEXT</a>            | Sets an attribute value.                                                       |
| <a href="#">PKG_UTIL.UTILITY_FORMAT_CALL_STACK</a>      | Displays the call stack of a stored procedure.                                 |



| API                                                     | Description                                              |
|---------------------------------------------------------|----------------------------------------------------------|
| <a href="#">PKG_UTIL.UTILITY_FORMAT_ERROR_BACKTRACE</a> | Displays the error stack of a stored procedure.          |
| <a href="#">PKG_UTIL.UTILITY_FORMAT_ERROR_STACK</a>     | Displays the error information about a stored procedure. |
| <a href="#">PKG_UTIL.UTILITY_GET_TIME</a>               | Displays the Unix timestamp of the system.               |

- **PKG\_UTIL.LOB\_GET\_LENGTH**

Obtains the length of the input data.

The prototype of the PKG\_UTIL.LOB\_GET\_LENGTH function is as follows:

```
PKG_UTIL.LOB_GET_LENGTH(
lob IN anyelement
)
RETURN INTEGER;
```

**Table 10-19** PKG\_UTIL.LOB\_GET\_LENGTH parameters

| Parameter | Type        | Input/Output Parameter | Can Be Empty | Description                                          |
|-----------|-------------|------------------------|--------------|------------------------------------------------------|
| lob       | clob / blob | IN                     | No           | Indicates the object whose length is to be obtained. |

- **PKG\_UTIL.LOB\_READ**

Reads an object and returns the specified part.

The prototype of the PKG\_UTIL.LOB\_READ function is as follows:

```
PKG_UTIL.LOB_READ(
lob IN anyelement,
len IN int,
start IN int,
mode IN int
)
RETURN ANYELEMENT
```

**Table 10-20** PKG\_UTIL.LOB\_READ parameters

| Parameter | Type       | Input/Output Parameter | Can Be Empty | Description                  |
|-----------|------------|------------------------|--------------|------------------------------|
| lob       | clob/ blob | IN                     | No           | Specifies CLOB or BLOB data. |

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                                        |
|-----------|------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| len       | int  | IN                     | No           | Specifies the length of the returned result.                                                                                                       |
| start     | int  | IN                     | No           | Specifies the offset to the first character.                                                                                                       |
| mode      | int  | IN                     | No           | Specifies the type of the read operation. <b>0</b> indicates <b>READ</b> , <b>1</b> indicates <b>TRIM</b> , and <b>2</b> indicates <b>SUBSTR</b> . |

- **PKG\_UTIL.LOB\_WRITE**

Writes the source object to the target object based on the specified parameters and returns the target object.

The prototype of the PKG\_UTIL.LOB\_WRITE function is as follows:

```
PKG_UTIL.LOB_WRITE(
dest_lob INOUT blob,
src_lob IN raw
len IN int,
start_pos IN int
)
RETURN BLOB;
PKG_UTIL.LOB_WRITE(
dest_lob INOUT clob,
src_lob IN varchar2
len IN int,
start_pos IN int
)
RETURN CLOB;
```

**Table 10-21** PKG\_UTIL.LOB\_WRITE parameters

| Parameter | Type          | Input/Output Parameter | Can Be Empty | Description                                               |
|-----------|---------------|------------------------|--------------|-----------------------------------------------------------|
| dest_lob  | clob/<br>blob | INOUT                  | No           | Specifies the target object that data will be written to. |
| src_lob   | clob/<br>blob | IN                     | No           | Specifies the source object to be written.                |
| len       | int           | IN                     | No           | Specifies the write length of the source object.          |
| start_pos | int           | IN                     | No           | Specifies the write start position of the target object.  |

- **PKG\_UTIL.LOB\_APPEND**

Appends the source object to the target BLOB/CLOB and returns the target BLOB/CLOB.

The prototype of the PKG\_UTIL.LOB\_APPEND function is as follows:

```
PKG_UTIL.LOB_APPEND(
dest_lob INOUT blob,
src_lob IN blob,
len IN int default NULL
)
RETURN BLOB;

PKG_UTIL.LOB_APPEND(
dest_lob INOUT clob,
src_lob IN clob,
len IN int default NULL
)
RETURN CLOB;
```

**Table 10-22** PKG\_UTIL.LOB\_APPEND parameters

| Parameter | Type        | Input/Output Parameter | Can Be Empty | Description                                                                                                              |
|-----------|-------------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------|
| dest_lob  | blob / clob | INOUT                  | No           | Specifies the target BLOB/CLOB that data will be written to.                                                             |
| src_lob   | blob / clob | IN                     | No           | Specifies the source BLOB/CLOB to be written.                                                                            |
| len       | int         | IN                     | Yes          | Length of the source object to be written. If the value is <b>NULL</b> , the entire source object is written by default. |

- **PKG\_UTIL.LOB\_COMPARE**

Checks whether objects are the same based on the specified start position and size. If **lob1** is larger, **1** is returned. If **lob2** is larger, **-1** is returned. If **lob1** is equal to **lob2**, **0** is returned.

The prototype of the PKG\_UTIL.LOB\_COMPARE function is as follows:

```
PKG_UTIL.LOB_COMPARE(
lob1 IN anyelement,
lob2 IN anyelement,
len IN int default 1073741771,
start_pos1 IN int default 1,
start_pos2 IN int default 1
)
RETURN INTEGER;
```

**Table 10-23** PKG\_UTIL.LOB\_COMPARE parameters

| Parameter  | Type      | Input / Output Parameter | Can Be Empty | Description                                    |
|------------|-----------|--------------------------|--------------|------------------------------------------------|
| lob1       | clob/blob | IN                       | No           | Indicates the character string for comparison. |
| lob2       | clob/blob | IN                       | No           | Indicates the character string for comparison. |
| len        | int       | IN                       | No           | Indicates the length to be compared.           |
| start_pos1 | int       | IN                       | No           | Specifies the start offset of <b>lob1</b> .    |
| start_pos2 | int       | IN                       | No           | Specifies the start offset of <b>lob2</b> .    |

- PKG\_UTIL.LOB\_MATCH

Returns the position where a pattern is displayed in a LOB for the *match\_nth* time.

The prototype of the PKG\_UTIL.LOB\_MATCH function is as follows:

```
PKG_UTIL.LOB_MATCH(
lob IN anyelement,
pattern IN anyelement,
start IN int,
match_nth IN int default 1
)
RETURN INTEGER;
```

**Table 10-24** PKG\_UTIL.LOB\_MATCH parameters

| Parameter | Type      | Input / Output Parameter | Can Be Empty | Description                                      |
|-----------|-----------|--------------------------|--------------|--------------------------------------------------|
| lob       | clob/blob | IN                       | No           | Indicates the character string for comparison.   |
| pattern   | clob/blob | IN                       | No           | Specifies the pattern to be matched.             |
| start     | int       | IN                       | No           | Specifies the start position for LOB comparison. |

| Parameter | Type | Input / Output Parameter | Can Be Empty | Description                   |
|-----------|------|--------------------------|--------------|-------------------------------|
| match_nth | int  | IN                       | No           | Specifies the matching times. |

- **PKG\_UTIL.LOB\_RESET**

Clears a character string and resets the string to the value of **value**.

The prototype of the PKG\_UTIL.LOB\_RESET function is as follows:

```
PKG_UTIL.LOB_RESET(
lob INOUT blob,
len INOUT int,
start IN int DEFAULT 1,
value IN int default 0
)
RETURN record;
```

**Table 10-25** PKG\_UTIL.LOB\_RESET parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                     |
|-----------|------|------------------------|--------------|-------------------------------------------------|
| lob       | blob | IN                     | No           | Indicates the character string for reset.       |
| len       | int  | IN                     | No           | Specifies the length of the string to be reset. |
| start     | int  | IN                     | No           | Specifies the start position for reset.         |
| value     | int  | IN                     | Yes          | Sets characters. Default value: '0'.            |

- **PKG\_UTIL.IO\_PRINT**

Outputs a string.

The prototype of the PKG\_UTIL.IO\_PRINT function is as follows:

```
PKG_UTIL.IO_PRINT(
format IN text,
is_one_line IN boolean
)
RETURN void;
```

**Table 10-26** PKG\_UTIL.IO\_PRINT parameters

| Parameter   | Type    | Input/Output Parameter | Can Be Empty | Description                                       |
|-------------|---------|------------------------|--------------|---------------------------------------------------|
| format      | text    | IN                     | No           | Indicates the character string to be output.      |
| is_one_line | boolean | IN                     | No           | Specifies whether to output the string as a line. |

- PKG\_UTIL.RAW\_GET\_LENGTH

Obtains the length of RAW data.

The prototype of the PKG\_UTIL.RAW\_GET\_LENGTH function is as follows:

```
PKG_UTIL.RAW_GET_LENGTH(
value IN raw
)
RETURN integer;
```

**Table 10-27** PKG\_UTIL.RAW\_GET\_LENGTH parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                          |
|-----------|------|------------------------|--------------|------------------------------------------------------|
| raw       | raw  | IN                     | No           | Indicates the object whose length is to be obtained. |

- PKG\_UTIL.RAW\_CAST\_FROM\_VARCHAR2

Converts VARCHAR2 data to RAW data.

The prototype of the PKG\_UTIL.RAW\_CAST\_FROM\_VARCHAR2 function is as follows:

```
PKG_UTIL.RAW_CAST_FROM_VARCHAR2(
str IN varchar2
)
RETURN raw;
```

**Table 10-28** PKG\_UTIL.RAW\_CAST\_FROM\_VARCHAR2 parameters

| Parameter | Type     | Input/Output Parameter | Can Be Empty | Description                                |
|-----------|----------|------------------------|--------------|--------------------------------------------|
| str       | varchar2 | IN                     | No           | Specifies the source data to be converted. |

- PKG\_UTIL.RAW\_CAST\_FROM\_BINARY\_INTEGER

Converts BIGINT data to RAW data.

The prototype of the `PKG_UTIL.RAW_CAST_FROM_BINARY_INTEGER` function is as follows:

```
PKG_UTIL.RAW_CAST_FROM_BINARY_INTEGER(
value IN BIGINT,
endianess IN INTEGER
)
RETURN RAW;
```

**Table 10-29** `PKG_UTIL.RAW_CAST_FROM_BINARY_INTEGER` parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                    |
|-----------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------|
| value     | BIGINT  | IN                     | No           | Specifies the source data to be converted.                                                                                     |
| endianess | INTEGER | IN                     | No           | The value is an integer in lexicographic order. Currently, the value can be <b>1</b> (BIG_ENDIAN) or <b>2</b> (LITTLE_ENDIAN). |

- `PKG_UTIL.RAW_CAST_TO_BINARY_INTEGER`

Converts RAW data into BINARY\_INTEGER.

The prototype of the `PKG_UTIL.RAW_CAST_TO_BINARY_INTEGER` function is as follows:

```
PKG_UTIL.RAW_CAST_TO_BINARY_INTEGER(
value IN RAW,
endianess IN INTEGER
)
RETURN INTEGER;
```

**Table 10-30** `PKG_UTIL.RAW_CAST_TO_BINARY_INTEGER` parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                    |
|-----------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------|
| value     | RAW     | IN                     | No           | Specifies the source data to be converted.                                                                                     |
| endianess | INTEGER | IN                     | No           | The value is an integer in lexicographic order. Currently, the value can be <b>1</b> (BIG_ENDIAN) or <b>2</b> (LITTLE_ENDIAN). |

- `PKG_UTIL.RANDOM_SET_SEED`

Sets a random seed.

The prototype of the **`PKG_UTIL.RANDOM_SET_SEED`** function is as follows:

```
PKG_UTIL.RANDOM_SET_SEED(
seed IN int
)
RETURN integer;
```

**Table 10-31** PKG\_UTIL.RANDOM\_SET\_SEED parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description         |
|-----------|------|------------------------|--------------|---------------------|
| seed      | int  | IN                     | No           | Sets a random seed. |

- PKG\_UTIL.RANDOM\_GET\_VALUE

Returns a 15-digit random number ranging from 0 to 1.

The prototype of the PKG\_UTIL.RANDOM\_GET\_VALUE function is as follows:

```
PKG_UTIL.RANDOM_GET_VALUE(
)
RETURN numeric;
```

- PKG\_UTIL.FILE\_SET\_DIRNAME

Sets the directory to be operated. It must be called to set directory for each operation involving a single directory.

The prototype of the **PKG\_UTIL.FILE\_SET\_DIRNAME** function is as follows:

```
PKG_UTIL.FILE_SET_DIRNAME(
dir IN text
)
RETURN bool
```

**Table 10-32** PKG\_UTIL.FILE\_SET\_DIRNAME parameters

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dirname   | Directory of a file. It is a string, indicating an object name.<br><b>NOTE</b><br>File directories need to be added to the system catalog <b>PG_DIRECTORY</b> . If the input path does not match the path in <b>PG_DIRECTORY</b> , an error indicating that the path does not exist is reported. Functions that involve <b>location</b> as parameters also comply with this rule. |

- PKG\_UTIL.FILE\_OPEN

Opens a file. A maximum of 50 files can be opened at a time. This function returns a handle of the INTEGER type.

The prototype of the PKG\_UTIL.FILE\_OPEN function is as follows:

```
PKG_UTIL.FILE_OPEN(
file_name IN text,
open_mode IN integer)
```

**Table 10-33** PKG\_UTIL.FILE\_OPEN parameters

| Parameter | Description                                                                                                                                                                                                          |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file_name | File name with an extension (file type), excluding the path name. A path contained in a file name is ignored in the OPEN function. In Unix, the file name cannot end with the combination of a slash and a dot (/.). |



| Parameter | Description                                                                                                                                                                                                                               |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| open_mode | File opening mode, including <b>r</b> (read), <b>w</b> (write), and <b>a</b> (append).<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits. |

- **PKG\_UTIL.FILE\_SET\_MAX\_LINE\_SIZE**

Sets the maximum length of a line to be written to a file.

The prototype of the PKG\_UTIL.FILE\_SET\_MAX\_LINE\_SIZE function is as follows:

```
PKG_UTIL.FILE_SET_MAX_LINE_SIZE(
max_line_size in integer)
RETURN BOOL
```

**Table 10-34** PKG\_UTIL.FILE\_SET\_MAX\_LINE\_SIZE parameters

| Parameter     | Description                                                                                                                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| max_line_size | Maximum number of characters in each line, including newline characters. The minimum value is <b>1</b> and the maximum is <b>32767</b> . If this parameter is not specified, the default value <b>1024</b> is used. |

- **PKG\_UTIL.FILE\_IS\_CLOSE**

Checks whether a file handle is closed.

The prototype of the PKG\_UTIL.FILE\_IS\_CLOSE function is as follows:

```
PKG_UTIL.FILE_IS_CLOSE(
file in integer
)
RETURN BOOL
```

**Table 10-35** PKG\_UTIL.FILE\_IS\_CLOSE parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

- **PKG\_UTIL.FILE\_READ**

Reads a line of data from an open file handle based on the specified length.

The prototype of the PKG\_UTIL.FILE\_READ function is as follows:

```
PKG_UTIL.FILE_READ(
file IN integer,
buffer OUT text,
len IN bigint default 1024)
```

**Table 10-36** PKG\_UTIL.FILE\_READ parameters

| Parameter | Description                                                                                                                                         |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| file      | File handle opened by calling the OPEN function. The file must be opened in read mode. Otherwise, the <b>INVALID_OPERATION</b> exception is thrown. |
| buffer    | Buffer used to receive data.                                                                                                                        |
| len       | Number of bytes read from a file.                                                                                                                   |

- PKG\_UTIL.FILE\_READLINE

Reads a line of data from an open file handle based on the specified length.

The prototype of the PKG\_UTIL.FILE\_READLINE function is as follows:

```
PKG_UTIL.FILE_READLINE(
file IN integer,
buffer OUT text,
len IN integer default NULL)
```

**Table 10-37** PKG\_UTIL.FILE\_READLINE parameters

| Parameter | Description                                                                                                                                                           |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file      | File handle opened by calling the OPEN function. The file must be opened in read mode. Otherwise, the <b>INVALID_OPERATION</b> exception is thrown.                   |
| buffer    | Buffer used to receive data.                                                                                                                                          |
| len       | Number of bytes read from a file. The default value is <b>NULL</b> . If the default value <b>NULL</b> is used, <b>max_line_size</b> is used to specify the line size. |

- PKG\_UTIL.FILE\_WRITE

Writes the data specified in the buffer to a file.

The prototype of the PKG\_UTIL.FILE\_WRITE function is as follows:

```
PKG_UTIL.FILE_WRITE(
file in integer,
buffer in text
)
RETURN BOOL
```

**Table 10-38** PKG\_UTIL.FILE\_WRITE parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

| Parameter | Description                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffer    | Text data to be written to a file. The maximum buffer size is 32767 bytes. If no value is specified, the default value is 1024 bytes. Before the writing is performed, the buffer occupied by <b>PUT</b> operations cannot exceed 32767 bytes.<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits. |

- **PKG\_UTIL.FILE\_NEWLINE**

Writes a line terminator to an open file. The line terminator is related to the platform.

The prototype of the PKG\_UTIL.FILE\_NEWLINE function is as follows:

```
PKG_UTIL.FILE_NEWLINE(
file in integer
)
RETURN BOOL
```

**Table 10-39** PKG\_UTIL.FILE\_NEWLINE parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

- **PKG\_UTIL.FILE\_WRITELINE**

Writes a line to a file.

The prototype of the PKG\_UTIL.FILE\_WRITELINE function is as follows:

```
PKG_UTIL.FILE_WRITELINE(
file in integer,
buffer in text
)
RETURN BOOL
```

**Table 10-40** PKG\_UTIL.FILE\_WRITELINE parameters

| Parameter | Description            |
|-----------|------------------------|
| file      | Opened file handle.    |
| buffer    | Content to be written. |

- **PKG\_UTIL.FILE\_READ\_RAW**

Reads binary data of a specified length from an open file handle and returns the read binary data. The return type is RAW.

The prototype of the PKG\_UTIL.FILE\_READ\_RAW function is as follows:

```
PKG_UTIL.FILE_READ_RAW(
file in integer,
length in integer default NULL
)
RETURN raw
```

**Table 10-41** PKG\_UTIL.FILE\_READ\_RAW parameters

| Parameter | Description                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------|
| file      | Opened file handle.                                                                                                                   |
| length    | Length of the data to be read. The default value is <b>NULL</b> . By default, all data in the file is read. The maximum size is 1 GB. |

- PKG\_UTIL.FILE\_WRITE\_RAW

Writes the input binary object of the RAW type to an open file. If the insertion is successful, **true** is returned.

The prototype of the PKG\_UTIL.FILE\_WRITE\_RAW function is as follows:

```
PKG_UTIL.FILE_WRITE_RAW(
file in integer,
r in raw
)
RETURN BOOL
```

**Table 10-42** PKG\_UTIL.FILE\_NEWLINE parameters

| Parameter | Description                                                                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file      | Opened file handle.                                                                                                                                                                |
| r         | Data to be written to the file.<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits. |

- PKG\_UTIL.FILE\_FLUSH

Data in a file handle must be written into a physical file. Data in the buffer must have a line terminator. Refresh is important if a file must be read when it is opened. For example, debugging information can be refreshed to a file so that it can be read immediately.

The prototype of the PKG\_UTIL.FILE\_FLUSH function is as follows:

```
PKG_UTIL.FILE_FLUSH (
file in integer
)
RETURN VOID
```

**Table 10-43** PKG\_UTIL.FILE\_FLUSH parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

- PKG\_UTIL.FILE\_CLOSE

Closes an open file handle.

The prototype of the PKG\_UTIL.FILE\_CLOSE function is as follows:

```
PKG_UTIL.FILE_CLOSE (
file in integer
```

```
)
RETURN BOOL
```

**Table 10-44** PKG\_UTIL.FILE\_CLOSE parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

- PKG\_UTIL.FILE\_REMOVE

Deletes a disk file. To perform this operation, you must have required permissions.

The prototype of the PKG\_UTIL.FILE\_REMOVE function is as follows:

```
PKG_UTIL.FILE_REMOVE(
file_name in text
)
RETURN VOID
```

**Table 10-45** PKG\_UTIL.FILE\_REMOVE parameters

| Parameter | Description                    |
|-----------|--------------------------------|
| filen_ame | Name of the file to be deleted |

- PKG\_UTIL.FILE\_RENAME

The function renames files on the disk, similar to **mv** in Unix.

The prototype of the PKG\_UTIL.FILE\_RENAME function is as follows:

```
PKG_UTIL.FILE_RENAME(
src_dir in text,
src_file_name in text,
dest_dir in text,
dest_file_name in text,
overwrite boolean default false)
```

**Table 10-46** PKG\_UTIL.FILE\_RENAME parameters

| Parameter     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| src_dir       | Source file directory (case-sensitive)<br><b>NOTE</b> <ul style="list-style-type: none"><li>• File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist is reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>• When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| src_file_name | Source file name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Parameter      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_dir       | Target file directory (case-sensitive)<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist is reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| dest_file_name | Target file name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| overwrite      | The default value is <b>false</b> . If a file with the same name exists in the destination directory, the file will not be rewritten.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- **PKG\_UTIL.FILE\_SIZE**

Returns the size of a specified file.

The prototype of the PKG\_UTIL.FILE\_SIZE function is as follows:

```
bigint PKG_UTIL.FILE_SIZE(
file_name in text
)return bigint
```

**Table 10-47** PKG\_UTIL.FILE\_SIZE parameters

| Parameter | Description |
|-----------|-------------|
| file_name | File name   |

- **PKG\_UTIL.FILE\_BLOCK\_SIZE**

Returns the number of blocks contained in a specified file.

The prototype of the PKG\_UTIL.FILE\_BLOCK\_SIZE function is as follows:

```
bigint PKG_UTIL.FILE_BLOCK_SIZE(
file_name in text
)return bigint
```

**Table 10-48** PKG\_UTIL.FILE\_BLOCK\_SIZE parameters

| Parameter | Description |
|-----------|-------------|
| file_name | File name   |

- **PKG\_UTIL.FILE\_EXISTS**

Checks whether a file exists.

The prototype of the PKG\_UTIL.FILE\_EXISTS function is as follows:

```
PKG_UTIL.FILE_EXISTS(
file_name in text
)
RETURN BOOL
```

**Table 10-49** PKG\_UTIL.FILE\_EXISTS parameters

| Parameter | Description |
|-----------|-------------|
| file_name | File name   |

- PKG\_UTIL.FILE\_GETPOS

Specifies the offset of a returned file, in bytes.

The prototype of the PKG\_UTIL.FILE\_GETPOS function is as follows:

```
PKG_UTIL.FILE_GETPOS(
file in integer
)
RETURN BIGINT
```

**Table 10-50** PKG\_UTIL.FILE\_GETPOS parameters

| Parameter | Description         |
|-----------|---------------------|
| file      | Opened file handle. |

- PKG\_UTIL.FILE\_SEEK

Adjusts the position of a file pointer forward or backward based on the specified number of bytes.

The prototype of the PKG\_UTIL.FILE\_SEEK function is as follows:

```
void PKG_UTIL.FILE_SEEK(
file in integer,
start in bigint
)
RETURN VOID
```

**Table 10-51** PKG\_UTIL.FILE\_SEEK parameters

| Parameter | Description            |
|-----------|------------------------|
| file      | Opened file handle.    |
| start     | File offset, in bytes. |

- PKG\_UTIL.FILE\_CLOSE\_ALL

Closes all file handles opened in a session.

The prototype of the PKG\_UTIL.FILE\_CLOSE\_ALL function is as follows:

```
PKG_UTIL.FILE_CLOSE_ALL(
)
RETURN VOID
```

**Table 10-52** PKG\_UTIL.FILE\_CLOSE\_ALL parameters

| Parameter | Description |
|-----------|-------------|
| None      | None        |

- **PKG\_UTIL.EXCEPTION\_REPORT\_ERROR**

Throws an exception.

The prototype of the PKG\_UTIL.EXCEPTION\_REPORT\_ERROR function is as follows:

```
PKG_UTIL.EXCEPTION_REPORT_ERROR(
code integer,
log text,
flag boolean DEFAULT false
)
RETURN INTEGER
```

**Table 10-53** PKG\_UTIL.EXCEPTION\_REPORT\_ERROR parameters

| Parameter | Description                                    |
|-----------|------------------------------------------------|
| code      | Error code displayed when an exception occurs. |
| log       | Logs displayed when an exception occurs.       |
| flag      | Reserved. The default value is <b>false</b> .  |

- **PKG\_UTIL.APP\_READ\_CLIENT\_INFO**

Reads the client information.

The prototype of the PKG\_UTIL.APP\_READ\_CLIENT\_INFO function is as follows:

```
PKG_UTIL.APP_READ_CLIENT_INFO(
OUT buffer text
)return text
```

**Table 10-54** PKG\_UTIL.APP\_READ\_CLIENT\_INFO parameters

| Parameter | Description                  |
|-----------|------------------------------|
| buffer    | Client information returned. |

- **PKG\_UTIL.APP\_SET\_CLIENT\_INFO**

Sets the client information.

The prototype of the PKG\_UTIL.APP\_SET\_CLIENT\_INFO function is as follows:

```
PKG_UTIL.APP_SET_CLIENT_INFO(
str text
)
```

**Table 10-55** PKG\_UTIL.APP\_SET\_CLIENT\_INFO parameters

| Parameter | Description                   |
|-----------|-------------------------------|
| str       | Client information to be set. |

- **PKG\_UTIL.LOB\_CONVERTTOBLOB**

Converts a CLOB to a BLOB. **amount** indicates the conversion length.

The prototype of the PKG\_UTIL.LOB\_CONVERTTOBLOB function is as follows:



```
PKG_UTIL.LOB_CONVERTTOBLOB(
dest_lob blob,
src_clob clob,
amount integer,
dest_offset integer,
src_offset integer
)return raw
```

**Table 10-56** PKG\_UTIL.LOB\_CONVERTTOBLOB parameters

| Parameter   | Description                        |
|-------------|------------------------------------|
| dest_lob    | Target LOB.                        |
| src_clob    | CLOB to be converted.              |
| amount      | Conversion length.                 |
| dest_offset | Start position of the target LOB.  |
| src_offset  | Start position of the source CLOB. |

- PKG\_UTIL.LOB\_CONVERTTOCLOB

Converts a BLOB to a CLOB. **amount** indicates the conversion length.

The prototype of the PKG\_UTIL.LOB\_CONVERTTOCLOB function is as follows:

```
PKG_UTIL.LOB_CONVERTTOCLOB(
dest_lob clob,
src_blob blob,
amount integer,
dest_offset integer,
src_offset integer
)return text
```

**Table 10-57** PKG\_UTIL.LOB\_CONVERTTOCLOB parameters

| Parameter   | Description                        |
|-------------|------------------------------------|
| dest_lob    | Target LOB.                        |
| src_blob    | BLOB to be converted.              |
| amount      | Conversion length.                 |
| dest_offset | Start position of the target LOB.  |
| src_offset  | Start position of the source CLOB. |

- PKG\_UTIL.LOB\_RAWTOTEXT

Converts RAW data to TEXT data.

The prototype of the PKG\_UTIL.LOB\_RAWTOTEXT function is as follows:

```
PKG_UTIL.LOB_RAWTOTEXT(
src_lob IN BLOB
)
RETURN TEXT
```

**Table 10-58** PKG\_UTIL.LOB\_RAWTOTEXT parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description          |
|-----------|------|------------------------|--------------|----------------------|
| src_lob   | LOB  | IN                     | No           | LOB to be converted. |

- PKG\_UTIL.LOB\_TEXTTORAW

Converts from the text type to the raw type.

The prototype of the PKG\_UTIL.LOB\_TEXTTORAW function is as follows:

```
PKG_UTIL.LOB_TEXTTORAW(
src_lob clob
)
RETURN raw
```

**Table 10-59** PKG\_UTIL.LOB\_TEXTTORAW parameters

| Parameter | Description          |
|-----------|----------------------|
| src_lob   | LOB to be converted. |

- PKG\_UTIL.MATCH\_EDIT\_DISTANCE\_SIMILARITY

Calculates the difference between two character strings.

The prototype of the PKG\_UTIL.MATCH\_EDIT\_DISTANCE\_SIMILARITY function is as follows:

```
PKG_UTIL.MATCH_EDIT_DISTANCE_SIMILARITY(
str1 text,
str2 text
)
RETURN INTEGER
```

**Table 10-60** PKG\_UTIL.MATCH\_EDIT\_DISTANCE\_SIMILARITY parameters

| Parameter | Description              |
|-----------|--------------------------|
| str1      | First character string.  |
| str2      | Second character string. |

- PKG\_UTIL.RAW\_CAST\_TO\_VARCHAR2

Converts the raw type to the varchar2 type.

The prototype of the PKG\_UTIL.RAW\_CAST\_TO\_VARCHAR2 function is as follows:

```
PKG_UTIL.RAW_CAST_TO_VARCHAR2(
 str raw
)
RETURN varchar2
```

**Table 10-61** PKG\_UTIL.RAW\_CAST\_TO\_VARCHAR2 parameters

| Parameter | Description        |
|-----------|--------------------|
| str       | Hexadecimal string |

- PKG\_UTIL.SESSION\_CLEAR\_CONTEXT

Clears the session context.

The prototype of the PKG\_UTIL.SESSION\_CLEAR\_CONTEXT function is as follows:

```
PKG_UTIL.SESSION_CLEAR_CONTEXT(
 namespace text,
 client_identifier text,
 attribute text
)
RETURN INTEGER
```

**Table 10-62** PKG\_UTIL.SESSION\_CLEAR\_CONTEXT parameters

| Parameter         | Description                                                                                                                                                     |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| namespace         | Namespace of an attribute.                                                                                                                                      |
| client_identifier | Usually the value of this parameter is the same as that of <b>namespace</b> . If this parameter is set to <b>null</b> , all namespaces are modified by default. |
| attribute         | Attribute to be cleared.                                                                                                                                        |

- PKG\_UTIL.SESSION\_SEARCH\_CONTEXT

Searches for an attribute value.

The prototype of the PKG\_UTIL.SESSION\_SEARCH\_CONTEXT function is as follows:

```
PKG_UTIL.SESSION_SEARCH_CONTEXT(
 namespace text,
 attribute text
)
RETURN INTEGER
```

**Table 10-63** PKG\_UTIL.SESSION\_SEARCH\_CONTEXT parameters

| Parameter | Description                |
|-----------|----------------------------|
| namespace | Namespace of an attribute. |
| attribute | Attribute to be cleared.   |

- **PKG\_UTIL.SESSION\_SET\_CONTEXT**

Sets the attribute value.

The prototype of the PKG\_UTIL.SESSION\_SET\_CONTEXT function is as follows:

```
PKG_UTIL.SESSION_SET_CONTEXT(
namespace text,
attribute text,
value text
)
RETURN INTEGER
```

**Table 10-64** PKG\_UTIL.SESSION\_SET\_CONTEXT parameters

| Parameter | Description                |
|-----------|----------------------------|
| namespace | Namespace of an attribute. |
| attribute | Attribute to be set.       |
| value     | Attribute value.           |

- **PKG\_UTIL.UTILITY\_GET\_TIME**

This function prints the Unix timestamp.

The prototype of the PKG\_UTIL.UTILITY\_GET\_TIME function is as follows:

```
PKG_UTIL.UTILITY_GET_TIME()
RETURN bigint
```

- **PKG\_UTIL.UTILITY\_FORMAT\_ERROR\_BACKTRACE**

Displays the error stack of a stored procedure.

The prototype of the PKG\_UTIL.UTILITY\_FORMAT\_ERROR\_BACKTRACE function is as follows:

```
PKG_UTIL.UTILITY_FORMAT_ERROR_BACKTRACE()
RETURN text
```

- **PKG\_UTIL.UTILITY\_FORMAT\_ERROR\_STACK**

Displays the error information about a stored procedure.

The prototype of the PKG\_UTIL.UTILITY\_FORMAT\_ERROR\_STACK function is as follows:

```
PKG_UTIL.UTILITY_FORMAT_ERROR_STACK()
RETURN text
```

- **PKG\_UTIL.UTILITY\_FORMAT\_CALL\_STACK**

Displays the call stack of a stored procedure.

The prototype of the PKG\_UTIL.UTILITY\_FORMAT\_CALL\_STACK function is as follows:

```
PKG_UTIL.UTILITY_FORMAT_CALL_STACK()
RETURN text
```

## 10.12.2 Secondary Encapsulation APIs (Recommended)

### 10.12.2.1 DBE\_LOB

#### API Description

**Table 10-65** lists all APIs supported by the **DBE\_LOB** package.

**Table 10-65** DBE\_LOB

| API                             | Description                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>DBE_LOB.GET_LENGTH</b>       | Obtains and returns the specified length of a LOB.                                                                                                             |
| <b>DBE_LOB.OPEN</b>             | Opens a LOB and returns a LOB descriptor.                                                                                                                      |
| <b>DBE_LOB.READ</b>             | Loads a part of LOB content to the buffer based on the specified length and initial position offset.                                                           |
| <b>DBE_LOB.WRITE</b>            | Copies content in the buffer to a LOB based on the specified length and initial position offset.                                                               |
| <b>DBE_LOB.WRITE_APPEND</b>     | Copies content in the buffer to the end part of a LOB based on the specified length.                                                                           |
| <b>DBE_LOB.COPY</b>             | Copies content in a BLOB to another BLOB based on the specified length and initial position offset.                                                            |
| <b>DBE_LOB.ERASE</b>            | Deletes content in a BLOB based on the specified length and initial position offset.                                                                           |
| <b>DBE_LOB.CLOSE</b>            | Closes a LOB descriptor.                                                                                                                                       |
| <b>DBE_LOB.MATCH</b>            | Returns the position of the <i>N</i> th occurrence of a character string in a LOB.                                                                             |
| <b>DBE_LOB.COMPARE</b>          | Compares two LOBs or a certain part of two LOBs.                                                                                                               |
| <b>DBE_LOB.SUBSTR</b>           | Reads the substring of a LOB and returns the number of read bytes or the number of characters.                                                                 |
| <b>DBE_LOB.STRIP</b>            | Truncates the LOB of a specified length. After the execution is complete, the length of the LOB is set to the length specified by the <b>newlen</b> parameter. |
| <b>DBE_LOB.CREATE_TEMPORARY</b> | Creates a temporary BLOB or CLOB.                                                                                                                              |
| <b>DBE_LOB.APPEND</b>           | Adds the content of a LOB to another LOB.                                                                                                                      |
| <b>DBE_LOB.FREETEMPORARY</b>    | Deletes a temporary BLOB or CLOB.                                                                                                                              |
| <b>DBE_LOB.FILEOPEN</b>         | Opens a database-external file and returns a file descriptor.                                                                                                  |
| <b>DBE_LOB.FILECLOSE</b>        | Closes an external file opened by <b>FILEOPEN</b> .                                                                                                            |
| <b>DBE_LOB.LOADFROMFILE</b>     | Reads a database-external file to a BLOB file.                                                                                                                 |

| API                                      | Description                                    |
|------------------------------------------|------------------------------------------------|
| <a href="#">DBE_LOB.LOADBLOBFROMFILE</a> | Reads a database-external file to a BLOB file. |
| <a href="#">DBE_LOB.LOADCLOBFROMFILE</a> | Reads a database-external file to a CLOB file. |
| <a href="#">DBE_LOB.CONVERTTOBLOB</a>    | Converts a CLOB file to a BLOB file.           |
| <a href="#">DBE_LOB.CONVERTTOCLOB</a>    | Converts a BLOB file to a CLOB file.           |

- DBE\_LOB.GET\_LENGTH

The stored procedure GET\_LENGTH obtains and returns the size of a specified LOB. The maximum size is 2 GB.

The prototype of the DBE\_LOB.GET\_LENGTH function is as follows:

```
DBE_LOB.GET_LENGTH (
lob IN BLOB)
RETURN INTEGER;
```

```
DBE_LOB.GET_LENGTH (
lob IN CLOB)
RETURN INTEGER;
```

**Table 10-66** DBE\_LOB.GET\_LENGTH parameters

| Parameter | Description                              |
|-----------|------------------------------------------|
| lob       | BLOB/CLOB whose length is to be obtained |

- DBE\_LOB.OPEN

Opens a LOB and returns a LOB descriptor. This procedure is meaningless and is used only for compatibility.

The prototype of the DBE\_LOB.OPEN function is as follows:

```
DBE_LOB.OPEN (
lob INOUT BLOB
);
```

```
DBE_LOB.OPEN (
lob INOUT CLOB
);
```

```
DBE_LOB.OPEN (
bfile db_lob.bfile,
open_mode text DEFAULT 'null':text
);
```

**Table 10-67** DBE\_LOB.OPEN parameters

| Parameter | Description                  |
|-----------|------------------------------|
| lob       | BLOB or CLOB that is opened. |

- DBE\_LOB.READ

The stored procedure READ loads a part of LOB content to the buffer based on the specified length and initial position offset.

The prototype of the DBE\_LOB.READ function is as follows:

```
DBE_LOB.READ (
lob IN BLOB,
len IN INTEGER,
start IN INTEGER,
buffer OUT RAW);

DBE_LOB.READ (
lob IN CLOB,
len INOUT INTEGER,
start IN INTEGER,
buffer OUT VARCHAR2);
```

**Table 10-68** DBE\_LOB.READ parameters

| Parameter | Description                                                                                                                                                                                 |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lob       | BLOB/CLOB to be read                                                                                                                                                                        |
| len       | Read size. A maximum of 32767 characters are supported.<br><b>NOTE</b><br>If the length is negative, the error message "ERROR: argument 2 is null, invalid, or out of range." is displayed. |
| start     | Indicates where to start reading the LOB content, that is, the offset bytes to initial position of LOB content.                                                                             |
| buffer    | Target buffer to store the read LOB content.                                                                                                                                                |

The value of **lob** can be a CLOB with a size greater than 1 GB. The maximum size is 2 GB. The maximum value of **buffer** is 32 KB and cannot be greater than the value of **len**.

- DBE\_LOB.WRITE

The stored procedure WRITE copies content in the buffer to LOB variables based on the specified length and initial position.

The prototype of the DBE\_LOB.WRITE function is as follows:

```
DBE_LOB.WRITE (
dest_lob INOUT BLOB,
len IN INTEGER,
start IN INTEGER,
src_lob IN RAW);

DBE_LOB.WRITE (
dest_lob INOUT CLOB,
len IN INTEGER,
```

```
start IN INTEGER,
src_lob IN VARCHAR2);
```

**Table 10-69** DBE\_LOB.WRITE parameters

| Parameter | Description                                                                                                                                                                                                                                                                                    |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob  | BLOB/CLOB to which the buffer content is written                                                                                                                                                                                                                                               |
| len       | Length to write, up to 32767 characters.<br><b>NOTE</b><br>If the length to write is less than 1 or greater than the length of the content to be written, an error is reported.                                                                                                                |
| start     | Indicates where to start writing the LOB content, that is, the offset bytes to initial position of LOB content.<br><b>NOTE</b><br>If the offset value is less than 1, an error is reported. If the offset value is greater than the maximum length of LOB type contents, no error is reported. |
| src_lob   | Content to be written                                                                                                                                                                                                                                                                          |

The value of **dest\_lob** can be a CLOB with a size greater than 1 GB. The maximum size is 2 GB. The maximum value of **src\_lob** is 1 GB.

- DBE\_LOB.WRITE\_APPEND

The stored procedure WRITE\_APPEND copies content in the buffer to the end part of a LOB based on the specified length.

The prototype of the DBE\_LOB.WRITE\_APPEND function is as follows:

```
DBE_LOB.WRITE_APPEND (
dest_lob INOUT BLOB,
len IN INTEGER,
src_lob IN RAW);

DBE_LOB.WRITE_APPEND (
dest_lob INOUT CLOB,
len IN INTEGER,
src_lob IN VARCHAR2);
```

**Table 10-70** DBE\_LOB.WRITE\_APPEND parameters

| Parameter | Description                                                                                                                                                                     |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob  | BLOB/CLOB to which the buffer content is written                                                                                                                                |
| len       | Length to write, up to 32767 characters.<br><b>NOTE</b><br>If the length to write is less than 1 or greater than the length of the content to be written, an error is reported. |
| src_lob   | Content to be written                                                                                                                                                           |

- DBE\_LOB.COPY



The stored procedure COPY copies content in a BLOB to another BLOB based on the specified length and initial position offset.

The prototype of the DBE\_LOB.COPY function is as follows:

```
DBE_LOB.COPY (
dest_lob INOUT BLOB,
src_lob IN BLOB,
len IN INTEGER,
dest_start IN INTEGER DEFAULT 1,
src_start IN INTEGER DEFAULT 1);
```

**Table 10-71** DBE\_LOB.COPY parameters

| Parameter  | Description                                                                                                                                                                                                                          |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob   | BLOB to be pasted.                                                                                                                                                                                                                   |
| src_lob    | BLOB to be copied.                                                                                                                                                                                                                   |
| len        | Length of copied content<br><b>NOTE</b><br>If the length of copied content is shorter than 1 or longer than the maximum length of BLOB, an error is reported.                                                                        |
| dest_start | Indicates where to start pasting the BLOB content, that is, the offset bytes to initial position of BLOB content.<br><b>NOTE</b><br>If the offset is shorter than 1 or longer than the maximum length of BLOB, an error is reported. |
| src_start  | Indicates where to start copying the BLOB content, that is, the offset bytes to initial position of BLOB content.<br><b>NOTE</b><br>If the offset is shorter than 1 or longer than the length of source BLOB, an error is reported.  |

- DBE\_LOB.ERASE

The stored procedure ERASE deletes content in BLOB based on the specified length and initial position offset.

The prototype of the DBE\_LOB.ERASE function is as follows:

```
DBE_LOB.ERASE (
lob INOUT BLOB,
len INOUT INTEGER,
start IN INTEGER DEFAULT 1);
```

**Table 10-72** DBE\_LOB.ERASE parameters

| Parameter | Description                                                                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lob       | BLOB whose content is to be deleted.                                                                                                                                   |
| len       | Length of content to be deleted.<br><b>NOTE</b><br>If the length of deleted content is shorter than 1 or longer than the maximum length of BLOB, an error is reported. |

| Parameter | Description                                                                                                                                                                                                                           |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| start     | Indicates where to start deleting the BLOB content, that is, the offset bytes to initial position of BLOB content.<br><b>NOTE</b><br>If the offset is shorter than 1 or longer than the maximum length of BLOB, an error is reported. |

- DBE\_LOB.CLOSE

Closes the LOB descriptor that has been opened.

The prototype of the DBE\_LOB.CLOSE function is as follows:

```
DBE_LOB.CLOSE(
lob IN BLOB);

DBE_LOB.CLOSE (
lob IN CLOB);

DBE_LOB.CLOSE (
file IN INTEGER);
```

**Table 10-73** DBE\_LOB.CLOSE parameters

| Parameter | Description               |
|-----------|---------------------------|
| lob       | BLOB/CLOB to be disabled. |

- DBE\_LOB.MATCH

This function returns the *N*th occurrence position of **pattern** in a LOB. **NULL** is returned for any of the following conditions: offset < 1 or offset > **LOBMAXSIZE**; *n*th < 1 or *n*th > **LOBMAXSIZE**.

The prototype of the DBE\_LOB.MATCH function is as follows:

```
DBE_LOB.MATCH (
lob IN BLOB,
pattern IN RAW,
start_index IN INTEGER DEFAULT 1,
match_index IN INTEGER DEFAULT 1)
RETURN INTEGER;

DBE_LOB.MATCH (
lob IN CLOB,
pattern IN VARCHAR2,
start_index IN INTEGER DEFAULT 1,
match_index IN INTEGER DEFAULT 1)
RETURN INTEGER;
```

**Table 10-74** DBE\_LOB.MATCH API parameters

| Parameter | Description                                             |
|-----------|---------------------------------------------------------|
| lob       | BLOB/CLOB descriptor to be searched for.                |
| pattern   | Pattern to match. It is RAW for BLOB and TEXT for CLOB. |

| Parameter   | Description                                                                                                                               |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| start_index | For BLOB, the absolute offset is in the unit of byte. For CLOB, the offset is in the unit of character. The matching start position is 1. |
| match_index | Number of pattern matching times. The minimum value is 1.                                                                                 |

- DBE\_LOB.COMPARE

This function compares two LOBs or a certain part of two LOBs.

- If the compared objects are equal, **0** is returned. Otherwise, a non-zero value is returned.
- If the first LOB is smaller than the second, **-1** is returned. If the first LOB is larger than the second, **1** is returned.
- If any of the *len*, *start1*, and *start2* parameters is invalid, **NULL** is returned. The valid offset range is 1 to *LOBMAXSIZE*.

The prototype of the DBE\_LOB.COMPARE function is as follows:

```
DBE_LOB.COMPARE (
lob1 IN BLOB,
lob2 IN BLOB,
len IN INTEGER DEFAULT DBE_LOB.LOBMAXSIZE,
start1 IN INTEGER DEFAULT 1,
start2 IN INTEGER DEFAULT 1)
RETURN INTEGER;

DBE_LOB.COMPARE (
lob1 IN CLOB,
lob2 IN CLOB,
len IN INTEGER DEFAULT DBE_LOB.LOBMAXSIZE,
start1 IN INTEGER DEFAULT 1,
start2 IN INTEGER DEFAULT 1)
RETURN INTEGER;
```

**Table 10-75** DBE\_LOB.COMPARE parameters

| Parameter | Description                                                                                    |
|-----------|------------------------------------------------------------------------------------------------|
| lob1      | First BLOB/CLOB to be compared.                                                                |
| lob2      | Second BLOB/CLOB to be compared.                                                               |
| len       | Number of characters or bytes to be compared. The maximum value is <b>DBE_LOB.LOBMAXSIZE</b> . |
| start1    | Offset of the first LOB descriptor. The initial position is 1.                                 |
| start2    | Offset of the second LOB descriptor. The initial position is 1.                                |

- DBE\_LOB.SUBSTR

This function reads the substring of a LOB and returns the number of read bytes or the number of characters. **NULL** is returned for any of the following conditions: value of **amount** < 1 or value of **amount** > 32767; value of **offset** < 1 or value of **offset** > value of **LOBMAXSIZE**

The prototype of the DBE\_LOB.SUBSTR function is as follows:

```
DBE_LOB.SUBSTR (
lob IN BLOB,
len IN INTEGER DEFAULT 32767,
start IN INTEGER DEFAULT 1)
RETURN RAW;

DBE_LOB.SUBSTR (
lob IN CLOB,
len IN INTEGER DEFAULT 32767,
start IN INTEGER DEFAULT 1)
RETURN VARCHAR2;
```

**Table 10-76** DBE\_LOB.SUBSTR parameters

| Parameter | Description                                                                                                                                                 |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lob       | LOB descriptor of the substring to be read. For BLOB, the return value is the number of read bytes. For CLOB, the return value is the number of characters. |
| len       | Number of bytes or characters to be read.                                                                                                                   |
| start     | Number of characters or bytes offset from the start position.                                                                                               |

- DBE\_LOB.STRIP

Truncates a LOB based on a specified length. After this stored procedure is executed, the length of the LOB is set to the length specified by the **newlen** parameter. If an empty LOB is truncated, no execution result is displayed. If the specified length is longer than the length of the LOB, an exception occurs.

The prototype of the DBE\_LOB.STRIP function is as follows:

```
DBE_LOB.STRIP (
lob IN OUT BLOB,
newlen IN INTEGER);

DBE_LOB.STRIP (
lob IN OUT CLOB,
newlen IN INTEGER);
```

**Table 10-77** DBE\_LOB.STRIP parameters

| Parameter | Description                                                                 |
|-----------|-----------------------------------------------------------------------------|
| lob       | BLOB to be read.                                                            |
| newlen    | New length after truncation, in bytes for BLOBs or in characters for CLOBs. |

- DBE\_LOB.CREATE\_TEMPORARY

Creates a temporary BLOB or CLOB. This API is used only for syntax compatibility.

The prototype of the DBE\_LOB.CREATE\_TEMPORARY function is as follows:

```
DBE_LOB.CREATE_TEMPORARY (
locator INOUT BLOB,
```

```
cache IN BOOLEAN,
keep_alive_time IN INTEGER);

DBE_LOB.CREATE_TEMPORARY (
locator INOUT CLOB,
cache IN BOOLEAN,
keep_alive_time IN INTEGER);
```

**Table 10-78** DBE\_LOB.CREATE\_TEMPORARY parameters

| Parameter       | Description                         |
|-----------------|-------------------------------------|
| locator         | LOB descriptor                      |
| cache           | Used only for syntax compatibility. |
| keep_alive_time | Used only for syntax compatibility. |

- DBE\_LOB.APPEND

The stored procedure APPEND loads a part of BLOB content to the buffer based on the specified length and initial position offset.

The prototype of the DBE\_LOB.APPEND function is as follows:

```
DBE_LOB.APPEND (
dest_lob INOUT BLOB,
src_lob IN BLOB);

DBE_LOB.APPEND (
dest_lob INOUT CLOB,
src_lob IN CLOB);
```

**Table 10-79** DBE\_LOB.APPEND parameters

| Parameter | Description                                            |
|-----------|--------------------------------------------------------|
| dest_lob  | BLOB/CLOB to which the buffer content is to be written |
| src_lob   | BLOB/CLOB from which the buffer content is to be read  |

The value of **dest\_lob** can be a CLOB with a size greater than 1 GB. The maximum size is 2 GB. The maximum value of **src\_lob** is 1 GB.

- DBE\_LOB.FREETEMPORARY

Frees LOB files created by CREATE\_TEMPORARY.

The prototype of the DBE\_LOB.FREETEMPORARY function is as follows:

```
DBE_LOB.FREETEMPORARY (
lob_loc INOUT BLOB);

DBE_LOB.FREETEMPORARY (
lob_loc INOUT CLOB);
```

**Table 10-80** DBE\_LOB.FREETEMPORARY parameters

| Parameter | Description            |
|-----------|------------------------|
| lob_loc   | BLOB/CLOB to be freed. |

- DBE\_LOB.FILEOPEN

This function is used to open a database-external file of the BFILE type and return the file descriptor corresponding to the file.

The BFILE type is defined as follows:

```
DBE_LOB.BFILE (
 directory text,
 filename text);
```

The prototype of the DBE\_LOB.FILEOPEN function is as follows:

```
DBE_LOB.FILEOPEN (
 file IN DBE_LOB.BFILE,
 open_mode IN text)
RETURN integer;
```

**Table 10-81** DBE\_LOB.FILEOPEN parameters

| Parameter | Description                                                                                             |
|-----------|---------------------------------------------------------------------------------------------------------|
| file      | Specifies the database-external file to be opened. The BFILE type contains the file path and file name. |
| open_mode | Specifies the file open mode ( <b>w</b> , <b>r</b> , or <b>a</b> ).                                     |

- DBE\_LOB.FILECLOSE

Closes an external BFILE file.

The prototype of the DBE\_LOB.FILECLOSE function is as follows:

```
DBE_LOB.FILECLOSE (
 file IN integer);
```

**Table 10-82** DBE\_LOB.FILECLOSE parameters

| Parameter | Description                                                       |
|-----------|-------------------------------------------------------------------|
| file      | External database file to be closed (that is opened by FILEOPEN). |

- DBE\_LOB.LOADFROMFILE

This is used to read an external BFILE file to a BLOB file.

The prototype of the DBE\_LOB.LOADFROMFILE function is as follows:

```
DBE_LOB.LOADFROMFILE (
 dest_lob IN BLOB,
 src_file IN INTEGER,
 amount IN INTEGER,
 dest_offset IN INTEGER,
 src_offset IN INTEGER)
RETURN raw;
```

**Table 10-83** DBE\_LOB.LOADFROMFILE parameters

| Parameter | Description                                                 |
|-----------|-------------------------------------------------------------|
| dest_lob  | Target BLOB file. The BFILE file will be read to this file. |

| Parameter   | Description                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| src_bfile   | Source BFILE file to be read.                                                                                                                                      |
| amount      | Size of a BLOB file. If the size of a file exceeds this threshold, the file will not be saved to the BLOB file.                                                    |
| dest_offset | Offset length of the BLOB file. If <b>dest_offset</b> is set to <b>1</b> , data is loaded from the start position of the file. The rest may be deduced by analogy. |
| src_offset  | Offset length of the BFILE file. If <b>src_offset</b> is set to <b>1</b> , data is read from the start position of the file. The rest may be deduced by analogy.   |

The size of the file to be imported can exceed 1 GB. The maximum value of **amount** is the value of **LOBMAXSIZE** (4 GB). When the value of **src\_offset** or **dest\_offset** is greater than that of **LOBMAXSIZE**, or the value of **amount** is greater than that of **LOBMAXSIZE**, an error is reported for the advanced package.

- DBE\_LOB.LOADBLOBFROMFILE

This is used to read an external BFILE file to a BLOB file.

The prototype of the DBE\_LOB.LOADBLOBFROMFILE function is as follows:

```
DBE_LOB.LOADBLOBFROMFILE (
dest_lob IN BLOB,
src_file IN INTEGER,
amount IN INTEGER,
dest_offset IN INTEGER,
src_offset IN INTEGER)
RETURN raw;
```

**Table 10-84** DBE\_LOB.LOADBLOBFROMFILE parameters

| Parameter   | Description                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob    | Target BLOB file. The BFILE file will be read to this file.                                                                                                        |
| src_bfile   | Source BFILE file to be read.                                                                                                                                      |
| amount      | Length of a BLOB file. If the length of a file exceeds this threshold, the file will not be saved to the BLOB file.                                                |
| dest_offset | Offset length of the BLOB file. If <b>dest_offset</b> is set to <b>1</b> , data is loaded from the start position of the file. The rest may be deduced by analogy. |
| src_offset  | Offset length of the BFILE file. If <b>src_offset</b> is set to <b>1</b> , data is read from the start position of the file. The rest may be deduced by analogy.   |

- DBE\_LOB.LOADCLOBFROMFILE

This is used to read an external BFILE file to a CLOB file.

The prototype of the DBE\_LOB.LOADCLOBFROMFILE function is as follows:

```
DBE_LOB.LOADCLOBFROMFILE (
dest_lob IN CLOB,
src_file IN INTEGER,
amount IN INTEGER,
dest_offset IN INTEGER,
src_offset IN INTEGER)
RETURN row;
```

**Table 10-85** DBE\_LOB.LOADCLOBFROMFILE parameters

| Parameter   | Description                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob    | Target CLOB file. The BFILE file will be read to this file.                                                                                                        |
| src_bfile   | Source BFILE file to be read.                                                                                                                                      |
| amount      | Size of a CLOB file. If the size of a file exceeds this threshold, the file will not be saved to the CLOB file.                                                    |
| dest_offset | Offset length of the CLOB file. If <b>dest_offset</b> is set to <b>1</b> , data is loaded from the start position of the file. The rest may be deduced by analogy. |
| src_offset  | Offset length of the BFILE file. If <b>src_offset</b> is set to <b>1</b> , data is read from the start position of the file. The rest may be deduced by analogy.   |

- DBE\_LOB.CONVERTTOBLOB

This function is used to convert a CLOB file to a BLOB file.

The prototype of the DBE\_LOB.CONVERTTOBLOB function is as follows:

```
DBE_LOB.CONVERTTOBLOB(
dest_blob IN BLOB,
src_clob IN CLOB,
amount IN INTEGER default 32767,
dest_offset IN INTEGER default 1,
src_offset IN INTEGER default 1)
RETURN row;
```

**Table 10-86** DBE\_LOB.CONVERTTOBLOB parameters

| Parameter   | Description                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob    | Target BLOB file, which is converted from a CLOB file.                                                                                                             |
| src_bfile   | Source CLOB file to be read.                                                                                                                                       |
| amount      | Size of a BLOB file. If the size of a file exceeds this threshold, the file will not be saved to the BLOB file.                                                    |
| dest_offset | Offset length of the BLOB file. If <b>dest_offset</b> is set to <b>1</b> , data is loaded from the start position of the file. The rest may be deduced by analogy. |
| src_offset  | Offset length of the CLOB file. If <b>src_offset</b> is set to <b>1</b> , data is read from the start position of the file. The rest may be deduced by analogy.    |



- DBE\_LOB.CONVERTTOCLOB

This function is used to convert a BLOB file to a CLOB file.

The prototype of the DBE\_LOB.CONVERTTOCLOB function is as follows:

```
DBE_LOB.CONVERTTOCLOB(
dest_clob IN CLOB,
src_blob IN BLOB,
amount IN INTEGER default 32767,
dest_offset IN INTEGER default 1,
src_offset IN INTEGER default 1)
RETURN text;
```

**Table 10-87** DBE\_LOB.CONVERTTOCLOB parameters

| Parameter   | Description                                                                                                                                                        |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_lob    | Target CLOB file, which is converted from a BLOB file.                                                                                                             |
| src_bfile   | Source BLOB file to be read.                                                                                                                                       |
| amount      | Size of a CLOB file. If the size of a file exceeds this threshold, the file will not be saved to the CLOB file.                                                    |
| dest_offset | Offset length of the CLOB file. If <b>dest_offset</b> is set to <b>1</b> , data is loaded from the start position of the file. The rest may be deduced by analogy. |
| src_offset  | Offset length of the BLOB file. If <b>src_offset</b> is set to <b>1</b> , data is read from the start position of the file. The rest may be deduced by analogy.    |

## Examples

```
-- Obtain the length of a string.
SELECT DBE_LOB.GET_LENGTH('12345678');
get_length

8
(1 row)
DECLARE
myraw RAW(100);
amount INTEGER :=2;
buffer INTEGER :=1;
begin
DBE_LOB.READ('123456789012345',amount,buffer,myraw);
dbe_output.print_line(myraw);
end;
/
0123
ANONYMOUS BLOCK EXECUTE

CREATE TABLE blob_Table (t1 blob);
CREATE TABLE blob_Table_bak (t2 blob);
INSERT INTO blob_Table VALUES('abcdef');
INSERT INTO blob_Table_bak VALUES('22222');

DECLARE
str varchar2(100) := 'abcdef';
source raw(100);
dest blob;
copyto blob;
amount int;
PSV_SQL varchar2(100);
```

```
PSV_SQL1 varchar2(100);
a int :=1;
len int;
BEGIN
source := dbc_raw.cast_from_varchar2_to_raw(str);
amount := dbc_raw.get_length(source);

PSV_SQL := 'select * from blob_Table for update';
PSV_SQL1 := 'select * from blob_Table_bak for update';

EXECUTE IMMEDIATE PSV_SQL into dest;
EXECUTE IMMEDIATE PSV_SQL1 into copyto;

DBC_LOB.WRITE(dest, amount, 1, source);
DBC_LOB.WRITE_APPEND(dest, amount, source);

DBC_LOB.ERASE(dest, a, 1);
DBC_OUTPUT.PRINT_LINE(a);
DBC_LOB.COPY(copyto, dest, amount, 10, 1);
perform DBC_LOB.CLOSE(dest);
RETURN;
END;
/

-- Delete the table.
DROP TABLE blob_Table;
DROP TABLE
DROP TABLE blob_Table_bak;
DROP TABLE
```

### 10.12.2.2 DBE\_RANDOM

#### API Description

**Table 10-88** lists all APIs supported by the **DBE\_RANDOM** package.

**Table 10-88** DBE\_RANDOM API parameters

| API                         | Description                                                                            |
|-----------------------------|----------------------------------------------------------------------------------------|
| <b>DBE_RANDOM.SET_SEED</b>  | Sets a seed for a random number.                                                       |
| <b>DBE_RANDOM.GET_VALUE</b> | Generates a random number between a specified <b>low</b> and a specified <b>high</b> . |

- DBE\_RANDOM.SET\_SEED

The stored procedure SEED is used to set a seed for a random number. The prototype of the DBE\_RANDOM.SET\_SEED function is as follows:

```
DBE_RANDOM.SET_SEED (seed IN INTEGER);
```

**Table 10-89** DBE\_RANDOM.SET\_SEED API parameters

| Parameter | Description                           |
|-----------|---------------------------------------|
| seed      | Generates a seed for a random number. |

- DBE\_RANDOM.GET\_VALUE

The GET\_VALUE function generates a random number between a specified **low** and a specified **high**. The prototype of the DBE\_RANDOM.GET\_VALUE function is as follows:

```
DBE_RANDOM.GET_VALUE(
min IN NUMBER default 0,
max IN NUMBER default 1)
RETURN NUMBER;
```

**Table 10-90** DBE\_RANDOM.GET\_VALUE API parameters

| Parameter | Description                                                                                                  |
|-----------|--------------------------------------------------------------------------------------------------------------|
| min       | Sets the low bound for a random number. The generated random number is greater than or equal to <b>min</b> . |
| max       | Sets the high bound for a random number. The generated random number is less than <b>max</b> .               |

 **NOTE**

- The only requirement is that the parameter type is numeric regardless of the right and left bound values.
- DBE\_RANDOM implements pseudo-random numbers. Therefore, if the initial value (seed) remains unchanged, the sequence of the pseudo-random numbers also remains unchanged.
- The generated random number contains 15 valid digits.

## Examples

```
-- Generate a random number between 0 and 1.
SELECT DBE_RANDOM.GET_VALUE(0,1);
get_value

.917468812743886
(1 row)
-- For integers within a specified range, add the arguments min and max, and truncate the decimals from
-- the result (the maximum value is not included as a possible value). Therefore, for integers from 0 to 99, you
-- can use the following code:
SELECT TRUNC(DBE_RANDOM.GET_VALUE(0,100));
trunc

26
(1 row)
```

### 10.12.2.3 DBE\_OUTPUT

#### API Description

[Table 10-91](#) provides all APIs supported by the **DBE\_OUTPUT** package.

**Table 10-91** DBE\_OUTPUT

| API                                        | Description                                                                                                                                                                                                                        |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">DBE_OUTPUT.PRINT_LINE</a>      | Outputs the specified text with newline characters.                                                                                                                                                                                |
| <a href="#">DBE_OUTPUT.PRINT</a>           | Outputs the specified text without newline characters.                                                                                                                                                                             |
| <a href="#">DBE_OUTPUT.SET_BUFFER_SIZE</a> | Sets the size of the output buffer. If the size is not specified, the buffer can contain a maximum of 20000 bytes. If the size is set to a value less than or equal to 2000 bytes, the buffer can contain a maximum of 2000 bytes. |

- [DBE\\_OUTPUT.PRINT\\_LINE](#)

The stored procedure **PRINT\_LINE** writes a row of text carrying a line end symbol in the buffer. The function prototype of [DBE\\_OUTPUT.PRINT\\_LINE](#) is as follows:

```
DBE_OUTPUT.PRINT_LINE (
format IN VARCHAR2);
```

**Table 10-92** DBE\_OUTPUT.PRINT\_LINE API parameters

| Parameter | Description                                        |
|-----------|----------------------------------------------------|
| format    | Specifies the text that was written to the buffer. |

- [DBE\\_OUTPUT.PRINT](#)

The stored procedure **PRINT** outputs the specified text to the front of the specified text without adding a newline character. The prototype of the [DBE\\_OUTPUT.PRINT](#) function is as follows:

```
DBE_OUTPUT.PRINT (
format IN VARCHAR2);
```

**Table 10-93** DBE\_OUTPUT.PRINT API parameters

| Parameter | Description                                                |
|-----------|------------------------------------------------------------|
| format    | Specifies the text that was written to the specified text. |

- [DBE\\_OUTPUT.SET\\_BUFFER\\_SIZE](#)

The stored procedure **SET\_BUFFER\_SIZE** sets the output buffer size. If the size is not specified, it contains a maximum of 20000 bytes. The function prototype of [DBE\\_OUTPUT.SET\\_BUFFER\\_SIZE](#) is as follows:

```
DBE_OUTPUT.SET_BUFFER_SIZE (
size IN INTEGER default 20000);
```

**Table 10-94** DBE\_OUTPUT.SET\_BUFFER\_SIZE API parameters

| Parameter | Description                  |
|-----------|------------------------------|
| size      | Sets the output buffer size. |

## Examples

```
BEGIN
 DBE_OUTPUT.SET_BUFFER_SIZE(50);
 DBE_OUTPUT.PRINT('hello, ');
 DBE_OUTPUT.PRINT_LINE('database!');-- Output "hello, database!".
END;
/
```

### 10.12.2.4 DBE\_RAW

#### API Description

**Table 10-95** provides all APIs supported by the **DBE\_RAW** package.

**Table 10-95** DBE\_RAW

| API                                                     | Description                                                                                 |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------|
| <a href="#">DBE_RAW.CAST_FROM_BINARY_INTEGER_TO_RAW</a> | Converts a value of the <b>INTEGER</b> type to a binary representation ( <b>RAW</b> type).  |
| <a href="#">DBE_RAW.CAST_FROM_RAW_TO_BINARY_INTEGER</a> | Converts a binary representation ( <b>RAW</b> type) to a value of the <b>INTEGER</b> type.  |
| <a href="#">DBE_RAW.GET_LENGTH</a>                      | Obtains the length of a <b>RAW</b> object.                                                  |
| <a href="#">DBE_RAW.CAST_FROM_VARCHAR2_TO_RAW</a>       | Converts a value of the <b>VARCHAR2</b> type to a binary representation ( <b>RAW</b> type). |
| <a href="#">DBE_RAW.CAST_TO_VARCHAR2</a>                | Converts a value of the <b>RAW</b> type to a value of the <b>VARCHAR2</b> type.             |
| <a href="#">DBE_RAW.BIT_OR</a>                          | Performs the bitwise OR operation on raw data.                                              |
| <a href="#">DBE_RAW.SUBSTR</a>                          | Returns the substring of the <b>RAW</b> type.                                               |

#### NOTICE

RAW data is represented as hexadecimal characters externally, and stored as binary characters internally. For example, the representation of RAW data 11001011 is 'CB', that is, the input for type conversion is 'CB'.

- DBE\_RAW.CAST\_FROM\_BINARY\_INTEGER\_TO\_RAW

The stored procedure **CAST\_FROM\_BINARY\_INTEGER\_TO\_RAW** converts a value of the **INTEGER** type to a binary representation (**RAW** type).

The prototype of the DBE\_RAW.CAST\_FROM\_BINARY\_INTEGER\_TO\_RAW function is as follows:

```
DBE_RAW.CAST_FROM_BINARY_INTEGER_TO_RAW (
value IN INTEGER,
endianess IN INTEGER DEFAULT 1)
RETURN RAW;
```

**Table 10-96** DBE\_RAW.CAST\_FROM\_BINARY\_INTEGER\_TO\_RAW parameters

| Parameter | Description                                                                                                                                                          |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| value     | Specifies the <b>INTEGER</b> value to be converted to the <b>RAW</b> value.                                                                                          |
| endianess | Specifies the <b>INTEGER</b> value <b>1</b> or <b>2</b> for the byte sequence. ( <b>1</b> indicates <b>BIG_ENDIAN</b> and <b>2</b> indicates <b>LITTLE-ENDIAN</b> .) |

- DBE\_RAW.CAST\_FROM\_RAW\_TO\_BINARY\_INTEGER

The stored procedure **CAST\_FROM\_RAW\_TO\_BINARY\_INTEGER** converts a binary representation (**RAW** type) to a value of the **INTEGER** type.

The prototype of the DBE\_RAW.CAST\_FROM\_RAW\_TO\_BINARY\_INTEGER function is as follows:

```
DBE_RAW.CAST_FROM_RAW_TO_BINARY_INTEGER (
value IN RAW,
endianess IN INTEGER DEFAULT 1)
RETURN BINARY_INTEGER;
```

**Table 10-97** DBE\_RAW.CAST\_FROM\_RAW\_TO\_BINARY\_INTEGER parameters

| Parameter | Description                                                                                                                                                          |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| value     | Specifies an <b>INTEGER</b> value in a binary representation ( <b>RAW</b> type).                                                                                     |
| endianess | Specifies the <b>INTEGER</b> value <b>1</b> or <b>2</b> for the byte sequence. ( <b>1</b> indicates <b>BIG_ENDIAN</b> and <b>2</b> indicates <b>LITTLE-ENDIAN</b> .) |

- DBE\_RAW.GET\_LENGTH

The stored procedure **GET\_LENGTH** returns the length of a **RAW** object.

The prototype of the DBE\_RAW.GET\_LENGTH function is as follows:

```
DBE_RAW.GET_LENGTH(
value IN RAW)
RETURN INTEGER;
```

**Table 10-98** DBE\_RAW.GET\_LENGTH parameters

| Parameter | Description                    |
|-----------|--------------------------------|
| value     | Specifies a <b>RAW</b> object. |

- DBE\_RAW.CAST\_FROM\_VARCHAR2\_TO\_RAW

The stored procedure **CAST\_FROM\_VARCHAR2\_TO\_RAW** converts a **VARCHAR2** object to a **RAW** object.

The prototype of the DBE\_RAW.CAST\_FROM\_VARCHAR2\_TO\_RAW function is as follows:

```
DBE_RAW.CAST_TO_RAW(
str IN VARCHAR2)
RETURN RAW;
```

**Table 10-99** DBE\_RAW.CAST\_FROM\_VARCHAR2\_TO\_RAW parameters

| Parameter | Description                                         |
|-----------|-----------------------------------------------------|
| c         | Specifies a <b>VARCHAR2</b> object to be converted. |

- DBE\_RAW.CAST\_TO\_VARCHAR2

The stored procedure **CAST\_TO\_VARCHAR2** converts a **RAW** object to a **VARCHAR2** object.

The prototype of the DBE\_RAW.CAST\_TO\_VARCHAR2 function is as follows:

```
DBE_RAW.CAST_TO_VARCHAR2(
str IN RAW)
RETURN VARCHAR2;
```

**Table 10-100** DBE\_RAW.CAST\_TO\_VARCHAR2 parameters

| Parameter | Description                                    |
|-----------|------------------------------------------------|
| str       | Specifies a <b>RAW</b> object to be converted. |

- DBE\_RAW.BIT\_OR

The stored procedure **BIT\_OR** calculates the bitwise OR result of two raw data records.

The prototype of the DBE\_RAW.BIT\_OR function is as follows:

```
DBE_RAW.BIT_OR(
str1 IN RAW,
str2 IN RAW)
RETURN RAW;
```

**Table 10-101** DBE\_RAW.BIT\_OR parameters

| Parameter | Description                                         |
|-----------|-----------------------------------------------------|
| str1      | First character string of the bitwise OR operation  |
| str2      | Second character string of the bitwise OR operation |

- DBE\_RAW.SUBSTR

The stored procedure **SUBSTR** truncates an object of the **RAW** type based on the start bit and length.

The prototype of the DBE\_RAW.SUBSTR function is as follows:

```
DBE_RAW.SUBSTR(
 IN lob_loc raw,
 IN off_set integer default 1,
 IN amount integer default 32767)
RETURN RAW;
```

**Table 10-102** DBE\_RAW.SUBSTR parameters

| Parameter | Description                                                      |
|-----------|------------------------------------------------------------------|
| lob_loc   | Source raw character string.                                     |
| off_set   | Start position of the substring. The default value is <b>1</b> . |
| amount    | Substring length. The default value is <b>32767</b> .            |

## Examples

```
-- Perform operations on RAW data in a stored procedure.
CREATE OR REPLACE PROCEDURE proc_raw
AS
 str varchar2(100) := 'abcdef';
 source raw(100);
 amount integer;
BEGIN
 source := dbe_raw.cast_from_varchar2_to_raw(str);-- Convert the type.
 amount := dbe_raw.get_length(source);-- Obtain the length.
 dbe_output.print_line(amount);
END;
/

-- Call the stored procedure.
CALL proc_raw();

-- Delete the stored procedure.
DROP PROCEDURE proc_raw;
```

### 10.12.2.5 DBE\_TASK

## API Description

**Table 10-103** lists all APIs supported by the **DBE\_TASK** package.

**Table 10-103** DBE\_TASK

| API                    | Description                                                                   |
|------------------------|-------------------------------------------------------------------------------|
| <b>DBE_TASK.SUBMIT</b> | Commits a scheduled job. The job ID is automatically generated by the system. |



| API                                 | Description                                                                                                                      |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">DBE_TASK.JOB_SUBMIT</a> | Same as <a href="#">DBE_TASK.SUBMIT</a> . However, It provides syntax compatibility parameters.                                  |
| <a href="#">DBE_TASK.ID_SUBMIT</a>  | Commits a scheduled job. The job ID is specified by the user.                                                                    |
| <a href="#">DBE_TASK.CANCEL</a>     | Removes a scheduled task by job ID.                                                                                              |
| <a href="#">DBE_TASK.RUN</a>        | Executes a scheduled task.                                                                                                       |
| <a href="#">DBE_TASK.FINISH</a>     | Disables or enables scheduled task execution.                                                                                    |
| <a href="#">DBE_TASK.UPDATE</a>     | Modifies user-definable attributes of a scheduled task, including the task content, next-execution time, and execution interval. |
| <a href="#">DBE_TASK.CHANGE</a>     | Same as <a href="#">DBE_TASK.UPDATE</a> . However, It provides syntax compatibility parameters.                                  |
| <a href="#">DBE_TASK.CONTENT</a>    | Modifies the job content attribute of a scheduled job.                                                                           |
| <a href="#">DBE_TASK.NEXT_TIME</a>  | Modifies the next-execution time attribute of a scheduled task.                                                                  |
| <a href="#">DBE_TASK.INTERVAL</a>   | Modifies the execution interval attribute of a scheduled task.                                                                   |

- **DBE\_TASK.SUBMIT**

The stored procedure SUBMIT commits a scheduled task provided by the system.

The prototype of the DBE\_TASK.SUBMIT function is as follows:

```
DBE_TASK.SUBMIT(
 what IN TEXT,
 next_time IN TIMESTAMP DEFAULT sysdate,
 interval_time IN TEXT DEFAULT 'null',
 id OUT INTEGER
)RETURN INTEGER;
```

#### NOTE

When a scheduled task is created (using **DBE\_TASK**), the system binds the current database and the username to the task by default. This function can be called by using the CALL or SELECT statement. If you call this function by using the SELECT statement, there is no need to specify output parameters. To call this function within a stored procedure, use **perform**. If the committed SQL statement task uses a non-public schema, specify the schema to a table schema or a function schema, or add **set current\_schema = xxx** before the SQL statement.

**Table 10-104** DBE\_TASK.SUBMIT parameters

| Parameter     | Type      | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                     |
|---------------|-----------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| what          | Text      | IN                     | No           | SQL statement to be executed. One or multiple DDL operations (excluding database-related operations), DML operations, anonymous blocks, and statements for calling stored procedures, or all four combined are supported.                                                                                       |
| next_time     | Timestamp | IN                     | No           | Specifies the next time the job will be executed. The default value is the current system time ( <b>sysdate</b> ). If the specified time has past, the job is executed at the time it is committed.                                                                                                             |
| interval_time | Text      | IN                     | Yes          | Calculates the next time to execute the job. It can be an interval expression, or <b>sysdate</b> followed by a numeric value, for example, <b>sysdate+1.0/24</b> . If this parameter is left empty or set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. |
| id            | Integer   | OUT                    | No           | Specifies the job ID. The value ranges from 1 to 32767. When SELECT is used for calling, this parameter cannot be added. When CALL is used for calling, this parameter must be added.                                                                                                                           |

**NOTICE**

When you create a user using the **what** parameter, the plaintext password of the user is recorded in the log. Therefore, you are advised not to do so.

**Example:**

```
openGauss=# SELECT DBE_TASK.SUBMIT('call pro_xxx();', to_date('20180101','yyyymmdd'),'sysdate+1');
submit

31031
(1 row)

openGauss=# SELECT DBE_TASK.SUBMIT('call pro_xxx();', to_date('20180101','yyyymmdd'),'sysdate+1.0/24');
submit

512
(1 row)
```

```
openGauss=# DECLARE
openGauss=# jobid int;
openGauss=# BEGIN
openGauss$$ PERFORM DBE_TASK.SUBMIT('call pro_xxx();', sysdate, 'interval "5 minute"', jobid);
openGauss$# END;
openGauss$# /
ANONYMOUS BLOCK EXECUTE
```

- DBE\_TASK.JOB\_SUBMIT

The stored procedure SUBMIT commits a scheduled task provided by the system. In addition, it provides additional compatibility parameters.

The prototype of the DBE\_TASK.JOB\_SUBMIT function is as follows:

```
DBE_TASK.JOB_SUBMIT(
job OUT INTEGER,
what IN TEXT,
next_date IN TIMESTAMP DEFAULT sysdate,
job_interval IN TEXT DEFAULT 'null',
no_parse IN BOOLEAN DEFAULT false,
instance IN INTEGER DEFAULT 0,
force IN BOOLEAN DEFAULT false
)RETURN INTEGER;
```

**Table 10-105** DBE\_TASK.JOB\_SUBMIT parameters

| Parameter    | Type      | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                     |
|--------------|-----------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job          | Integer   | OUT                    | No           | Specifies the job ID. The value ranges from 1 to 32767. When dbe.job_submit is called by using the SELECT statement, this parameter can be omitted.                                                                                                                                                             |
| what         | Text      | IN                     | No           | SQL statement to be executed. One or multiple DDL operations (excluding database-related operations), DML operations, anonymous blocks, and statements for calling stored procedures, or all four combined are supported.                                                                                       |
| next_date    | Timestamp | IN                     | Yes          | Specifies the next time the job will be executed. The default value is the current system time ( <b>sysdate</b> ). If the specified time has past, the job is executed at the time it is committed.                                                                                                             |
| job_interval | Text      | IN                     | Yes          | Calculates the next time to execute the job. It can be an interval expression, or <b>sysdate</b> followed by a numeric value, for example, <b>sysdate+1.0/24</b> . If this parameter is left empty or set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                      |
|-----------|---------|------------------------|--------------|----------------------------------------------------------------------------------|
| no_parse  | Boolean | IN                     | Yes          | The default value is <b>false</b> , which is used only for syntax compatibility. |
| instance  | Integer | IN                     | Yes          | The default value is <b>0</b> , which is used only for syntax compatibility.     |
| force     | Boolean | IN                     | Yes          | The default value is <b>false</b> , which is used only for syntax compatibility. |

**Example:**

```
openGauss=# DECLARE
openGauss-# id integer;
openGauss-# BEGIN
openGauss$$ id = DBE_TASK.JOB_SUBMIT(
openGauss$$ what => 'insert into t1 values (1, 2)',
openGauss$$ job_interval => 'sysdate + 1' --daily
openGauss$$);
openGauss$$ END;
openGauss$$ /
ANONYMOUS BLOCK EXECUTE
```

- **DBE\_TASK.ID\_SUBMIT**

ID\_SUBMIT has the same syntax function as SUBMIT, but the first parameter of ID\_SUBMIT is an input parameter, that is, a specified job ID. In contrast, that last parameter of ID\_SUBMIT is an output parameter, indicating the job ID automatically generated by the system.

```
DBE_TASK.ID_SUBMIT(
id IN BIGINT,
what IN TEXT,
next_time IN TIMESTAMP DEFAULT sysdate,
interval_time IN TEXT DEFAULT 'null');
```

**Example:**

```
openGauss=# CALL dbe_task.id_submit(101, 'insert_msg_statistic1;', sysdate, 'sysdate+3.0/24');
id_submit

(1 row)
```

- **DBE\_TASK.CANCEL**

The stored procedure CANCEL deletes a specified task.

The prototype of the DBE\_TASK.CANCEL function is as follows:

```
CANCEL(id IN INTEGER);
```

**Table 10-106** DBE\_TASK.CANCEL parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description           |
|-----------|---------|------------------------|--------------|-----------------------|
| id        | Integer | IN                     | No           | Specifies the job ID. |

Example:

```
openGauss=# CALL dbe_task.cancel(101);
cancel

(1 row)
```

- **DBE\_TASK.RUN**

The stored procedure runs a scheduled task.

The prototype of the DBE\_TASK.RUN function is as follows:

```
DBE_TASK.RUN(
job IN BIGINT,
force IN BOOLEAN DEFAULT FALSE);
```

**Table 10-107** DBE\_TASK.RUN parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                         |
|-----------|---------|------------------------|--------------|-------------------------------------|
| job       | bigint  | IN                     | No           | Specifies the job ID.               |
| force     | Boolean | IN                     | Yes          | Used only for syntax compatibility. |

Example:

```
openGauss=# BEGIN
openGauss$$ DBE_TASK.ID_SUBMIT(12345, 'insert_msg_statistic1;', sysdate, 'sysdate+3.0/24');
openGauss$$ DBE_TASK.RUN(12345);
openGauss$$ END;
openGauss$$ /
ANONYMOUS BLOCK EXECUTE
```

- **DBE\_TASK.FINISH**

The stored procedure FINISH disables or enables a scheduled task.

The prototype of the DBE\_TASK.FINISH function is as follows:

```
DBE_TASK.FINISH(
id IN INTEGER,
broken IN BOOLEAN,
next_time IN TIMESTAMP DEFAULT sysdate);
```

**Table 10-108** DBE\_TASK.FINISH parameters

| Parameter | Type      | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|-----------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id        | Integer   | IN                     | No           | Specifies the job ID.                                                                                                                                                                                                                                                                                                                                                                                                          |
| broken    | Boolean   | IN                     | No           | Specifies the status flag, <b>true</b> for broken and <b>false</b> for not broken. The current job is updated based on the parameter value <b>true</b> or <b>false</b> . If the parameter is left empty, the job status remains unchanged.                                                                                                                                                                                     |
| next_time | Timestamp | IN                     | Yes          | Specifies the next execution time. The default value is the current system time. If <b>broken</b> is set to <b>true</b> , <b>next_time</b> is updated to '4000-1-1'. If <b>broken</b> is set to <b>false</b> and <b>next_time</b> is not empty, <b>next_time</b> is updated for the job. If <b>next_time</b> is empty, it will not be updated. This parameter can be omitted, and its default value will be used in this case. |

**Example:**

```
openGauss=# CALL dbe_task.id_submit(101, 'insert_msg_statistic1;', sysdate, 'sysdate+3.0/24');
id_submit
```

```

(1 row)
```

```
openGauss=# CALL dbe_task.finish(101, true);
finish
```

```

(1 row)
```

- **DBE\_TASK.UPDATE**

The stored procedure UPDATE modifies user-definable attributes of a task, including the task content, next-execution time, and execution interval.

The prototype of the DBE\_TASK.UPDATE function is as follows:

```
DBE_TASK.UPDATE(
id IN INTEGER,
content IN TEXT,
next_time IN TIMESTAMP,
interval_time IN TEXT);
```

**Table 10-109** DBE\_TASK.UPDATE parameters

| Parameter     | Type      | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|-----------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id            | Integer   | IN                     | No           | Specifies the job ID.                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| content       | Text      | IN                     | Yes          | Specifies the name of the stored procedure or SQL statement block that is executed. If this parameter is left empty, the system does not update the <b>content</b> parameter for the specified job. Otherwise, the system updates the <b>content</b> parameter for the specified job.                                                                                                                                                                            |
| next_time     | Timestamp | IN                     | Yes          | Specifies the next execution time. If this parameter is left empty, the system does not update the <b>next_time</b> parameter for the specified job. Otherwise, the system updates the <b>next_time</b> parameter for the specified job.                                                                                                                                                                                                                         |
| interval_time | Text      | IN                     | Yes          | Specifies the time expression for calculating the next time the job will be executed. If this parameter is left empty, the system does not update the <b>interval_time</b> parameter for the specified job. Otherwise, the system updates the <b>interval_time</b> parameter for the specified job after necessary validity check. If this parameter is set to <b>null</b> , the job will be executed only once, and the job state will change to 'd' afterward. |

**Example:**

```
openGauss=# CALL dbe_task.update(101, 'call userproc();', sysdate, 'sysdate + 1.0/1440');
update

```

(1 row)

```
openGauss=# CALL dbe_task.update(101, 'insert into tbl_a values(sysdate);', sysdate, 'sysdate + 1.0/1440');
update

```

(1 row)

- DBE\_TASK.CHANGE

The stored procedure UPDATE modifies user-definable attributes of a task, including the task content, next-execution time, and execution interval.

The prototype of the DBE\_TASK.CHANGE function is as follows:

```
DBE_TASK.CHANGE(
job IN INTEGER,
what IN TEXT DEFAULT NULL,
next_date IN TIMESTAMP DEFAULT NULL,
job_interval IN TEXT DEFAULT NULL,
instance IN INTEGER DEFAULT NULL,
force IN BOOLEAN DEFAULT false);
```

**Table 10-110** DBE\_TASK.CHANGE parameters

| Parameter    | Type      | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------|-----------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job          | Integer   | IN                     | No           | Specifies the job ID.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| what         | Text      | IN                     | Yes          | Specifies the name of the stored procedure or SQL statement block that is executed. If this parameter is left empty, the system does not update the <b>what</b> parameter for the specified job. Otherwise, the system updates the <b>what</b> parameter for the specified job.                                                                                                                                                                                |
| next_date    | Timestamp | IN                     | Yes          | Specifies the next execution time. If this parameter is left empty, the system does not update the <b>next_date</b> parameter for the specified job. Otherwise, the system updates the <b>next_date</b> parameter for the specified job.                                                                                                                                                                                                                       |
| job_interval | Text      | IN                     | Yes          | Specifies the time expression for calculating the next time the job will be executed. If this parameter is left empty, the system does not update the <b>job_interval</b> parameter for the specified job. Otherwise, the system updates the <b>job_interval</b> parameter for the specified job after necessary validity check. If this parameter is set to <b>null</b> , the job will be executed only once, and the job state will change to 'd' afterward. |
| instance     | Integer   | IN                     | Yes          | Used only for syntax compatibility.                                                                                                                                                                                                                                                                                                                                                                                                                            |



| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                         |
|-----------|---------|------------------------|--------------|-------------------------------------|
| force     | Boolean | IN                     | No           | Used only for syntax compatibility. |

Example:

```
openGauss=# BEGIN
openGauss$$ DBE_TASK.CHANGE(
openGauss$$ job => 101,
openGauss$$ what => 'insert into t2 values (2);'
openGauss$$);
openGauss$$ END;
openGauss$$ /
ANONYMOUS BLOCK EXECUTE
```

- DBE\_TASK.CONTENT

The stored procedure CONTENT modifies the procedures to be executed by a specified task.

The prototype of the DBE\_TASK.CONTENT function is as follows:

```
DBE_TASK.CONTENT(
id IN INTEGER,
content IN TEXT);
```

**Table 10-111** DBE\_TASK.CONTENT parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                  |
|-----------|---------|------------------------|--------------|----------------------------------------------------------------------------------------------|
| id        | Integer | IN                     | No           | Specifies the job ID.                                                                        |
| content   | Text    | IN                     | No           | Specifies the stored procedure call, SQL statement block, or program block that is executed. |

 **NOTE**

- If the value specified by the **content** parameter is one or multiple executable SQL statements, program blocks, or stored procedures, this procedure can be executed successfully; otherwise, it will fail to be executed.
- If the value specified by the **content** parameter is a simple statement such as **INSERT** and **UPDATE**, a schema name must be added in front of the table name.

Example:

```
openGauss=# CALL db_e_task.content(101, 'call userproc();');
content

```

```
(1 row)

openGauss=# CALL db_task.content(101, 'insert into tbl_a values(sysdate);');
content

(1 row)
```

- DBE\_TASK.NEXT\_TIME

The stored procedure NEXT\_TIME modifies the next-execution time attribute of a task.

The prototype of the DBE\_TASK.NEXT\_TIME function is as follows:

```
DBE_TASK.NEXT_TIME(
id IN BIGINT,
next_time IN TEXT);
```

**Table 10-112** DBE\_TASK.NEXT\_TIME parameters

| Parameter | Type   | Input/Output Parameter | Can Be Empty | Description                        |
|-----------|--------|------------------------|--------------|------------------------------------|
| id        | bigint | IN                     | No           | Specifies the job ID.              |
| next_time | Text   | IN                     | No           | Specifies the next execution time. |

 **NOTE**

If the specified **next\_time** value is earlier than the current date, the job is executed once immediately.

Example:

```
openGauss=# CALL db_task.next_time(101, sysdate);
next_time

(1 row)
```

- DBE\_TASK.INTERVAL

The stored procedure INTERVAL modifies the execution interval attribute of a task.

The prototype of the DBE\_TASK.INTERVAL function is as follows:

```
DBE_TASK.INTERVAL(
id IN INTEGER,
interval_time IN TEXT);
```

**Table 10-113** DBE\_TASK.INTERVAL parameters

| Parameter     | Type    | Input / Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                        |
|---------------|---------|--------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| id            | Integer | IN                       | No           | Specifies the job ID.                                                                                                                                                                                                                                                                              |
| interval_time | Text    | IN                       | Yes          | Specifies the time expression for calculating the next time the job will be executed. If this parameter is left empty or set to <b>null</b> , the job will be executed only once, and the job status will change to 'd' afterward. The job interval must be of a valid time type or interval type. |

Example:

```
openGauss=# CALL dbe_task.interval(101, 'sysdate + 1.0/1440');
interval

(1 row)

openGauss=# CALL dbe_task.cancel(101);
cancel

(1 row)
```

#### NOTE

For a job that is currently running (that is, **job\_status** is 'r'), it is not allowed to use **cancel**, **update**, **next\_time**, **content**, or **interval** to delete or modify job parameters.

## Constraints

1. You can create, update, and delete jobs only using the procedures provided by the **DBE\_TASK** package. These procedures synchronize job information between different primary database nodes and associate primary keys between the pg\_job and pg\_job\_proc catalogs. If you use DML statements to add, delete, or modify records in the pg\_job catalog, job information will become inconsistent between primary database nodes and system catalogs may fail to be associated, compromising internal job management.
2. Each task created by a user is bound to a primary database node. If the primary database node fails while a task is being executed, the task status cannot be updated in real time and will stay at 'r'. The task status will be updated to 's' only after the primary database node recovers. When a primary database node fails, all tasks on this primary database node cannot be scheduled or executed until the primary database node is restored manually, or deleted and then replaced.

- For each job, the bound primary database node updates the real-time job information (including the job status, last execution start time, last execution end time, next execution start time, the number of execution failures [if any]) to the pg\_job system catalog, and synchronizes the information to other primary database nodes, ensuring consistent job information between different primary database nodes. In the case of primary database node failures, job information synchronization is reattempted by the hosting primary database nodes, which increases job execution time. Although job information fails to be synchronized between primary database nodes, job information can still be properly updated in the pg\_job table on the hosting primary database nodes, and jobs can be executed successfully. After a primary database node recovers, job information such as job execution time and status in its pg\_job table may be incorrect and will be updated only after the jobs are executed again on related primary database nodes.
- For each job, a thread is established to execute it. If multiple jobs are triggered concurrently as scheduled, the system will need some time to start the required threads, resulting in a latency of 0.1 ms in job execution.

### 10.12.2.6 DBE\_SCHEDULER

#### API Description

The advanced package **DBE\_SCHEDULER** supports more flexible creation of scheduled jobs through scheduling and programing. For details about all the supported APIs, see [Table 10-114](#).

**Table 10-114** DBE\_SCHEDULER

| API                                              | Description                            |
|--------------------------------------------------|----------------------------------------|
| <a href="#">DBE_SCHEDULER.CREATE_JOB</a>         | Creates a scheduled job.               |
| <a href="#">DBE_SCHEDULER.DROP_JOB</a>           | Deletes a scheduled job.               |
| <a href="#">DBE_SCHEDULER.DROP_SINGLE_JOB</a>    | Deletes a single scheduled job.        |
| <a href="#">DBE_SCHEDULER.SET_ATTRIBUTE</a>      | Sets object attributes.                |
| <a href="#">DBE_SCHEDULER.RUN_JOB</a>            | Executes a scheduled job.              |
| <a href="#">DBE_SCHEDULER.RUN_BACKGROUND_JOB</a> | Runs a scheduled job in the backend.   |
| <a href="#">DBE_SCHEDULER.RUN_FOREGROUND_JOB</a> | Runs a scheduled job in the frontend.  |
| <a href="#">DBE_SCHEDULER.STOP_JOB</a>           | Stops a scheduled job.                 |
| <a href="#">DBE_SCHEDULER.STOP_SINGLE_JOB</a>    | Stops a single scheduled job.          |
| <a href="#">DBE_SCHEDULER.GENERATE_JOB_NAME</a>  | Generates the name of a scheduled job. |

| API                                                     | Description                                  |
|---------------------------------------------------------|----------------------------------------------|
| <a href="#">DBE_SCHEDULER.CREATE_PROGRAM</a>            | Creates a program.                           |
| <a href="#">DBE_SCHEDULER.DEFINE_PROGRAM_ARGUMENT</a>   | Defines program parameters.                  |
| <a href="#">DBE_SCHEDULER.DROP_PROGRAM</a>              | Deletes a program.                           |
| <a href="#">DBE_SCHEDULER.DROP_SINGLE_PROGRAM</a>       | Deletes a single program.                    |
| <a href="#">DBE_SCHEDULER.SET_JOB_ARGUMENT_VALUE</a>    | Sets the parameters of a scheduled job.      |
| <a href="#">DBE_SCHEDULER.CREATE_SCHEDULE</a>           | Creates a schedule.                          |
| <a href="#">DBE_SCHEDULER.DROP_SCHEDULE</a>             | Deletes a schedule.                          |
| <a href="#">DBE_SCHEDULER.DROP_SINGLE_SCHEDULE</a>      | Deletes a single schedule.                   |
| <a href="#">DBE_SCHEDULER.CREATE_JOB_CLASS</a>          | Creates the class of a scheduled job.        |
| <a href="#">DBE_SCHEDULER.DROP_JOB_CLASS</a>            | Deletes the class of a scheduled job.        |
| <a href="#">DBE_SCHEDULER.DROP_SINGLE_JOB_CLASS</a>     | Deletes the class of a single scheduled job. |
| <a href="#">DBE_SCHEDULER.GRANT_USER_AUTHORIZATION</a>  | Grants special permissions to a user.        |
| <a href="#">DBE_SCHEDULER.REVOKE_USER_AUTHORIZATION</a> | Revokes special permissions from a user.     |
| <a href="#">DBE_SCHEDULER.CREATE_CREDENTIAL</a>         | Creates a certificate.                       |
| <a href="#">DBE_SCHEDULER.DROP_CREDENTIAL</a>           | Destroys a certificate.                      |
| <a href="#">DBE_SCHEDULER.ENABLE</a>                    | Enables an object.                           |
| <a href="#">DBE_SCHEDULER.ENABLE_SINGLE</a>             | Enables a single object.                     |
| <a href="#">DBE_SCHEDULER.DISABLE</a>                   | Disables an object.                          |
| <a href="#">DBE_SCHEDULER.DISABLE_SINGLE</a>            | Disables a single object.                    |

| API                                                    | Description                         |
|--------------------------------------------------------|-------------------------------------|
| <a href="#">DBE_SCHEDULER.EVAL_CALENDAR_STRING</a>     | Analyzes the scheduling job period. |
| <a href="#">DBE_SCHEDULER.EVALUATE_CALENDAR_STRING</a> | Analyzes the scheduling job period. |

- **DBE\_SCHEDULER.CREATE\_JOB**

Creates a scheduled job.

The prototypes of the DBE\_SCHEDULER.CREATE\_JOB function are as follows:

```
-- Scheduled jobs of an inline schedule and a program.
DBE_SCHEDULER.CREATE_JOB(
job_name TEXT,
job_type TEXT,
job_action TEXT,
number_of_arguments INTEGER DEFAULT 0,
start_date TIMESTAMP WITH TIME ZONE DEFAULT NULL,
repeat_interval TEXT DEFAULT NULL,
end_date TIMESTAMP WITH TIME ZONE DEFAULT NULL,
job_class TEXT DEFAULT 'DEFAULT_JOB_CLASS',
enabled BOOLEAN DEFAULT FALSE,
auto_drop BOOLEAN DEFAULT TRUE,
comments TEXT DEFAULT NULL,
credential_name TEXT DEFAULT NULL,
destination_name TEXT DEFAULT NULL
)

-- Reference the created scheduled jobs of the schedule and the program.
DBE_SCHEDULER.CREATE_JOB(
job_name TEXT,
program_name TEXT,
schedule_name TEXT,
job_class TEXT DEFAULT 'DEFAULT_JOB_CLASS',
enabled BOOLEAN DEFAULT FALSE,
auto_drop BOOLEAN DEFAULT TRUE,
comments TEXT DEFAULT NULL,
job_style TEXT DEFAULT 'REGULAR',
credential_name TEXT DEFAULT NULL,
destination_name TEXT DEFAULT NULL
)

-- Reference the created program and the scheduled job of the inline schedule.
DBE_SCHEDULER.CREATE_JOB(
job_name text,
program_name TEXT,
start_date TIMESTAMP WITH TIME ZONE DEFAULT NULL,
repeat_interval TEXT DEFAULT NULL,
end_date TIMESTAMP WITH TIME ZONE DEFAULT NULL,
job_class TEXT DEFAULT 'DEFAULT_JOB_CLASS',
enabled BOOLEAN DEFAULT FALSE,
auto_drop BOOLEAN DEFAULT TRUE,
comments TEXT DEFAULT NULL,
job_style TEXT DEFAULT 'REGULAR',
credential_name TEXT DEFAULT NULL,
destination_name TEXT DEFAULT NULL
)

-- Reference the created schedule and the scheduled job of the inline program.
DBE_SCHEDULER.CREATE_JOB(
job_name TEXT,
schedule_name TEXT,
job_type TEXT,
```

```
job_action TEXT,
number_of_arguments INTEGER DEFAULT 0,
job_class TEXT DEFAULT 'DEFAULT_JOB_CLASS',
enabled BOOLEAN DEFAULT FALSE,
auto_drop BOOLEAN DEFAULT TRUE,
comments TEXT DEFAULT NULL,
credential_name TEXT DEFAULT NULL,
destination_name TEXT DEFAULT NULL
)
```

 **NOTE**

The scheduled job created through **DBE\_SCHEDULER** does not conflict with the scheduled job in **DBE\_TASK**.

The scheduled job created by **DBE\_SCHEDULER** generates the corresponding **job\_id**. However, **job\_id** is meaningless.

**Table 10-115** DBE\_SCHEDULER.CREATE\_JOB API parameters

| Parameter           | Type                     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                            |
|---------------------|--------------------------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name            | text                     | IN                     | No           | Name of a scheduled job.                                                                                                                                                                                                                                                                               |
| job_type            | text                     | IN                     | No           | Inline program type of a scheduled job. The options are as follows: <ul style="list-style-type: none"><li>• <b>'PLSQL_BLOCK'</b>: fast anonymous stored procedure.</li><li>• <b>'STORED_PROCEDURE'</b>: stored procedure that is saved.</li><li>• <b>'EXTERNAL_SCRIPT'</b>: external script.</li></ul> |
| job_action          | text                     | IN                     | No           | Content executed by an inline program of a scheduled job.                                                                                                                                                                                                                                              |
| number_of_arguments | integer                  | IN                     | No           | Number of inline program parameters of a scheduled job.                                                                                                                                                                                                                                                |
| program_name        | text                     | IN                     | No           | Name of the program referenced by a scheduled job.                                                                                                                                                                                                                                                     |
| start_date          | timestamp with time zone | IN                     | Yes          | Inline scheduling start time of a scheduled job.                                                                                                                                                                                                                                                       |

| Parameter        | Type                     | Input/Output Parameter | Can Be Empty | Description                                                              |
|------------------|--------------------------|------------------------|--------------|--------------------------------------------------------------------------|
| repeat_interval  | text                     | IN                     | Yes          | Inline scheduling period of a scheduled job.                             |
| end_date         | timestamp with time zone | IN                     | Yes          | Inline scheduling expiration time of a scheduled job.                    |
| schedule_name    | text                     | IN                     | No           | Name of the schedule referenced by a scheduled job.                      |
| job_class        | text                     | IN                     | No           | Class name of a scheduled job.                                           |
| enabled          | boolean                  | IN                     | No           | Status of a scheduled job.                                               |
| auto_drop        | boolean                  | IN                     | No           | Automatic deletion of a scheduled job.                                   |
| comments         | text                     | IN                     | Yes          | Comments.                                                                |
| job_style        | text                     | IN                     | No           | Behavior pattern of a scheduled job. Only <b>'REGULAR'</b> is supported. |
| credential_name  | text                     | IN                     | Yes          | Certificate name of a scheduled job.                                     |
| destination_name | text                     | IN                     | Yes          | Target name of a scheduled job.                                          |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 3, false, 'test');
create_program
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule1', NULL, 'sysdate', NULL, 'test');
create_schedule
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.create_job(job_name=>'job1', program_name=>'program1',
schedule_name=>'schedule1');
create_job
```



```

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule1');
drop_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

(1 row)
```

#### NOTICE

To create a scheduled job of the EXTERNAL\_SCRIPT type, the administrator must assign related permissions and certificates and the user who starts the database must have the read permission on the external script.

- **DBE\_SCHEDULER.DROP\_JOB**

Deletes a scheduled job.

The prototype of the DBE\_SCHEDULER.DROP\_JOB function is as follows:

```
DBE_SCHEDULER.drop_job(
job_name text,
force boolean default false,
defer boolean default false,
commit_semantics text default 'STOP_ON_FIRST_ERROR'
)
```

#### NOTE

DBE\_SCHEDULER.DROP\_JOB can specify one or more jobs, or specify a job class to delete a scheduled job.

**Table 10-116** DBE\_SCHEDULER.DROP\_JOB API parameters

| Parameter        | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name         | text    | IN                     | No           | Name or class of a scheduled job. You can specify one or more scheduled jobs. If you specify multiple scheduled jobs, separate them with commas (.).                                                                                                                                                                                                                                                                                                                     |
| force            | boolean | IN                     | No           | Specifies whether to delete a scheduled job. <ul style="list-style-type: none"><li>• <b>true</b>: The current scheduled job is stopped and then deleted.</li><li>• <b>false</b>: The scheduled job fails to be deleted if it is running.</li></ul>                                                                                                                                                                                                                       |
| defer            | boolean | IN                     | No           | Specifies whether to delete a scheduled job. <ul style="list-style-type: none"><li>• <b>true</b>: A scheduled job can be deleted after it is complete.</li><li>• <b>false</b>: The scheduled job cannot be executed and is deleted.</li></ul>                                                                                                                                                                                                                            |
| commit_semantics | text    | IN                     | No           | Commit rules: <ul style="list-style-type: none"><li>• <b>'STOP_ON_FIRST_ERROR'</b>: The deletion operation performed before the first error is reported is committed.</li><li>• <b>'TRANSACTIONAL'</b>: Transaction-level commit. The deletion operation performed before an error is reported will be rolled back.</li><li>• <b>'ABSORB_ERRORS'</b>: The system attempts to bypass an error and commit the deletion operation that is performed successfully.</li></ul> |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 3, false, 'test');
create_program
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule1', NULL, 'sysdate', NULL, 'test');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_job(job_name=>'job1', program_name=>'program1',
schedule_name=>'schedule1');
create_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule1');
drop_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

(1 row)
```

**NOTICE**

The **TRANSACTIONAL** option in **commit\_semantic** takes effect only when **force** is set to **false**.

- DBE\_SCHEDULER.DROP\_SINGLE\_JOB

Deletes a scheduled job.

The prototype of the DBE\_SCHEDULER.DROP\_SINGLE\_JOB function is as follows:

```
DBE_SCHEDULER.drop_single_job(
job_name text,
force boolean default false,
defer boolean default false
)
```

**Table 10-117** DBE\_SCHEDULER.DROP\_SINGLE\_JOB API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                                   |
|-----------|------|------------------------|--------------|---------------------------------------------------------------|
| job_name  | text | IN                     | No           | Specifies the name of a scheduled job or scheduled job class. |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                          |
|-----------|---------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| force     | boolean | IN                     | No           | Specifies whether to delete a scheduled job.<br><ul style="list-style-type: none"> <li><b>true:</b> The current scheduled job is stopped and then deleted.</li> <li><b>false:</b> The scheduled job fails to be deleted if it is running.</li> </ul> |
| defer     | boolean | IN                     | No           | Specifies whether to delete a scheduled job.<br><ul style="list-style-type: none"> <li><b>true:</b> A scheduled job can be deleted after it is complete.</li> <li><b>false:</b> The scheduled job cannot be executed and is deleted.</li> </ul>      |

#### Example:

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 0, false, 'test');
create_program

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_job('job1', 'program1', '2021-07-20', 'interval "3 minute"',
'2121-07-20', 'DEFAULT_JOB_CLASS', false, false, 'test', 'style', NULL, NULL);
create_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_single_job('job1', false, false);
drop_single_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

(1 row)
```

- DBE\_SCHEDULER.SET\_ATTRIBUTE

Modifies the attributes of a scheduled job.

The prototypes of the DBE\_SCHEDULER.SET\_ATTRIBUTE function are as follows:

```
DBE_SCHEDULER.set_attribute(
name text,
attribute text,
value boolean
)
```

```
DBE_SCHEDULER.set_attribute(
name text,
attribute text,
value text
)

DBE_SCHEDULER.set_attribute(
name text,
attribute text,
value timestamp
)

DBE_SCHEDULER.set_attribute(
name text,
attribute text,
value timestamp with time zone
)

DBE_SCHEDULER.set_attribute(
name text,
attribute text,
value text,
value2 text default NULL
)
```

 **NOTE**

**name** specifies any object in **DBE\_SCHEDULER**.

**Table 10-118** DBE\_SCHEDULER.SET\_ATTRIBUTE API parameters

| Parameter | Type                                                 | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------|------------------------------------------------------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | text                                                 | IN                     | No           | Object name.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| attribute | text                                                 | IN                     | No           | Attribute name.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| value     | boolean/date/timestamp/timestamp with time zone/text | IN                     | No           | Attribute value. The options are as follows:<br><br>Scheduled-job-related: job_type, job_action, number_of_arguments, start_date, repeat_interval, end_date, job_class, enabled, auto_drop, comments, credential_name, destination_name, program_name, schedule_name, and job_style.<br><br>Scheduling-related: program_action, program_type, number_of_arguments, comments.<br><br>Program-related: start_date, repeat_interval, end_date, comments. |

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                       |
|-----------|------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------|
| value2    | text | IN                     | Yes          | Additional attribute value. Reserved parameter bit. Currently, the target attribute with extra attribute values is not supported. |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 3, false, 'test');
create_program

```

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.set_attribute('program1', 'number_of_arguments', 0);
set_attribute

```

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.set_attribute('program1', 'program_type',
'STORED_PROCEDURE');
set_attribute

```

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

```

(1 row)

**NOTICE**

Do not use DBE\_SCHEDULER.SET\_ATTRIBUTE to leave the parameters empty. The object name cannot be changed using DBE\_SCHEDULER.SET\_ATTRIBUTE.

- **DBE\_SCHEDULER.RUN\_JOB**

Executes a scheduled job.

The prototype of the DBE\_SCHEDULER.RUN\_JOB function is as follows:

```
DBE_SCHEDULER.run_job(
job_name text,
use_current_session boolean default true
)
```

 NOTE

DBE\_SCHEDULER.RUN\_JOB is used to run scheduled jobs immediately. It is independent of the scheduling of scheduled jobs and can even run at the same time.

**Table 10-119** DBE\_SCHEDULER.RUN\_JOB API parameters

| Parameter           | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                     |
|---------------------|---------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name            | text    | IN                     | No           | One or more scheduled job names. If you specify multiple scheduled jobs, separate them with commas (,).                                                                                                                                                                                         |
| use_current_session | boolean | IN                     | No           | Specifies whether to run a scheduled job. <ul style="list-style-type: none"><li>• <b>true</b>: Use the current session to check whether the scheduled job can run properly.</li><li>• <b>false</b>: Start the scheduled job in the backend. The execution result is recorded in logs.</li></ul> |

## Example:

```
openGauss=# SELECT dbe_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1
values(12); end;','0,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.run_job('job1', false);
run_job
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job
```

```

(1 row)
```

**NOTICE**

**use\_current\_session** cannot be used to record execution results.

- DBE\_SCHEDULER.RUN\_BACKEND\_JOB

Runs a scheduled job in the backend.

The prototype of the DBE\_SCHEDULER.RUN\_BACKEND\_JOB function is as follows:

```
DBE_SCHEDULER.run_backend_job(
job_name text
)
```

**Table 10-120** DBE\_SCHEDULER.RUN\_BACKEND\_JOB API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description              |
|-----------|------|------------------------|--------------|--------------------------|
| job_name  | text | IN                     | No           | Name of a scheduled job. |

Example:

```
openGauss=# SELECT db_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1
values(12); end;','0,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.run_backend_job('job1');
run_backend_job
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job
```

-----

(1 row)

- DBE\_SCHEDULER.RUN\_FOREGROUND\_JOB

Executes a scheduled job in the current session.

Only external jobs can be executed.

Return value: text

The prototype of the DBE\_SCHEDULER.RUN\_FOREGROUND\_JOB function is as follows:

```
DBE_SCHEDULER.run_foreground_job(
job_name text
)return text
```

**Table 10-121** DBE\_SCHEDULER.RUN\_FOREGROUND\_JOB API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description              |
|-----------|------|------------------------|--------------|--------------------------|
| job_name  | text | IN                     | No           | Name of a scheduled job. |



**Example:**

```
openGauss=# create user test1 identified by '*****';
NOTICE: The encrypted password contains MD5 ciphertext, which is not secure.
CREATE ROLE
openGauss=# select DBE_SCHEDULER.create_credential('cre_1', 'test1', '*****');
create_credential

(1 row)

openGauss=# select DBE_SCHEDULER.create_job(job_name=>'job1', job_type=>'EXTERNAL_SCRIPT',
job_action=>'/usr/bin/pwd', enabled=>true, auto_drop=>false, credential_name => 'cre_1');
create_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.run_foreground_job('job1');
run_foreground_job

Host key verification failed.\r+
(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_credential('cre_1', false);
drop_credential

(1 row)

openGauss=# drop user test1;
DROP ROLE
```

- **DBE\_SCHEDULER.STOP\_JOB**

Stops a scheduled job.

The prototype of the DBE\_SCHEDULER.STOP\_JOB function is as follows:

```
DBE_SCHEDULER.stop_job(
job_name text,
force boolean default false,
commit_semantics text default 'STOP_ON_FIRST_ERROR'
)
```

**Table 10-122** DBE\_SCHEDULER.STOP\_JOB API parameters

| Parameter        | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                              |
|------------------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name         | text    | IN                     | No           | Name or class of a scheduled job. You can specify one or more scheduled jobs. If you specify multiple scheduled jobs, separate them with commas (,).                                                                                                                     |
| force            | boolean | IN                     | No           | Specifies whether to delete a scheduled job.<br><b>true</b> : The scheduler sends a termination signal to end the job thread immediately.<br><b>false</b> : The scheduler attempts to use the interrupt signal to terminate the scheduled job thread.                    |
| commit_semantics | text    | IN                     | No           | Commit rules:<br><b>'STOP_ON_FIRST_ERROR'</b> : The interrupt operation performed before the first error is reported is committed.<br><b>'ABSORB_ERRORS'</b> : The system attempts to bypass an error and commit the interrupt operation that is performed successfully. |

**Example:**

```
openGauss=# SELECT dbe_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1
values(12); end;','0,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.stop_job('job1', true, 'STOP_ON_FIRST_ERROR');
stop_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)
```

- DBE\_SCHEDULER.STOP\_SINGLE\_JOB

Stops a single scheduled job.

The prototype of the DBE\_SCHEDULER.STOP\_SINGLE\_JOB function is as follows:

```
DBE_SCHEDULER.stop_single_job(
 job_name text,
 force boolean default false
)
```

**Table 10-123** DBE\_SCHEDULER.STOP\_SINGLE\_JOB API parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                  |
|-----------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_name  | text    | IN                     | No           | Specifies the name of a scheduled job or scheduled job class.                                                                                                                                                                                                                                                |
| force     | boolean | IN                     | No           | <p>Specifies whether to delete a scheduled job.</p> <ul style="list-style-type: none"><li><b>true:</b> The scheduler sends a termination signal to end the job thread immediately.</li><li><b>false:</b> The scheduler attempts to use the interrupt signal to terminate the scheduled job thread.</li></ul> |

Example:

```
openGauss=# SELECT db_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1
values(12); end;','0,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.stop_single_job('job1', true);
stop_single_job

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)
```

- DBE\_SCHEDULER.GENERATE\_JOB\_NAME

Generates the name of a scheduled job.

The prototype of the DBE\_SCHEDULER.GENERATE\_JOB\_NAME function is as follows:

```
DBE_SCHEDULER.generate_job_name(
 prefix text default 'JOB$_
)return text
```

**Table 10-124** DBE\_SCHEDULER.GENERATE\_JOB\_NAME API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                          |
|-----------|------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| prefix    | text | IN                     | No           | Prefix of the generated name. The default value is 'JOB\$_. Scheduled jobs that are repeatedly executed are named as follows: <b>job\$_1</b> , <b>job\$_2</b> , <b>job\$_3</b> , ... |

Example:

```
openGauss=# CALL DBE_SCHEDULER.generate_job_name();
generate_job_name

JOB$_1
(1 row)

openGauss=# CALL DBE_SCHEDULER.generate_job_name();
generate_job_name

JOB$_2
(1 row)

openGauss=# CALL DBE_SCHEDULER.generate_job_name('job');
generate_job_name

job3
(1 row)

openGauss=# CALL DBE_SCHEDULER.generate_job_name('job');
generate_job_name

job4
(1 row)
```

**NOTICE**

When DBE\_SCHEDULER.GENERATE\_JOB\_NAME is executed for the first time, a temporary sequence is created in **public** to store the sequence number of the current name. A common user does not have the create permission in **public**. Therefore, if a common user calls the function for the first time in the current database, the function fails to be called. In this case, you need to grant the create permission in **public** to the common user or call the API as a user with the create permission to create a temporary sequence.

- DBE\_SCHEDULER.CREATE\_PROGRAM

Creates a program.

The prototype of the DBE\_SCHEDULER.CREATE\_PROGRAM function is as follows:

```
DBE_SCHEDULER.create_program(
program_name text,
program_type text,
program_action text,
number_of_arguments integer default 0,
enabled boolean default false,
comments text default NULL
)
```

**Table 10-125** DBE\_SCHEDULER.CREATE\_PROGRAM API parameters

| Parameter           | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                       |
|---------------------|---------|------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| program_name        | text    | IN                     | No           | Program name.                                                                                                                                                                                                                                     |
| program_type        | text    | IN                     | No           | Program type. The options are as follows: <ul style="list-style-type: none"><li>'PLSQL_BLOCK': fast anonymous stored procedure.</li><li>'STORED_PROCEDURE': stored procedure that is saved.</li><li>'EXTERNAL_SCRIPT': external script.</li></ul> |
| program_action      | text    | IN                     | No           | Program operation.                                                                                                                                                                                                                                |
| number_of_arguments | integer | IN                     | No           | Number of parameters used by the program.                                                                                                                                                                                                         |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                               |
|-----------|---------|------------------------|--------------|-------------------------------------------|
| enabled   | boolean | IN                     | No           | Specifies whether the program is enabled. |
| comments  | text    | IN                     | Yes          | Comments.                                 |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 3, false, 'test');
create_program
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program
```

```

(1 row)
```

- DBE\_SCHEDULER.DEFINE\_PROGRAM\_ARGUMENT

Defines program parameters.

An API with the default value **default\_value** will not convert characters to lowercase letters by default. In this version, characters are case-sensitive.

The prototype of the DBE\_SCHEDULER.DEFINE\_PROGRAM\_ARGUMENT function is as follows:

```
DBE_SCHEDULER.define_program_argument(
program_name text,
argument_position integer,
argument_name text default NULL,
argument_type text,
out_argument boolean default false
)
```

-- With a default value --

```
DBE_SCHEDULER.define_program_argument(
program_name text,
argument_position integer,
argument_name text default NULL,
argument_type text,
default_value text,
out_argument boolean default false
)
```

**Table 10-126** DBE\_SCHEDULER.DEFINE\_PROGRAM\_ARGUMENT API parameters

| Parameter         | Type    | Input/Output Parameter | Can Be Empty | Description         |
|-------------------|---------|------------------------|--------------|---------------------|
| program_name      | text    | IN                     | No           | Program name.       |
| argument_position | integer | IN                     | No           | Parameter location. |
| argument_name     | text    | IN                     | No           | Parameter name.     |
| argument_type     | text    | IN                     | No           | Parameter type.     |
| default_value     | text    | IN                     | No           | Default value.      |
| out_argument      | boolean | IN                     | No           | Reserved parameter. |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 2, false, 'test');
create_program

(1 row)

openGauss=# CALL DBE_SCHEDULER.define_program_argument('program1', 1, 'pa1', 'type1', false);
define_program_argument

(1 row)

openGauss=# CALL DBE_SCHEDULER.define_program_argument('program1', 1, 'pa1', 'type1', 'value1',
false);
define_program_argument

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

(1 row)
```

- **DBE\_SCHEDULER.DROP\_PROGRAM**  
Deletes a program.

The prototype of the DBE\_SCHEDULER.DROP\_PROGRAM function is as follows:

```
DBE_SCHEDULER.drop_program(
program_name text,
force boolean default false
)
```

**Table 10-127** DBE\_SCHEDULER.DROP\_PROGRAM API parameters

| Parameter    | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                      |
|--------------|---------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| program_name | text    | IN                     | No           | Program name.                                                                                                                                                                                                                                                                    |
| force        | boolean | IN                     | No           | Specifies whether to delete a program. <ul style="list-style-type: none"><li>• <b>true</b>: Before a program is deleted, all jobs that use the program are disabled.</li><li>• <b>false</b>: The program cannot be referenced by any job. Otherwise, an error is sent.</li></ul> |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 2, false, 'test');
create_program
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program
```

-----

(1 row)

- **DBE\_SCHEDULER.DROP\_SINGLE\_PROGRAM**

Deletes a single program.

The prototype of the DBE\_SCHEDULER.DROP\_SINGLE\_PROGRAM function is as follows:

```
DBE_SCHEDULER.drop_single_program(
program_name text,
force boolean default false
)
```



**Table 10-128** DBE\_SCHEDULER.DROP\_SINGLE\_PROGRAM API parameters

| Parameter    | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                      |
|--------------|---------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| program_name | text    | IN                     | No           | Program name.                                                                                                                                                                                                                                                                    |
| force        | boolean | IN                     | No           | Specifies whether to delete a program. <ul style="list-style-type: none"><li>• <b>true</b>: Before a program is deleted, all jobs that use the program are disabled.</li><li>• <b>false</b>: The program cannot be referenced by any job. Otherwise, an error is sent.</li></ul> |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'STORED_PROCEDURE', 'select
pg_sleep(1);', 2, false, 'test');
create_program
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_single_program('program1', false);
drop_single_program
```

```

(1 row)
```

- **DBE\_SCHEDULER.SET\_JOB\_ARGUMENT\_VALUE**

Sets the parameters of a scheduled job. **argument\_value** can be left empty.

The prototype of the DBE\_SCHEDULER.SET\_JOB\_ARGUMENT\_VALUE function is as follows:

```
DBE_SCHEDULER.set_job_argument_value(
job_name text,
argument_position integer,
argument_value text
)
```

```
DBE_SCHEDULER.set_job_argument_value(
job_name text,
argument_name text,
argument_value text
)
```

**Table 10-129** DBE\_SCHEDULER.SET\_JOB\_ARGUMENT\_VALUE API parameters

| Parameter         | Type    | Input/<br>Output<br>Parameter | Can Be<br>Empty | Description              |
|-------------------|---------|-------------------------------|-----------------|--------------------------|
| job_name          | text    | IN                            | No              | Name of a scheduled job. |
| argument_position | integer | IN                            | No              | Parameter location.      |
| argument_name     | text    | IN                            | Yes             | Parameter name.          |
| argument_value    | text    | IN                            | Yes             | Parameter value.         |

Example:

```
openGauss=# CALL db_scheduler.create_job('job1','EXTERNAL_SCRIPT','begin insert into test1
values(12); end;','2,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job
```

-----  
(1 row)

```
openGauss=# CALL DBE_SCHEDULER.set_job_argument_value('job1', 1, 'value1');
set_job_argument_value
```

-----  
(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job
```

-----  
(1 row)

- **DBE\_SCHEDULER.CREATE\_SCHEDULE**

Creates a schedule.

The prototype of the DBE\_SCHEDULER.CREATE\_SCHEDULE function is as follows:

```
DBE_SCHEDULER.create_schedule(
schedule_name text,
start_date timestamp with time zone default NULL,
repeat_interval text,
end_date timestamp with time zone default NULL,
comments text default NULL
)
```

**Table 10-130** DBE\_SCHEDULER.CREATE\_SCHEDULE API parameters

| Parameter       | Type                     | Input/Output Parameter | Can Be Empty | Description                         |
|-----------------|--------------------------|------------------------|--------------|-------------------------------------|
| schedule_name   | text                     | IN                     | No           | Name of a schedule.                 |
| start_date      | timestamp with time zone | IN                     | Yes          | Start time of a schedule.           |
| repeat_interval | text                     | IN                     | No           | Repetition frequency of a schedule. |
| end_date        | timestamp with time zone | IN                     | Yes          | End time of a schedule.             |
| comments        | text                     | IN                     | Yes          | Comments.                           |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule1', sysdate, 'sysdate + 3 / (24 * 60 * 60)', null, 'test1');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule2', sysdate, 'FREQ=DAILY; BYHOUR=6;', null, 'test1');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule3', sysdate, 'FREQ=DAILY; BYHOUR=6;');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule1');
drop_single_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule2', false);
drop_single_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule3', true);
drop_single_schedule

(1 row)
```

- DBE\_SCHEDULER.DROP\_SCHEDULE

Deletes a schedule.

The prototype of the DBE\_SCHEDULER.DROP\_SCHEDULE function is as follows:

```
DBE_SCHEDULER.drop_schedule(
 schedule_name text,
 force boolean default false
)
```

**Table 10-131** DBE\_SCHEDULER.DROP\_SCHEDULE API parameters

| Parameter     | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                            |
|---------------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| schedule_name | text    | IN                     | No           | Name of a schedule.                                                                                                                                                                                                                                                                                    |
| force         | boolean | IN                     | No           | Specifies whether to delete a schedule. <ul style="list-style-type: none"><li><b>true</b>: Any jobs or windows that use this schedule are disabled before the schedule is deleted.</li><li><b>false</b>: The schedule cannot be referenced by any job or window. Otherwise, an error occurs.</li></ul> |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule1', sysdate, 'sysdate + 3 / (24 * 60 * 60)', null, 'test1');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule2', sysdate, 'FREQ=DAILY; BYHOUR=6;', null, 'test1');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule3', sysdate, 'FREQ=DAILY; BYHOUR=6;');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule1');
drop_single_schedule

```

```
(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule2', false);
drop_single_schedule

```

```
(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_schedule('schedule3', true);
drop_single_schedule

```

```
(1 row)
```

- **DBE\_SCHEDULER.DROP\_SINGLE\_SCHEDULE**

Deletes a single schedule.

The prototype of the DBE\_SCHEDULER.DROP\_SINGLE\_SCHEDULE function is as follows:

```
DBE_SCHEDULER.drop_single_schedule(
schedule_name text,
force boolean default false
)
```

**Table 10-132** DBE\_SCHEDULER.DROP\_SINGLE\_SCHEDULE API parameters

| Parameter     | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                |
|---------------|---------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| schedule_name | text    | IN                     | No           | Name of a schedule.                                                                                                                                                                                                                                                                                        |
| force         | boolean | IN                     | No           | Specifies whether to delete a schedule. <ul style="list-style-type: none"><li>• <b>true</b>: Any jobs or windows that use this schedule are disabled before the schedule is deleted.</li><li>• <b>false</b>: The schedule cannot be referenced by any job or window. Otherwise, an error occurs.</li></ul> |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule1', sysdate, 'sysdate + 3 / (24 * 60 * 60)', null, 'test1');
create_schedule

```

```
(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule2', sysdate, 'FREQ=DAILY; BYHOUR=6;', null, 'test1');
```

```
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.create_schedule('schedule3', sysdate, 'FREQ=DAILY;
BYHOUR=6;');
create_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_single_schedule('schedule1');
drop_single_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_single_schedule('schedule2', false);
drop_single_schedule

(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_single_schedule('schedule3', true);
drop_single_schedule

(1 row)
```

- **DBE\_SCHEDULER.CREATE\_JOB\_CLASS**

Creates the class of a scheduled job.

The prototype of the DBE\_SCHEDULER.CREATE\_JOB\_CLASS function is as follows:

```
DBE_SCHEDULER.create_job_class(
job_class_name text,
resource_consumer_group text default NULL,
service text default NULL,
logging_level integer default 0,
log_history integer default NULL,
comments text default NULL
)
```

**Table 10-133** DBE\_SCHEDULER.CREATE\_JOB\_CLASS API parameters

| Parameter               | Type    | Input/Output Parameter | Can Be Empty | Description                                              |
|-------------------------|---------|------------------------|--------------|----------------------------------------------------------|
| job_class_name          | text    | IN                     | No           | Name of a scheduled job class.                           |
| resource_consumer_group | text    | IN                     | Yes          | Inline resource consumer group of a scheduled job class. |
| service                 | text    | IN                     | Yes          | Inline database service of a scheduled job class.        |
| logging_level           | integer | IN                     | No           | Number of scheduled job records.                         |

| Parameter   | Type    | Input/Output Parameter | Can Be Empty | Description                                       |
|-------------|---------|------------------------|--------------|---------------------------------------------------|
| log_history | integer | IN                     | Yes          | Number of days for storing scheduled job records. |
| comments    | text    | IN                     | Yes          | Comments.                                         |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_job_class(job_class_name => 'jc1',
resource_consumer_group => '123');
create_job_class
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_job_class('jc1', false);
drop_job_class
```

```

(1 row)
```

- **DBE\_SCHEDULER.DROP\_JOB\_CLASS**

Deletes the class of a scheduled job.

The prototype of the DBE\_SCHEDULER.DROP\_JOB\_CLASS function is as follows:

```
DBE_SCHEDULER.drop_job_class(
job_class_name text,
force boolean default false
)
```

**Table 10-134** DBE\_SCHEDULER.DROP\_JOB\_CLASS API parameters

| Parameter      | Type | Input/Output Parameter | Can Be Empty | Description                    |
|----------------|------|------------------------|--------------|--------------------------------|
| job_class_name | text | IN                     | No           | Name of a scheduled job class. |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------|---------|------------------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| force     | boolean | IN                     | No           | <p>Specifies whether to delete a scheduled job class.</p> <ul style="list-style-type: none"><li>• <b>true:</b> Jobs of this class will be disabled, and another class will be set as the default class. Only when such setting is successful, this class will be deleted.</li><li>• <b>false:</b> The class to be deleted cannot be referenced by any job. Otherwise, an error occurs.</li></ul> |

**Example:**

```
openGauss=# CALL DBE_SCHEDULER.create_job_class(job_class_name => 'jc1',
resource_consumer_group => '123');
create_job_class
```

-----  
(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_job_class('jc1', false);
drop_job_class
```

-----  
(1 row)

- **DBE\_SCHEDULER.DROP\_SINGLE\_JOB\_CLASS**

Deletes the class of a single scheduled job.

The prototype of the DBE\_SCHEDULER.DROP\_SINGLE\_JOB\_CLASS function is as follows:

```
DBE_SCHEDULER.drop_single_job_class(
job_class_name text,
force boolean default false
)
```



**Table 10-135** DBE\_SCHEDULER.DROP\_SINGLE\_JOB\_CLASS API parameters

| Parameter      | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                               |
|----------------|---------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_class_name | text    | IN                     | No           | Name of a scheduled job class.                                                                                                                                                                                                                                                                                                                                                            |
| force          | boolean | IN                     | No           | Specifies whether to delete a scheduled job class. <ul style="list-style-type: none"><li>• <b>true</b>: Jobs of this class will be disabled, and another class will be set as the default class. Only when such setting is successful, this class will be deleted.</li><li>• <b>false</b>: The class to be deleted cannot be referenced by any job. Otherwise, an error occurs.</li></ul> |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_job_class(job_class_name => 'jc1',
resource_consumer_group => '123');
create_job_class
```

-----  
(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_single_job_class('jc1', false);
drop_job_class
```

-----  
(1 row)

- **DBE\_SCHEDULER.GRANT\_USER\_AUTHORIZATION**

Grants the scheduled job permissions to the database user. The user who calls this function must have the SYSADMIN permission.

The prototype of the DBE\_SCHEDULER.GRANT\_USER\_AUTHORIZATION function is as follows:

```
DBE_SCHEDULER.grant_user_authorization(
username text,
privilege text
)
```

**Table 10-136** DBE\_SCHEDULER.GRANT\_USER\_AUTHORIZATION API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description               |
|-----------|------|------------------------|--------------|---------------------------|
| username  | text | IN                     | No           | Database username.        |
| privilege | text | IN                     | No           | Scheduled job permission. |

Example:

```
openGauss=# create user user1 password '1*G*****';
NOTICE: The encrypted password contains MD5 ciphertext, which is not secure.
CREATE ROLE
openGauss=# CALL DBE_SCHEDULER.grant_user_authorization('user1', 'create job');
grant_user_authorization

(1 row)
openGauss=# drop user user1;
DROP ROLE
```

- **DBE\_SCHEDULER.REVOKE\_USER\_AUTHORIZATION**

Revokes the scheduled job permissions from the database user. The user who calls this function must have the SYSADMIN permission.

The prototype of the DBE\_SCHEDULER.REVOKE\_USER\_AUTHORIZATION function is as follows:

```
DBE_SCHEDULER.revoke_user_authorization(
username text,
privilege text
)
```

**Table 10-137** DBE\_SCHEDULER.REVOKE\_USER\_AUTHORIZATION API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description               |
|-----------|------|------------------------|--------------|---------------------------|
| username  | text | IN                     | No           | Database username.        |
| privilege | text | IN                     | No           | Scheduled job permission. |

Example:

```
openGauss=# create user user1 password '1*G*****';
NOTICE: The encrypted password contains MD5 ciphertext, which is not secure.
```

```
CREATE ROLE
openGauss=# CALL DBE_SCHEDULER.grant_user_authorization('user1', 'create job');
grant_user_authorization

(1 row)

openGauss=# CALL DBE_SCHEDULER.revoke_user_authorization('user1', 'create job');
revoke_user_authorization

(1 row)

openGauss=# drop user user1;
DROP ROLE
```

- **DBE\_SCHEDULER.CREATE\_CREDENTIAL**

Creates an authorization certificate. The user who calls this function must have the SYSADMIN permission.

The prototype of the DBE\_SCHEDULER.CREATE\_CREDENTIAL function is as follows:

```
DBE_SCHEDULER.create_credential(
credential_name text,
username text,
password text default NULL,
database_role text default NULL,
windows_domain text default NULL,
comments text default NULL
)
```

**Table 10-138** DBE\_SCHEDULER.CREATE\_CREDENTIAL API parameters

| Parameter       | Type | Input/Output Parameter | Can Be Empty | Description                             |
|-----------------|------|------------------------|--------------|-----------------------------------------|
| credential_name | text | IN                     | No           | Name of the authorization certificate.  |
| username        | text | IN                     | No           | Database username.                      |
| password        | text | IN                     | Yes          | User password.                          |
| database_role   | text | IN                     | Yes          | Database system permission.             |
| windows_domain  | text | IN                     | Yes          | Domain to which a Windows user belongs. |
| comments        | text | IN                     | Yes          | Comments.                               |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_credential('cre_1', 'user1', '');
create_credential

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_credential('cre_1', false);
drop_credential

(1 row)
```

**NOTICE**

The **password** parameter of DBE\_SCHEDULER.CREATE\_CREDENTIAL must be set to **NULL** or '\*\*\*\*\*'. This parameter is used only for compatibility and does not indicate any actual meaning. Do not use the OS username corresponding to the installation user to create a certificate.

- DBE\_SCHEDULER.DROP\_CREDENTIAL

Destroys an authorization certificate. The user who calls this function must have the SYSADMIN permission.

The prototype of the DBE\_SCHEDULER.DROP\_CREDENTIAL function is as follows:

```
DBE_SCHEDULER.drop_credential(
credential_name text,
force boolean default false
)
```

**Table 10-139** DBE\_SCHEDULER.DROP\_CREDENTIAL API parameters

| Parameter       | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                    |
|-----------------|---------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| credential_name | text    | IN                     | No           | Name of the authorization certificate.                                                                                                                                                                                                                                                         |
| force           | boolean | IN                     | No           | Specifies whether to delete the authorization certificate. <ul style="list-style-type: none"><li><b>true:</b> The certificate will be deleted no matter whether it is referenced by any job.</li><li><b>false:</b> No job can reference the certificate. Otherwise, an error occurs.</li></ul> |

Example:

```
openGauss=# CALL DBE_SCHEDULER.create_credential('cre_1', 'user1', '');
create_credential

```

```
(1 row)

openGauss=# CALL DBE_SCHEDULER.drop_credential('cre_1', false);
drop_credential

(1 row)
```

• DBE\_SCHEDULER.ENABLE

Enables an object.

The prototype of the DBE\_SCHEDULER.ENABLE function is as follows:

```
DBE_SCHEDULER.enable(
name text,
commit_semantics text default 'STOP_ON_FIRST_ERROR'
)
```

Table 10-140 DBE\_SCHEDULER.ENABLE API parameters

| Parameter        | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------|------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name             | text | IN                     | No           | Object name. You can specify one or more objects. If you specify multiple objects, separate them with commas (,).                                                                                                                                                                                                                                                                                                                                                                                           |
| commit_semantics | text | IN                     | No           | Commit rules. The following types are supported: <ul style="list-style-type: none"><li>• <b>'STOP_ON_FIRST_ERROR'</b>: The enabling operation performed before the first error is reported is committed.</li><li>• <b>'TRANSACTIONAL'</b>: Transaction-level commit. The enabling operation performed before an error is reported will be rolled back.</li><li>• <b>'ABSORB_ERRORS'</b>: The system attempts to bypass an error and commit the enabling operation that is performed successfully.</li></ul> |

Example:

```
openGauss=# CALL db_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1 values(12);
end;','0,null,null,null','DEFAULT_JOB_CLASS',false,false,null,null,null);
lse, 'test');
create_job
```

```

(1 row)
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'stored_procedure', 'insert into
tb_job_test(key) values(null);', 0, false, '');
create_program

(1 row)
openGauss=# CALL DBE_SCHEDULER.enable('job1');
enable

(1 row)
openGauss=# CALL DBE_SCHEDULER.enable('program1', 'STOP_ON_FIRST_ERROR');
enable

(1 row)
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

(1 row)
openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

(1 row)
```

- DBE\_SCHEDULER.ENABLE\_SINGLE

Enables a single object.

The prototype of the DBE\_SCHEDULER.ENABLE\_SINGLE function is as follows:

```
DBE_SCHEDULER.enable_single(
name text
)
```

**Table 10-141** DBE\_SCHEDULER.ENABLE\_SINGLE API parameters

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description  |
|-----------|------|------------------------|--------------|--------------|
| name      | text | IN                     | No           | Object name. |

Example:

```
openGauss=# CALL dbe_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1 values(12);
end;',0,null,null,null,'DEFAULT_JOB_CLASS',false,false,null,null,null);
lse, 'test');
create_job
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.enable_single('job1');
enable_single
```

```

(1 row)
```

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job
```

```

(1 row)
```

- **DBE\_SCHEDULER.DISABLE**

Disables multiple objects. The value of name is a character string separated by commas (,). Each character string separated by commas (,) is an object. Operations are synchronized only when operation synchronization is enabled.

The prototype of the DBE\_SCHEDULER.DISABLE function is as follows:

```
DBE_SCHEDULER.disable(
name text,
force boolean default false,
commit_semantics text default 'STOP_ON_FIRST_ERROR'
)
```

**Table 10-142** DBE\_SCHEDULER.DISABLE API parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                             |
|-----------|---------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | text    | IN                     | No           | Object name.                                                                                                                                                                                                                                                            |
| force     | boolean | IN                     | No           | Specifies whether to disable an object. <ul style="list-style-type: none"><li>• <b>true</b>: The object is disabled regardless of whether other objects depend on it.</li><li>• <b>false</b>: No object can depend on the object. Otherwise, an error occurs.</li></ul> |

| Parameter        | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------|------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| commit_semantics | text | IN                     | No           | Commit rules. The following types are supported: <ul style="list-style-type: none"><li>• <b>'STOP_ON_FIRST_ERROR'</b>: The disabling operation performed before the first error is reported is committed.</li><li>• <b>'TRANSACTIONAL'</b>: Transaction-level commit. The disabling operation performed before an error is reported will be rolled back.</li><li>• <b>'ABSORB_ERRORS'</b>: The system attempts to bypass an error and commit the disabling operation that is performed successfully.</li></ul> |

## Example:

```
openGauss=# CALL db_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1 values(12);
end;','0,null,null,null,'DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.create_program('program1', 'stored_procedure', 'insert into
tb_job_test(key) values(null);', 0, false, '');
create_program
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.disable('job1');
disable
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.disable('program1', false, 'STOP_ON_FIRST_ERROR');
disable
```

-----

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job
```

-----

(1 row)



```
openGauss=# CALL DBE_SCHEDULER.drop_program('program1', false);
drop_program

```

(1 row)

- DBE\_SCHEDULER.DISABLE\_SINGLE

Disables a single object.

The prototype of the DBE\_SCHEDULER.DISABLE\_SINGLE function is as follows:

```
DBE_SCHEDULER.disable_single(
name text,
force boolean default false
)
```

**Table 10-143** DBE\_SCHEDULER.DISABLE\_SINGLE API parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                             |
|-----------|---------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name      | text    | IN                     | No           | Object name.                                                                                                                                                                                                                                                            |
| force     | boolean | IN                     | No           | Specifies whether to disable an object. <ul style="list-style-type: none"><li>• <b>true</b>: The object is disabled regardless of whether other objects depend on it.</li><li>• <b>false</b>: No object can depend on the object. Otherwise, an error occurs.</li></ul> |

Example:

```
openGauss=# CALL db_scheduler.create_job('job1','PLSQL_BLOCK','begin insert into test1 values(12);
end;',0,null,null,null,'DEFAULT_JOB_CLASS',false,false,null,null,null);
create_job

```

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.disable_single('job1', false);
disable_single

```

(1 row)

```
openGauss=# CALL DBE_SCHEDULER.drop_job('job1', true, false, 'STOP_ON_FIRST_ERROR');
drop_job

```

(1 row)

- DBE\_SCHEDULER.EVAL\_CALENDAR\_STRING

Analyzes the scheduling job period.

Return type: timestamp with time zone

The prototype of the DBE\_SCHEDULER.EVAL\_CALENDAR\_STRING function is as follows:

```
DBE_SCHEDULER.evaluate_calendar_string(
 IN calendar_string text,
 IN start_date timestamp with time zone,
 IN return_date_after timestamp with time zone
)return timestamp with time zone
```

**Table 10-144** DBE\_SCHEDULER.EVAL\_CALENDAR\_STRING API parameters

| Parameter         | Type                     | Input/<br>Output<br>Param<br>eter | Can Be<br>Empty | Description                            |
|-------------------|--------------------------|-----------------------------------|-----------------|----------------------------------------|
| calendar_string   | text                     | IN                                | No              | Date string of a scheduled job.        |
| start_date        | timestamp with time zone | IN                                | No              | Start time of a scheduled job.         |
| return_date_after | timestamp with time zone | IN                                | No              | Date when a scheduled job is returned. |

Example:

```
openGauss=# CALL DBE_SCHEDULER.eval_calendar_string('FREQ=DAILY; BYHOUR=6;', sysdate,
eval_calendar_string

2023-09-16 06:05:55+08
(1 row)
```

- DBE\_SCHEDULER.EVALUATE\_CALENDAR\_STRING

Analyzes the scheduling job period.

Return type: timestamp with time zone

The prototype of the DBE\_SCHEDULER.EVALUATE\_CALENDAR\_STRING function is as follows:

```
DBE_SCHEDULER.evaluate_calendar_string(
 IN calendar_string text,
 IN start_date timestamp with time zone,
 IN return_date_after timestamp with time zone,
 OUT next_run_date timestamp with time zone
)return timestamp with time zone
```

**Table 10-145** DBE\_SCHEDULER.EVALUATE\_CALENDAR\_STRING API parameters

| Parameter         | Type                     | Input/Output Parameter | Can Be Empty | Description                                     |
|-------------------|--------------------------|------------------------|--------------|-------------------------------------------------|
| calendar_string   | text                     | IN                     | No           | Date string of a scheduled job.                 |
| start_date        | timestamp with time zone | IN                     | No           | Start time of a scheduled job.                  |
| return_date_after | timestamp with time zone | IN                     | No           | Date when a scheduled job is returned.          |
| next_run_date     | timestamp with time zone | OUT                    | No           | The next date when a scheduled job is returned. |

Example:

```
openGauss=# CREATE OR REPLACE PROCEDURE pr1(calendar_str text) as
DECLARE
 start_date timestamp with time zone;
 return_date_after timestamp with time zone;
 next_run_date timestamp with time zone;
BEGIN
 start_date := '2003-2-1 10:30:00.111111+8':timestamp with time zone;
 return_date_after := start_date;
 DBE_SCHEDULER.evaluate_calendar_string(
 calendar_str,
 start_date, return_date_after, next_run_date);
 DBE_OUTPUT.PRINT_LINE('next_run_date: ' || next_run_date);
 return_date_after := next_run_date;
END;
/
CREATE PROCEDURE
openGauss=# CALL pr1('FREQ=hourly;INTERVAL=2;BYHOUR=6,10;BYMINUTE=0;BYSECOND=0');
next_run_date: 2003-02-02 06:00:00+08
pr1

(1 row)
```

### 10.12.2.7 DBE\_SQL

#### API Description

**Table 10-146** lists APIs supported by the **DBE\_SQL** package.

**Table 10-146** DBE\_SQL

| API                             | Description     |
|---------------------------------|-----------------|
| <b>DBE_SQL.REGISTER_CONTEXT</b> | Opens a cursor. |

| API                                                  | Description                                                       |
|------------------------------------------------------|-------------------------------------------------------------------|
| <a href="#">DBE_SQL.SQL_UNREGISTER_CON<br/>TEXT</a>  | Closes an open cursor.                                            |
| <a href="#">DBE_SQL.SQL_SET_SQL</a>                  | Passes a set of SQL statements to a cursor.                       |
| <a href="#">DBE_SQL.SQL_RUN</a>                      | Performs a set of dynamically defined operations on a cursor.     |
| <a href="#">DBE_SQL.NEXT_ROW</a>                     | Reads a row of cursor data.                                       |
| <a href="#">DBE_SQL.SET_RESULT_TYPE</a>              | Dynamically defines a column.                                     |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_CHA<br/>R</a>    | Dynamically defines a column of the <b>CHAR</b> type.             |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_INT</a>          | Dynamically defines a column of the <b>INT</b> type.              |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_LON<br/>G</a>    | Dynamically defines a column of the <b>LONG</b> type.             |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_RAW</a>          | Dynamically defines a column of the <b>RAW</b> type.              |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_BYTE<br/>A</a>   | Dynamically defines a column of the bytea type.                   |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_TEXT</a>         | Dynamically defines a column of the <b>TEXT</b> type.             |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_UNK<br/>NOWN</a> | Dynamically defines a column of an unknown type.                  |
| <a href="#">DBE_SQL.GET_RESULT</a>                   | Reads a dynamically defined column value.                         |
| <a href="#">DBE_SQL.GET_RESULT_CHAR</a>              | Reads a dynamically defined column value of the <b>CHAR</b> type. |
| <a href="#">DBE_SQL.GET_RESULT_INT</a>               | Reads a dynamically defined column value of the <b>INT</b> type.  |
| <a href="#">DBE_SQL.GET_RESULT_LONG</a>              | Reads a dynamically defined column value of the <b>LONG</b> type. |
| <a href="#">DBE_SQL.GET_RESULT_RAW</a>               | Reads a dynamically defined column value of the <b>RAW</b> type.  |
| <a href="#">DBE_SQL.GET_RESULT_BYTEA</a>             | Reads a dynamically defined column value of the bytea type.       |
| <a href="#">DBE_SQL.GET_RESULT_TEXT</a>              | Reads a dynamically defined column value of the <b>TEXT</b> type. |
| <a href="#">DBE_SQL.GET_RESULT_UNKNOW<br/>N</a>      | Reads a dynamically defined column value of an unknown type.      |

| API                                             | Description                                                              |
|-------------------------------------------------|--------------------------------------------------------------------------|
| <a href="#">DBE_SQL.DBE_SQL_GET_RESULT_CHAR</a> | Reads a dynamically defined column value of the <b>CHAR</b> type.        |
| <a href="#">DBE_SQL.DBE_SQL_GET_RESULT_LONG</a> | Reads a dynamically defined column value of the <b>LONG</b> type.        |
| <a href="#">DBE_SQL.DBE_SQL_GET_RESULT_RAW</a>  | Reads a dynamically defined column value of the <b>RAW</b> type.         |
| <a href="#">DBE_SQL.IS_ACTIVE</a>               | Checks whether a cursor is opened.                                       |
| <a href="#">DBE_SQL.LAST_ROW_COUNT</a>          | Compatible API. This function is not supported currently.                |
| <a href="#">DBE_SQL.RUN_AND_NEXT</a>            | Reserved API. This function is not supported currently.                  |
| <a href="#">DBE_SQL.SQL_BIND_VARIABLE</a>       | Binds a value to a variable in a statement.                              |
| <a href="#">DBE_SQL.SQL_BIND_ARRAY</a>          | Binds a group of values to a variable in a statement.                    |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_INTS</a>    | Dynamically defines a column of the <b>INT</b> array type.               |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_TEXTS</a>   | Dynamically defines a column of the <b>TEXT</b> array type.              |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_RAWS</a>    | Dynamically defines a column of the <b>RAW</b> array type.               |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_BYTEAS</a>  | Dynamically defines a column of the <b>BYTEA</b> array type.             |
| <a href="#">DBE_SQL.SET_RESULT_TYPE_CHARS</a>   | Dynamically defines a column of the <b>CHAR</b> array type.              |
| <a href="#">DBE_SQL.SET_RESULTS_TYPE</a>        | Dynamically defines a column of the array type.                          |
| <a href="#">DBE_SQL.GET_RESULTS_INT</a>         | Reads a dynamically defined column value of the <b>INT</b> array type.   |
| <a href="#">DBE_SQL.GET_RESULTS_TEXT</a>        | Reads a dynamically defined column value of the <b>TEXT</b> array type.  |
| <a href="#">DBE_SQL.GET_RESULTS_RAW</a>         | Reads a dynamically defined column value of the <b>RAW</b> array type.   |
| <a href="#">DBE_SQL.GET_RESULTS_BYTEA</a>       | Reads a dynamically defined column value of the <b>BYTEA</b> array type. |
| <a href="#">DBE_SQL.GET_RESULTS_CHAR</a>        | Reads a dynamically defined column value of the <b>CHAR</b> array type.  |

| API                                              | Description                                                          |
|--------------------------------------------------|----------------------------------------------------------------------|
| <a href="#">DBE_SQL.GET_RESULTS</a>              | Reads a dynamically defined column value.                            |
| <a href="#">DBE_SQL.SQL_DESCRIBE_COLUMNS</a>     | Describes the column information read by the cursor.                 |
| <a href="#">DBE_SQL.DESC_REC</a>                 | Stores the type of the column information read by the cursor.        |
| <a href="#">DBE_SQL.DESC_TAB</a>                 | TABLE type of DESC_REC.                                              |
| <a href="#">DBE_SQL.DATE_TABLE</a>               | TABLE type of DATE.                                                  |
| <a href="#">DBE_SQL.NUMBER_TABLE</a>             | TABLE type of NUMBER.                                                |
| <a href="#">DBE_SQL.VARCHAR2_TABLE</a>           | TABLE type of VARCHAR2.                                              |
| <a href="#">DBE_SQL.BIND_VARIABLE</a>            | Binds parameters.                                                    |
| <a href="#">DBE_SQL.SQL_SET_RESULTS_TYPE_C</a>   | Dynamically defines a column of the array type.                      |
| <a href="#">DBE_SQL.SQL_GET_VALUES_C</a>         | Reads a dynamically defined column value.                            |
| <a href="#">DBE_SQL.GET_VARIABLE_RESULT</a>      | Reads the return value of an SQL statement.                          |
| <a href="#">DBE_SQL.GET_VARIABLE_RESULT_CHAR</a> | Reads the return value (of the char type) of an SQL statement.       |
| <a href="#">DBE_SQL.GET_VARIABLE_RESULT_RAW</a>  | Reads the return value (of the raw type) of an SQL statement.        |
| <a href="#">DBE_SQL.GET_VARIABLE_RESULT_TEXT</a> | Reads the return value (of the text type) of an SQL statement.       |
| <a href="#">DBE_SQL.GET_VARIABLE_RESULT_INT</a>  | Reads the return value (of the int type) of an SQL statement.        |
| <a href="#">DBE_SQL.GET_ARRAY_RESULT_TEXT</a>    | Reads the return value (of the text array type) of an SQL statement. |
| <a href="#">DBE_SQL.GET_ARRAY_RESULT_RAW</a>     | Reads the return value (of the raw array type) of an SQL statement.  |
| <a href="#">DBE_SQL.GET_ARRAY_RESULT_CHAR</a>    | Reads the return value (of the char array type) of an SQL statement. |
| <a href="#">DBE_SQL.GET_ARRAY_RESULT_INT</a>     | Reads the return value (of the int array type) of an SQL statement.  |

 NOTE

- You are advised to use **dbe\_sql.set\_result\_type** and **dbe\_sql.get\_result** to define columns.
- If the size of the result set is greater than the value of **work\_mem**, the result set will be spilled to a disk temporarily. The value of **work\_mem** must be no greater than 512 MB.

- DBE\_SQL.REGISTER\_CONTEXT

This function opens a cursor, which is the prerequisite for the subsequent **dbe\_sql** operations. This function does not transfer any parameter. It automatically generates cursor IDs in an ascending order and returns values to integer variables.

The prototype of the DBE\_SQL.REGISTER\_CONTEXT function is as follows:

```
DBE_SQL.REGISTER_CONTEXT(
)
RETURN INTEGER;
```

- DBE\_SQL.SQL\_UNREGISTER\_CONTEXT

This function closes a cursor, which is the end of each **dbe\_sql** operation. If this function is not called when the stored procedure ends, the memory is still occupied by the cursor. Therefore, remember to close a cursor when you do not need to use it. If an exception occurs, the stored procedure exits but the cursor is not closed. Therefore, you are advised to include this API in the exception handling of the stored procedure.

The prototype of the DBE\_SQL.SQL\_UNREGISTER\_CONTEXT function is as follows:

```
DBE_SQL.SQL_UNREGISTER_CONTEXT(
context_id IN INTEGER
)
RETURN INTEGER;
```

**Table 10-147** DBE\_SQL.SQL\_UNREGISTER\_CONTEXT parameters

| Parameter  | Description                    |
|------------|--------------------------------|
| context_id | ID of the cursor to be closed. |

- DBE\_SQL.SQL\_SET\_SQL

This function is used to parse the query statement of a specified cursor. The statement parameters can be specified only through the text type. The length cannot exceed 1 GB.

The prototype of the **DBE\_SQL.SQL\_SET\_SQL** function is as follows:

```
DBE_SQL.SQL_SET_SQL(
context_id IN INTEGER,
query_string IN TEXT,
language_flag IN INTEGER
)
RETURN BOOLEAN;
```

**Table 10-148** DBE\_SQL.SQL\_SET\_SQL parameters

| Parameter     | Description                                                     |
|---------------|-----------------------------------------------------------------|
| context_id    | ID of the cursor whose query statement is to be parsed.         |
| query_string  | Query statement to be parsed.                                   |
| language_flag | Version language number. Currently, only <b>1</b> is supported. |

- DBE\_SQL.SQL\_RUN

This function executes a given cursor. It receives a cursor ID first, and the data obtained after execution is used for subsequent operations.

The prototype of the **DBE\_SQL.SQL\_RUN** function is as follows:

```
DBE_SQL.SQL_RUN(
context_id IN INTEGER,
)
RETURN INTEGER;
```

**Table 10-149** DBE\_SQL.SQL\_RUN parameters

| Parameter  | Description                                             |
|------------|---------------------------------------------------------|
| context_id | ID of the cursor whose query statement is to be parsed. |

- DBE\_SQL.NEXT\_ROW

This function returns the number of data rows that meet query conditions. Each time the API is executed, the system obtains a set of new rows until all data is read.

The prototype of the **DBE\_SQL.NEXT\_ROW** function is as follows:

```
DBE_SQL.NEXT_ROW(
context_id IN INTEGER,
)
RETURN INTEGER;
```

**Table 10-150** DBE\_SQL.NEXT\_ROW parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be executed |

- DBE\_SQL.SET\_RESULT\_TYPE

This function defines columns returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE(
context_id IN INTEGER,
pos IN INTEGER,
```



```
column_ref IN ANYELEMENT,
maxsize IN INTEGER default 1024
)
RETURN INTEGER;
```

**Table 10-151** DBE\_SQL.SET\_RESULT\_TYPE parameters

| Parameter  | Description                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------|
| context_id | ID of the cursor to be executed                                                                                |
| pos        | Position of a dynamically defined column in the query                                                          |
| column_ref | Variable of any type. You can select an appropriate API to dynamically define columns based on variable types. |
| maxsize    | Length of a defined column.                                                                                    |

- DBE\_SQL.SET\_RESULT\_TYPE\_CHAR

This function defines columns of the **CHAR** type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_CHAR** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_CHAR(
context_id IN INTEGER,
pos IN INTEGER,
column_ref IN TEXT,
column_size IN INTEGER
)
RETURN INTEGER;
```

**Table 10-152** DBE\_SQL.SET\_RESULT\_TYPE\_CHAR parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be executed                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Parameter to be defined.                              |
| column_size | Length of a dynamically defined column.               |

- DBE\_SQL.SET\_RESULT\_TYPE\_INT

This function defines columns of the **INT** type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_INT** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_INT(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN INTEGER;
```

**Table 10-153** DBE\_SQL.SET\_RESULT\_TYPE\_INT parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.SET\_RESULT\_TYPE\_LONG

This function defines columns of a long type (not **LONG**) returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type. The maximum size of a long column is 1 GB.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_LONG** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_LONG(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN INTEGER;
```

**Table 10-154** DBE\_SQL.SET\_RESULT\_TYPE\_LONG parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.SET\_RESULT\_TYPE\_RAW

This function defines columns of the **RAW** type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_RAW** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_RAW(
context_id IN INTEGER,
pos IN INTEGER,
column_ref IN RAW,
column_size IN INTEGER
)
RETURN INTEGER;
```

**Table 10-155** DBE\_SQL.SET\_RESULT\_TYPE\_RAW parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be executed                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Parameter of the RAW type.                            |
| column_size | Column length.                                        |

- DBE\_SQL.SET\_RESULT\_TYPE\_BYTEA

Defines columns of the BYTEA type returned from a given cursor and can be used only for cursors defined by SELECT. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the DBE\_SQL.SET\_RESULT\_TYPE\_BYTEA function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_BYTEA(
context_id IN INTEGER,
pos IN INTEGER,
column_ref IN BYTEA,
column_size IN INTEGER
)
RETURN INTEGER;
```

**Table 10-156** DBE\_SQL.SET\_RESULT\_TYPE\_BYTEA parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be executed                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Parameter of the BYTEA type.                          |
| column_size | Column length.                                        |

- DBE\_SQL.SET\_RESULT\_TYPE\_TEXT

This function defines columns of the **TEXT** type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_TEXT** function is as follows:

```
DBE_SQL.DEFINE_COLUMN_CHAR(
context_id IN INTEGER,
pos IN INTEGER,
maxsize IN INTEGER
)
RETURN INTEGER;
```

**Table 10-157** DBE\_SQL.SET\_RESULT\_TYPE\_TEXT parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |
| maxsize    | Maximum length of the defined <b>TEXT</b> type        |

- DBE\_SQL.SET\_RESULT\_TYPE\_UNKNOWN

This function processes columns of unknown data types returned from a given cursor. It is used only for the system to report an error and exist when the type cannot be identified.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_UNKNOWN** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_UNKNOWN(
context_id IN INTEGER,
pos IN INTEGER,
col_type IN TEXT
)
RETURN INTEGER;
```

**Table 10-158** DBE\_SQL.SET\_RESULT\_TYPE\_UNKNOWN parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| posn       | Position of a dynamically defined column in the query |
| col_type   | Dynamically defined parameter.                        |

- DBE\_SQL.GET\_RESULT

This function returns the cursor element value in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULT** function is as follows:

```
DBE_SQL.GET_RESULT(
context_id IN INTEGER,
pos IN INTEGER,
column_value INOUT ANYELEMENT
)
RETURN ANYELEMENT;
```

**Table 10-159** DBE\_SQL.GET\_RESULT parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

| Parameter    | Description                       |
|--------------|-----------------------------------|
| column_value | Return value of a defined column. |

- DBE\_SQL.GET\_RESULT\_CHAR

This stored procedure returns the value of the **CHAR** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the DBE\_SQL.GET\_RESULT\_CHAR stored procedure is as follows:

```
DBE_SQL.GET_RESULT_CHAR(
 context_id IN INTEGER,
 pos IN INTEGER,
 tr INOUT CHARACTER,
 err INOUT NUMERIC,
 actual_length INOUT INTEGER
);
```

**Table 10-160** DBE\_SQL.GET\_RESULT\_CHAR parameters

| Parameter     | Description                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| context_id    | ID of the cursor to be executed                                                                                                                          |
| pos           | Position of a dynamically defined column in the query                                                                                                    |
| tr            | Return value                                                                                                                                             |
| err           | Error No. It is an output parameter. The input parameter must be a variable. Currently, the output value is <b>-1</b> regardless of the input parameter. |
| actual_length | Length of a return value.                                                                                                                                |

The overloading of the DBE\_SQL.GET\_RESULT\_CHAR stored procedure is as follows:

```
DBE_SQL.GET_RESULT_CHAR(
 context_id IN INTEGER,
 pos IN INTEGER,
 tr INOUT CHARACTER
);
```

**Table 10-161** DBE\_SQL.GET\_RESULT\_CHAR parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |
| tr         | Return value                                          |

- DBE\_SQL.GET\_RESULT\_INT

This function returns the value of the **INT** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**. The prototype of the **DBE\_SQL.GET\_RESULT\_INT** function is as follows:

```
DBE_SQL.GET_RESULT_INT(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN INTEGER;
```

**Table 10-162** DBE\_SQL.GET\_RESULT\_INT parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.GET\_RESULT\_LONG

This function returns the value of a long type (not **LONG** or **BIGINT**) in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULT\_LONG** function is as follows:

```
DBE_SQL.GET_RESULT_LONG(
context_id IN INTEGER,
pos IN INTEGER,
lgth IN INTEGER,
off_set IN INTEGER,
vl INOUT TEXT,
vl_length INOUT INTEGER
)
RETURN RECORD;
```

**Table 10-163** DBE\_SQL.GET\_RESULT\_LONG parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |
| lgth       | Length of a return value.                             |
| off_set    | Start position of a return value.                     |
| vl         | Return value                                          |
| vl_length  | Length of a return value.                             |

- DBE\_SQL.GET\_RESULT\_RAW

Returns the value of the **RAW** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULT\_RAW** stored procedure is as follows:

```
DBE_SQL.GET_RESULT_RAW(
context_id IN INTEGER,
pos IN INTEGER,
tr INOUT RAW,
err INOUT NUMERIC,
actual_length INOUT INTEGER
);
```

**Table 10-164** DBE\_SQL.GET\_RESULT\_RAW parameters

| Parameter     | Description                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| context_id    | ID of the cursor to be executed                                                                                                                          |
| pos           | Position of a dynamically defined column in the query                                                                                                    |
| tr            | Returned column value.                                                                                                                                   |
| err           | Error No. It is an output parameter. The input parameter must be a variable. Currently, the output value is <b>-1</b> regardless of the input parameter. |
| actual_length | Length of a return value. The value longer than this length will be truncated.                                                                           |

- The overloading of the DBE\_SQL.GET\_RESULT\_RAW stored procedure is as follows:

```
DBE_SQL.GET_RESULT_RAW(
context_id IN INTEGER,
pos IN INTEGER,
tr INOUT RAW
);
```

**Table 10-165** DBE\_SQL.GET\_RESULT\_RAW parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |
| tr         | Returned column value.                                |

- DBE\_SQL.GET\_RESULT\_BYTEA

Returns the value of the BYTEA type in a specified position of a cursor that is obtained by DBE\_SQL.NEXT\_ROW.

The prototype of the DBE\_SQL.GET\_RESULT\_BYTEA stored procedure is as follows:

```
DBE_SQL.GET_RESULT_BYTEA(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN BYTEA;
```

**Table 10-166** DBE\_SQL.GET\_RESULT\_BYTEA parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.GET\_RESULT\_TEXT

This function returns the value of the **TEXT** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULT\_TEXT** function is as follows:

```
DBE_SQL.GET_RESULT_TEXT(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN TEXT;
```

**Table 10-167** DBE\_SQL.GET\_RESULT\_TEXT parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.GET\_RESULT\_UNKNOWN

This function returns the value of an unknown type in a specified position of a cursor. It serves as an error handling API when the type is not unknown.

The prototype of the **DBE\_SQL.GET\_RESULT\_UNKNOWN** function is as follows:

```
DBE_SQL.GET_RESULT_UNKNOWN(
context_id IN INTEGER,
pos IN INTEGER,
col_type IN TEXT
)
RETURN TEXT;
```

**Table 10-168** DBE\_SQL.GET\_RESULT\_UNKNOWN parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |
| col_type   | Returned parameter type.                              |

- DBE\_SQL.DBE\_SQL\_GET\_RESULT\_CHAR

This function returns the value of the **CHAR** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**. Different



from **DBE\_SQL.GET\_RESULT\_CHAR**, the length of the return value is not set and the entire string is returned.

The prototype of the **DBE\_SQL.DBE\_SQL\_GET\_RESULT\_CHAR** function is as follows:

```
DBE_SQL.DBE_SQL_GET_RESULT_CHAR(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN CHARACTER;
```

**Table 10-169** DBE\_SQL.GET\_RESULT\_CHAR parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- **DBE\_SQL.DBE\_SQL\_GET\_RESULT\_LONG**

This function returns the value of a long type (not **LONG** or **BIGINT**) in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

Different from **DBE\_SQL.GET\_RESULT\_LONG**, the length of the return value is not set and the entire BIGINT value is returned.

The prototype of the **DBE\_SQL.DBE\_SQL\_GET\_RESULT\_LONG** function is as follows:

```
DBE_SQL.DBE_SQL_GET_RESULT_LONG(
context_id IN INTEGER,
pos IN INTEGER
)
RETURN BIGINT;
```

**Table 10-170** DBE\_SQL.DBE\_SQL\_GET\_RESULT\_LONG parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- **DBE\_SQL.DBE\_SQL\_GET\_RESULT\_RAW**

This function returns the value of the **RAW** type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

Different from **DBE\_SQL.GET\_RESULT\_RAW**, the length of the return value is not set and the entire string is returned.

The prototype of the **DBE\_SQL.DBE\_SQL\_GET\_RESULT\_RAW** function is as follows:

```
DBE_SQL.GET_RESULT_RAW(
context_id IN INTEGER,
pos IN INTEGER,
tr INOUT RAW
)
RETURN RAW;
```

**Table 10-171** DBE\_SQL.GET\_RESULT\_RAW parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be executed                       |
| pos        | Position of a dynamically defined column in the query |

- DBE\_SQL.IS\_ACTIVE

This function returns the status of a cursor. The status can be **open**, **parse**, **execute**, or **define**. If the status is **open**, the value is **TRUE**. If the status is unknown, an error is reported. In other cases, the value is **FALSE**.

The prototype of the **DBE\_SQL.IS\_ACTIVE** function is as follows:

```
DBE_SQL.IS_ACTIVE(
 context_id IN INTEGER
)
RETURN BOOLEAN;
```

**Table 10-172** DBE\_SQL.IS\_ACTIVE parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be queried. |

- DBE\_SQL.SQL\_BIND\_VARIABLE

This function is used to bind a parameter to an SQL statement. When an SQL statement is executed, the SQL statement is executed based on the bound value.

The prototype of the **DBE\_SQL.SQL\_BIND\_VARIABLE** function is as follows:

```
DBE_SQL.SQL_BIND_VARIABLE(
 context_id in int,
 query_string in text,
 language_flag in anyelement,
 out_value_size in int default null
)
RETURNS void;
```

**Table 10-173** DBE\_SQL.SQL\_BIND\_VARIABLE parameters

| Parameter      | Description                                            |
|----------------|--------------------------------------------------------|
| context_id     | ID of the cursor to be queried.                        |
| query_string   | Name of a bind variable.                               |
| language_flag  | Bound value.                                           |
| out_value_size | Size of the return value. Default value: <b>null</b> . |

- DBE\_SQL.SQL\_BIND\_ARRAY

This function is used to bind a set of parameters to an SQL statement. When an SQL statement is executed, the SQL statement is executed based on the bound array.

The prototype of the **DBE\_SQL.SQL\_BIND\_ARRAY** function is as follows:

```
DBE_SQL.SQL_BIND_ARRAY(
 IN context_id int,
 IN query_string text,
 IN value anyarray
)
RETURNS void;
DBE_SQL.SQL_BIND_ARRAY(
 IN context_id int,
 IN query_string text,
 IN value anyarray,
 IN lower_index int default 1,
 IN higher_index int default 1
)
RETURNS void;
```

**Table 10-174** DBE\_SQL.SQL\_BIND\_ARRAY parameters

| Parameter    | Description                       |
|--------------|-----------------------------------|
| context_id   | ID of the cursor to be queried.   |
| query_string | Name of a bind variable.          |
| value        | Bound array.                      |
| lower_index  | Minimum index of the bound array. |
| higher_index | Maximum index of the bound array. |

- **DBE\_SQL.SET\_RESULT\_TYPE\_INTS**

This function defines columns of the **INT** array type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_INTS** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_INTS(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int
)
RETURNS integer;
```

**Table 10-175** DBE\_SQL.SET\_RESULT\_TYPE\_INTS parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be queried.                       |
| pos        | Position of a dynamically defined column in the query |
| column_ref | Type of the returned array.                           |

| Parameter | Description                            |
|-----------|----------------------------------------|
| cnt       | Number of values obtained at a time.   |
| lower_bnd | Start index when an array is returned. |

- DBE\_SQL.SET\_RESULT\_TYPE\_TEXTS

This function defines columns of the **TEXT** array type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_TEXTS** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_TEXTS(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int,
 IN maxsize int
)
RETURNS integer;
```

**Table 10-176** DBE\_SQL.SET\_RESULT\_TYPE\_TEXTS parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be queried.                       |
| pos        | Position of a dynamically defined column in the query |
| column_ref | Type of the returned array.                           |
| cnt        | Number of values obtained at a time.                  |
| lower_bnd  | Start index when an array is returned.                |
| maxsize    | Maximum length of the defined <b>TEXT</b> type        |

- DBE\_SQL.SET\_RESULT\_TYPE\_RAWS

This function defines columns of the **RAW** array type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_RAWS** function is as follows:

```
DBE_SQL.set_result_type_raws(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int,
```

```
 IN column_size int
)
RETURNS integer;
```

**Table 10-177** DBE\_SQL.SET\_RESULT\_TYPE\_RAWS parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be queried.                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Type of the returned array.                           |
| cnt         | Number of values obtained at a time.                  |
| lower_bnd   | Start index when an array is returned.                |
| column_size | Column length.                                        |

- DBE\_SQL.SET\_RESULT\_TYPE\_BYTEAS

This function defines columns of the **BYTEA** array type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_BYTEAS** function is as follows:

```
DBE_SQL.set_result_type_byteas(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int,
 IN column_size int
)
RETURNS integer;
```

**Table 10-178** DBE\_SQL.SET\_RESULT\_TYPE\_BYTEAS parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be queried.                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Type of the returned array.                           |
| cnt         | Number of values obtained at a time.                  |
| lower_bnd   | Start index when an array is returned.                |
| column_size | Column length.                                        |

- DBE\_SQL.SET\_RESULT\_TYPE\_CHARS

This function defines columns of the **CHAR** array type returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULT\_TYPE\_CHARS** function is as follows:

```
DBE_SQL.SET_RESULT_TYPE_CHARS(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int,
 IN column_size int
)
RETURNS integer;
```

**Table 10-179** DBE\_SQL.SET\_RESULT\_TYPE\_CHARS parameters

| Parameter   | Description                                           |
|-------------|-------------------------------------------------------|
| context_id  | ID of the cursor to be queried.                       |
| pos         | Position of a dynamically defined column in the query |
| column_ref  | Type of the returned array.                           |
| cnt         | Number of values obtained at a time.                  |
| lower_bnd   | Start index when an array is returned.                |
| column_size | Column length.                                        |

- **DBE\_SQL.SET\_RESULTS\_TYPE**

This function defines columns returned from a given cursor and can be used only for the cursors defined by **SELECT**. The defined columns are identified by the relative positions in the query list. The data type of an input variable determines the corresponding column type.

The prototype of the **DBE\_SQL.SET\_RESULTS\_TYPE** function is as follows:

```
DBE_SQL.SET_RESULTS_TYPE(
 IN context_id int,
 IN pos int,
 IN column_ref anyarray,
 IN cnt int,
 IN lower_bnd int,
 IN maxsize int DEFAULT 1024
) returns void;
```

**Table 10-180** DBE\_SQL.SET\_RESULTS\_TYPE parameters

| Parameter  | Description                                           |
|------------|-------------------------------------------------------|
| context_id | ID of the cursor to be queried.                       |
| pos        | Position of a dynamically defined column in the query |

| Parameter  | Description                            |
|------------|----------------------------------------|
| column_ref | Type of the returned array.            |
| cnt        | Number of values obtained at a time.   |
| lower_bnd  | Start index when an array is returned. |
| maxsize    | Maximum length of the defined type.    |

- DBE\_SQL.GET\_RESULTS\_INT

This function returns the value of the **INT** array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULTS\_INT** function is as follows:

```
DBE_SQL.GET_RESULTS_INT(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```

**Table 10-181** DBE\_SQL.GET\_RESULTS\_INT parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- DBE\_SQL.GET\_RESULTS\_TEXT

This function returns the value of the **TEXT** array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULTS\_TEXT** function is as follows:

```
DBE_SQL.GET_RESULTS_TEXT(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```

**Table 10-182** DBE\_SQL.GET\_RESULTS\_TEXT parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- DBE\_SQL.GET\_RESULTS\_RAW

This function returns the value of the **RAW** array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULTS\_RAW** function is as follows:

```
DBE_SQL.GET_RESULTS_RAW(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```

**Table 10-183** DBE\_SQL.GET\_RESULTS\_RAW parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- **DBE\_SQL.GET\_RESULTS\_BYTEA**

This function returns the value of the **BYTEA** array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULTS\_BYTEA** function is as follows:

```
DBE_SQL.GET_RESULTS_BYTEA(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```

**Table 10-184** DBE\_SQL.GET\_RESULTS\_BYTEA parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- **DBE\_SQL.GET\_RESULTS\_CHAR**

This function returns the value of the **CHAR** array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

The prototype of the **DBE\_SQL.GET\_RESULTS\_CHAR** function is as follows:

```
DBE_SQL.GET_RESULTS_CHAR(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```



**Table 10-185** DBE\_SQL.GET\_RESULTS\_CHAR parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- DBE\_SQL.GET\_RESULTS

This function returns the value of the array type in a specified position of a cursor and accesses the data obtained by **DBE\_SQL.NEXT\_ROW**.

 **NOTE**

The bottom-layer mechanism of DBE\_SQL.GET\_RESULTS is implemented through arrays. When different arrays are used to obtain the return value of the same column, NULL values are filled in the array due to discontinuous internal indexes to ensure the continuity of array indexes. As a result, the length of the returned result array is different from that of the database A.

The prototype of the **DBE\_SQL.GET\_RESULTS** function is as follows:

```
DBE_SQL.GET_RESULTS(
 IN context_id int,
 IN pos int,
 INOUT column_value anyarray
);
```

**Table 10-186** DBE\_SQL.GET\_RESULTS parameters

| Parameter    | Description                                           |
|--------------|-------------------------------------------------------|
| context_id   | ID of the cursor to be queried.                       |
| pos          | Position of a dynamically defined column in the query |
| column_value | Return value                                          |

- DBE\_SQL.SQL\_DESCRIBE\_COLUMNS

This function is used to describe column information and can be used only for cursors defined by **SELECT**.

The prototype of the **DBE\_SQL.SQL\_DESCRIBE\_COLUMNS** function is as follows:

```
DBE_SQL.SQL_DESCRIBE_COLUMNS(
 context_id in int,
 col_cnt inout int,
 desc_t inout db_sql.desc_tab
)RETURNS record ;
```

**Table 10-187** DBE\_SQL.SQL\_DESCRIBE\_COLUMNS parameters

| Parameter  | Description                         |
|------------|-------------------------------------|
| context_id | ID of the cursor to be queried.     |
| col_cnt    | Number of columns returned.         |
| desc_t     | Description of the returned column. |

- DBE\_SQL.DESC\_REC

This type is a composite type and is used to store the description of the SQL\_DESCRIBE\_COLUMNS API.

The prototype of the DBE\_SQL.DESC\_REC function is as follows:

```
CREATE TYPE DBE_SQL.DESC_REC AS (
 col_type int,
 col_max_len int,
 col_name VARCHAR2(32),
 col_name_len int,
 col_schema_name VARCHAR2(32),
 col_schema_name_len int,
 col_precision int,
 col_scale int,
 col_charsetid int,
 col_charsetform int,
 col_null_ok BOOLEAN
);
```

- DBE\_SQL.DESC\_TAB

This type is the TABLE type of DESC\_REC and is implemented through the TABLE OF syntax.

The prototype of the DBE\_SQL.DESC\_TAB function is as follows:

```
CREATE TYPE DBE_SQL.DESC_TAB AS TABLE OF DBE_SQL.DESC_REC;
```

- DBE\_SQL.DATE\_TABLE

This type is the TABLE type of DATE and is implemented through the TABLE OF syntax.

The prototype of the DBE\_SQL.DATE\_TABLE function is as follows:

```
CREATE TYPE DBE_SQL.DATE_TABLE AS TABLE OF DATE;
```

- DBE\_SQL.NUMBER\_TABLE

This type is the TABLE type of NUMBER and is implemented through the TABLE OF syntax.

The prototype of the DBE\_SQL.NUMBER\_TABLE function is as follows:

```
CREATE TYPE DBE_SQL.NUMBER_TABLE AS TABLE OF NUMBER;
```

- DBE\_SQL.VARCHAR2\_TABLE

This type is the TABLE type of VARCHAR2 and is implemented through the TABLE OF syntax.

The prototype of the DBE\_SQL.VARCHAR2 function is as follows:

```
CREATE TYPE DBE_SQL.VARCHAR2_TABLE AS TABLE OF VARCHAR2(2000);
```

- DBE\_SQL.BIND\_VARIABLE

This function is used to bind parameters. You are advised to use **DBE\_SQL.SQL\_BIND\_VARIABLE**.

- DBE\_SQL.SQL\_SET\_RESULTS\_TYPE\_C

This function is used to dynamically define a column of the array type. You are advised not to use it.

The prototype of the DBE\_SQL.SQL\_SET\_RESULTS\_TYPE\_C function is as follows:

```
DBE_SQL.sql_set_results_type_c(
 context_id in int,
 pos in int,
 column_ref in anyarray,
 cnt in int,
 lower_bnd in int,
 col_type in anyelement,
 maxsize in int
)return integer;
```

**Table 10-188** DBE\_SQL.SQL\_SET\_RESULTS\_TYPE\_C parameters

| Parameter  | Description                                             |
|------------|---------------------------------------------------------|
| context_id | ID of the cursor to be queried.                         |
| pos        | Position of a dynamically defined column in the query   |
| column_ref | Type of the returned array.                             |
| cnt        | Number of values obtained at a time.                    |
| lower_bnd  | Start index when an array is returned.                  |
| col_type   | Variable type corresponding to the returned array type. |
| maxsize    | Maximum length of the defined type.                     |

- DBE\_SQL.SQL\_GET\_VALUES\_C

This function is used to read a dynamically defined column value. You are advised not to use it.

The prototype of the DBE\_SQL.SQL\_GET\_VALUES\_C function is as follows:

```
DBE_SQL.sql_get_values_c(
 context_id in int,
 pos in int,
 results_type inout anyarray,
 result_type in anyelement
)return anyarray;
```

**Table 10-189** DBE\_SQL.SQL\_GET\_VALUES\_C parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Parameter position.             |
| results_type | Obtained result.                |

| Parameter   | Description                  |
|-------------|------------------------------|
| result_type | Type of the obtained result. |

- DBE\_SQL.GET\_VARIABLE\_RESULT

This function is used to return the value of the bound OUT parameter and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_VARIABLE\_RESULT** function is as follows:

```
DBE_SQL.get_variable_result(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT column_value anyelement
);
```

**Table 10-190** DBE\_SQL.GET\_VARIABLE\_RESULT parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Name of the bound parameter.    |
| column_value | Return value                    |

- DBE\_SQL.GET\_VARIABLE\_RESULT\_CHAR

This function is used to return the value of the bound OUT parameter of the **CHAR** type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_VARIABLE\_RESULT\_CHAR** function is as follows:

```
DBE_SQL.get_variable_result_char(
 IN context_id int,
 IN pos VARCHAR2
)
RETURNS char
```

**Table 10-191** DBE\_SQL.GET\_VARIABLE\_RESULT\_CHAR parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be queried. |
| pos        | Name of the bound parameter.    |

- DBE\_SQL.GET\_VARIABLE\_RESULT\_RAW

This function is used to return the value of the bound OUT parameter of the **RAW** type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_VARIABLE\_RESULT\_RAW** function is as follows:

```
CREATE OR REPLACE FUNCTION DBE_SQL.get_variable_result_raw(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT value anyelement
```

```
)
RETURNS anyelement
```

**Table 10-192** DBE\_SQL.GET\_VARIABLE\_RESULT\_RAW parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be queried. |
| pos        | Name of the bound parameter.    |
| value      | Return value                    |

- DBE\_SQL.GET\_VARIABLE\_RESULT\_TEXT

This function is used to return the value of the bound OUT parameter of the **TEXT** type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_VARIABLE\_RESULT\_TEXT** function is as follows:

```
CREATE OR REPLACE FUNCTION DBE_SQL.get_variable_result_text(
 IN context_id int,
 IN pos VARCHAR2
)
RETURNS text
```

**Table 10-193** DBE\_SQL.GET\_VARIABLE\_RESULT\_TEXT parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be queried. |
| pos        | Name of the bound parameter.    |

- DBE\_SQL.GET\_VARIABLE\_RESULT\_INT

This function is used to return the value of the bound OUT parameter of the **INT** type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_VARIABLE\_RESULT\_INT** function is as follows:

```
DBE_SQL.get_variable_result_int(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT value anyelement
)
RETURNS anyelement
```

**Table 10-194** DBE\_SQL.GET\_VARIABLE\_RESULT\_INT parameters

| Parameter  | Description                     |
|------------|---------------------------------|
| context_id | ID of the cursor to be queried. |
| pos        | Name of the bound parameter.    |
| value      | Return value                    |

- DBE\_SQL.GET\_ARRAY\_RESULT\_TEXT

This function is used to return the value of the bound OUT parameter of the **TEXT** array type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_ARRAY\_RESULT\_TEXT** function is as follows:

```
DBE_SQL.get_array_result_text(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT column_value anyarray
)
```

**Table 10-195** DBE\_SQL.GET\_ARRAY\_RESULT\_TEXT parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Name of the bound parameter.    |
| column_value | Return value                    |

- DBE\_SQL.GET\_ARRAY\_RESULT\_RAW

This function is used to return the value of the bound OUT parameter of the **RAW** array type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_ARRAY\_RESULT\_RAW** function is as follows:

```
DBE_SQL.get_array_result_raw(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT column_value anyarray
)
```

**Table 10-196** DBE\_SQL.GET\_ARRAY\_RESULT\_RAW parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Name of the bound parameter.    |
| column_value | Return value                    |

- DBE\_SQL.GET\_ARRAY\_RESULT\_CHAR

This function is used to return the value of the bound OUT parameter of the **CHAR** array type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_ARRAY\_RESULT\_CHAR** function is as follows:

```
DBE_SQL.get_array_result_char(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT column_value anyarray
)
```

**Table 10-197** DBE\_SQL.GET\_ARRAY\_RESULT\_CHAR parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Name of the bound parameter.    |
| column_value | Return value                    |

- DBE\_SQL.GET\_ARRAY\_RESULT\_INT

This function is used to return the value of the bound OUT parameter of the **INT** array type and obtain the OUT parameter in a stored procedure.

The prototype of the **DBE\_SQL.GET\_ARRAY\_RESULT\_INT** function is as follows:

```
DBE_SQL.get_array_result_int(
 IN context_id int,
 IN pos VARCHAR2,
 INOUT column_value anyarray
)
```

**Table 10-198** DBE\_SQL.GET\_ARRAY\_RESULT\_INT parameters

| Parameter    | Description                     |
|--------------|---------------------------------|
| context_id   | ID of the cursor to be queried. |
| pos          | Name of the bound parameter.    |
| column_value | Return value                    |

## Examples

```
-- Perform operations on RAW data in a stored procedure.
openGauss=# create or replace procedure pro_dbe_sql_all_02(in_raw raw,v_in int,v_offset int)
as
context_id int;
v_id int;
v_info bytea :=1;
query varchar(2000);
execute_ret int;
define_column_ret_raw bytea :='1';
define_column_ret int;
begin
drop table if exists pro_dbe_sql_all_tb1_02 ;
create table pro_dbe_sql_all_tb1_02(a int ,b blob);
insert into pro_dbe_sql_all_tb1_02 values(1,HEXTORAW('DEADBEEE'));
insert into pro_dbe_sql_all_tb1_02 values(2,in_raw);
query := 'select * from pro_dbe_sql_all_tb1_02 order by 1';
-- Open a cursor.
context_id := dbe_sql.register_context();
-- Compile the cursor.
dbe_sql.sql_set_sql(context_id, query, 1);
-- Define columns.
define_column_ret:= dbe_sql.set_result_type(context_id,1,v_id);
define_column_ret_raw:= dbe_sql.set_result_type_raw(context_id,2,v_info,10);
-- Execute the cursor.
execute_ret := dbe_sql.sql_run(context_id);
loop
exit when (dbe_sql.next_row(cursorid) <= 0);
```

```
-- Obtain values.
dbe_sql.get_result(context_id,1,v_id);
dbe_sql.get_result_raw(context_id,2,v_info,v_in,v_offset);
-- Output the result.
dbe_output.print_line('id'|| v_id || ' info:' || v_info);
end loop;
-- Close the cursor.
dbe_sql.sql_unregister_context(context_id);
end;
/
-- Call the stored procedure.
openGauss=# call pro_dbe_sql_all_02(HEXTORAW('DEADBEEF'),0,1);

-- Drop the stored procedure.
openGauss=# DROP PROCEDURE pro_dbe_sql_all_02;
```

### 10.12.2.8 DBE\_FILE

The DBE\_FILE package provides the capabilities of reading and writing OS text files for stored procedures.

#### API Description

**Table 10-199** lists all APIs supported by the **DBE\_FILE** package.

**Table 10-199** DBE\_FILE

| API                          | Description                                                              |
|------------------------------|--------------------------------------------------------------------------|
| <b>DBE_FILE.OPEN</b>         | Opens a file based on the specified directory and file name.             |
| <b>DBE_FILE.IS_CLOSED</b>    | Checks whether a file is closed.                                         |
| <b>DBE_FILE.IS_OPEN</b>      | Checks whether a file is opened.                                         |
| <b>DBE_FILE.READ_LINE</b>    | Reads a line of a specified length from an open file handle.             |
| <b>DBE_FILE.WRITE</b>        | Writes data to an open file.                                             |
| <b>DBE_FILE.NEW_LINE</b>     | Writes one or more line terminators to an open file.                     |
| <b>DBE_FILE.WRITE_LINE</b>   | Writes data to an open file and automatically appends a line terminator. |
| <b>DBE_FILE.FORMAT_WRITE</b> | Writes data in a specified format to the buffer of an open file.         |
| <b>DBE_FILE.GET_RAW</b>      | Reads RAW data from an open file.                                        |
| <b>DBE_FILE.PUT_RAW</b>      | Writes RAW data to an open file.                                         |
| <b>DBE_FILE.FLUSH</b>        | Writes cached data to a physical file.                                   |
| <b>DBE_FILE.CLOSE</b>        | Closes an open file handle.                                              |



| API                                    | Description                                                                                                                      |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">DBE_FILE.CLOSE_ALL</a>     | Closes all file handles opened in a session.                                                                                     |
| <a href="#">DBE_FILE.REMOVE</a>        | Deletes a disk file. To perform this operation, you must have required permissions.                                              |
| <a href="#">DBE_FILE.RENAME</a>        | Renames a disk file, which is similar to the mv command of Unix.                                                                 |
| <a href="#">DBE_FILE.COPY</a>          | Copies data in a continuous area to a new file. If <b>start_line</b> and <b>end_line</b> are omitted, the entire file is copied. |
| <a href="#">DBE_FILE.GET_ATTR</a>      | Reads and returns the attributes of a disk file.                                                                                 |
| <a href="#">DBE_FILE.SEEK</a>          | Adjusts the position of a file pointer forward or backward based on the specified number of bytes.                               |
| <a href="#">DBE_FILE.GET_POSITIONS</a> | Returns the current offset of a file, in bytes.                                                                                  |

- DBE\_FILE.OPEN

This function opens a file. You can specify the maximum number of characters in each line. A maximum of 50 files can be opened at a time. This function returns a handle of the INTEGER type.

The prototype of the DBE\_FILE.OPEN function is as follows:

```
DBE_FILE.OPEN (
 dir IN VARCHAR2,
 file_name IN VARCHAR2,
 open_mode IN VARCHAR2,
 max_line_size IN INTEGER DEFAULT 1024)
RETURN INTEGER;
```

**Table 10-200** DBE\_FILE.OPEN API parameters

| Parameter     | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|----------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dir           | VARCHAR2 | IN                     | No           | Directory of a file. It is a string, indicating an object name.<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| file_name     | VARCHAR2 | IN                     | No           | File name with an extension (file type), excluding the path name. A path contained in a file name is ignored in the OPEN function. In Unix, the file name cannot end with the combination of a slash and a dot (/.).                                                                                                                                                                                                                                                                                                                                                                                                 |
| open_mode     | VARCHAR2 | IN                     | No           | File opening mode, including <b>r</b> (read), <b>w</b> (write), and <b>a</b> (append).<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits.                                                                                                                                                                                                                                                                                                                                                                            |
| max_line_size | INTEGER  | IN                     | Yes          | Maximum number of characters in each line, including newline characters. The minimum value is <b>1</b> and the maximum is <b>32767</b> . If this parameter is not specified, the default value <b>1024</b> is used.                                                                                                                                                                                                                                                                                                                                                                                                  |

- DBE\_FILE.IS\_CLOSE

Checks whether a file handle is closed. A Boolean value is returned. If an invalid file handle is detected, the **INVALID\_FILEHANDLE** exception is thrown. The prototype of the DBE\_FILE.IS\_CLOSE function is as follows:

```
DBE_FILE.IS_CLOSE (
file IN INTEGER)
RETURN BOOLEAN;
```

**Table 10-201** DBE\_FILE.IS\_CLOSE parameters

| Parameter          | Type                | Input/<br>Output<br>Parameter | Can<br>Be<br>Empty | Description                          |
|--------------------|---------------------|-------------------------------|--------------------|--------------------------------------|
| file IN<br>INTEGER | IN<br>TE<br>GE<br>R | IN                            | Yes                | Passes a file handle to be detected. |

- DBE\_FILE.READ\_LINE

Reads data from an open file and stores the read result to the buffer. It reads data to the end of each line excluding the line terminator, to the end of the file, or to the size specified by the **len** parameter. The length of the data to be read cannot exceed the value of **max\_line\_size** specified when the file is opened.

The prototype of the DBE\_FILE.READ\_LINE function is as follows:

```
DBE_FILE.READ_LINE (
file IN INTEGER,
buffer OUT VARCHAR2,
len IN INTEGER DEFAULT NULL)
```

**Table 10-202** DBE\_FILE.READ\_LINE parameters

| Parameter | Type                | Input/<br>Output<br>Parameter | Can<br>Be<br>Empty | Description                                                                                                                                  |
|-----------|---------------------|-------------------------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| file      | IN<br>TE<br>GE<br>R | IN                            | No                 | File handle opened by calling the OPEN function. The file must be opened in read mode. Otherwise, the INVALID_OPERATION exception is thrown. |

| Parameter | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                           |
|-----------|----------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffer    | VARCHAR2 | OUT                    | No           | Buffer used to receive data.                                                                                                                                          |
| len       | INTEGER  | IN                     | Yes          | Number of bytes read from a file. The default value is <b>NULL</b> . If the default value <b>NULL</b> is used, <b>max_line_size</b> is used to specify the line size. |

- DBE\_FILE.WRITE

Writes buffer data to the buffer corresponding to a file. The file must be opened in write mode. This operation does not write a line terminator.

The prototype of the **DBE\_FILE.WRITE** function is as follows:

```
DBE_FILE.WRITE (
file IN INTEGER,
buffer IN TEXT)
RETURN BOOLEAN;;
```

**Table 10-203** DBE\_FILE.WRITE parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                         |
|-----------|---------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------|
| file      | INTEGER | IN                     | No           | This stored procedure writes data in the buffer to a file. The file must be opened in write mode. Line terminators are not written. |

| Parameter | Type | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-----------|------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffer    | TEXT | IN                     | Yes          | <p>Text data to be written to a file. The maximum buffer size is 32767 bytes. If no value is specified in the open state, the default value is 1024 bytes. Before the writing is performed, the buffer occupied by WRITE operations cannot exceed 32767 bytes.</p> <p><b>NOTE</b></p> <p>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits.</p> |

- DBE\_FILE.NEW\_LINE

Writes one or more line terminators to the buffer corresponding to a file. The line terminators are related to the platform used.

The prototype of the **DBE\_FILE.NEW\_LINE** function is as follows:

```
DBE_FILE.NEW_LINE (
file IN INTEGER,
line_nums IN INTEGER := 1)
RETURN BOOLEAN;
```

**Table 10-204** DBE\_FILE.NEW\_LINE parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                      |
|-----------|---------|------------------------|--------------|------------------------------------------------------------------|
| line_nums | INTEGER | IN                     | Yes          | Number of terminators written to a file. The default value is 1. |

- DBE\_FILE.WRITE\_LINE

Writes buffer data to the buffer corresponding to a file. The file must be opened in write mode. This operation automatically adds a line terminator.

The prototype of the **DBE\_FILE.WRITE\_LINE** function is as follows:

```
DBE_FILE.WRITE_LINE(
file IN INTEGER,
buffer IN TEXT,
flush IN BOOLEAN DEFAULT FALSE)
RETURN BOOLEAN;
```

**Table 10-205** DBE\_FILE.WRITE\_LINE parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|---------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffer    | TEXT    | IN                     | Yes          | <p>Text data to be written to a file. The maximum buffer size is 32767 bytes. If no value is specified during the open state, the default value is 1024 bytes. Before the writing is performed, the buffer occupied by PUT operations cannot exceed 32767 bytes.</p> <p><b>NOTE</b><br/>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits.</p> |
| flush     | BOOLEAN | IN                     | Yes          | <p>Specifies whether to flush data to the disk after the writing.</p>                                                                                                                                                                                                                                                                                                                                                        |

- DBE\_FILE.FORMAT\_WRITE

Writes formatted data to the buffer corresponding to an open file. It is a DBE\_FILE.WRITE API that allows formatting.

The prototype of the **DBE\_FILE.FORMAT\_WRITE** function is as follows:

```
DBE_FILE.FORMAT_WRITE (
file IN INTEGER,
format IN VARCHAR2,
arg1 IN VARCHAR2 DEFAULT NULL,
...
arg6 IN VARCHAR2 DEFAULT NULL))
RETURN BOOLEAN;
```

**Table 10-206** DBE\_FILE.FORMAT\_WRITE parameters

| Parameter     | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                     |
|---------------|----------|------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file          | INTEGER  | IN                     | No           | File handle opened using OPEN.                                                                                                                                                                                                                  |
| format        | VARCHAR2 | IN                     | Yes          | A string to be formatted, containing the text and format characters \n and %s.                                                                                                                                                                  |
| [arg1...arg6] | VARCHAR2 | IN                     | Yes          | Six optional parameters. The parameters and the positions of characters to be formatted are in one-to-one correspondence. If the parameter corresponding to a character to be formatted is not provided, an empty string is used to replace %s. |

- DBE\_FILE.GET\_RAW

Reads RAW data from an open file, stores the read result in the buffer, and returns the result from *r*.

The prototype of the DBE\_FILE.GET\_RAW function is as follows:

```
DBE_FILE.GET_RAW (
file IN INTEGER,
r OUT RAW,
length IN INTEGER DEFAULT NULL);
```



**Table 10-207** DBE\_FILE.GET\_RAW parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                 |
|-----------|---------|------------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN.                                                                                              |
| r         | RAW     | OUTPUT                 | No           | Output binary data.                                                                                                         |
| length    | INTEGER | IN                     | Yes          | Length of the file to be read. The default value is <b>NULL</b> . All data in the file is read. The maximum length is 1 GB. |

- DBE\_FILE.PUT\_RAW

Writes raw data to the buffer corresponding to a file.

The prototype of the DBE\_FILE.PUT\_RAW function is as follows:

```
DBE_FILE.PUT_RAW (
file IN INTEGER,
r IN RAW,
flush IN BOOLEAN DEFAULT FALSE)
RETURN BOOLEAN;
```

**Table 10-208** DBE\_FILE.PUT\_RAW parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                            |
|-----------|---------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN.                                                                                                                                         |
| r         | RAW     | IN                     | No           | Output binary data.<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits. |
| flush     | BOOLEAN | IN                     | Yes          | Specifies whether to flush data to a file. The default value is <b>false</b> .                                                                                         |

- DBE\_FILE.FLUSH

Writes cached data to a physical file. The cached data must have a line terminator.

The prototype of the DBE\_FILE.FLUSH function is as follows:

```
DBE_FILE.FLUSH (
file IN INTEGER)
RETURN VOID;
```

**Table 10-209** DBE\_FILE.FLUSH parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

- DBE\_FILE.CLOSE

Closes an open file. When this function is called, if there is cached data to be written, exception information may be received.

The prototype of the DBE\_FILE.CLOSE function is as follows:

```
DBE_FILE.CLOSE (
file IN INTEGER
)RETURN INTEGER;
```

**Table 10-210** DBE\_FILE.CLOSE parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

- DBE\_FILE.CLOSE\_ALL

Closes all file handles opened in a session. This function can be used for emergency cleanup.

The prototype of the DBE\_FILE.CLOSE\_ALL function is as follows:

```
DBE_FILE.CLOSE_ALL()
RETRUN VOID;
```

**Table 10-211** DBE\_FILE.CLOSE\_ALL API parameters

| Parameter | Description |
|-----------|-------------|
| None.     | None.       |

- DBE\_FILE.REMOVE

Deletes a disk file. To use this function, you must have the required permission.

The prototype of the DBE\_FILE.REMOVE function is as follows:

```
DBE_FILE.REMOVE (
 dir IN VARCHAR2,
 file_name IN VARCHAR2)
RETURN VOID;
```

**Table 10-212** DBE\_FILE.REMOVE parameters

| Parameter | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------|----------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dir       | VARCHAR2 | IN                     | No           | Directory of a file. It is a string, indicating an object name.<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| file_name | VARCHAR2 | IN                     | No           | File name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

- DBE\_FILE.RENAME

Renames a disk file. This function is similar to the mv command of Unix.

The prototype of the DBE\_FILE.RENAME function is as follows:

```
DBE_FILE.RENAME (
src_dir IN VARCHAR2,
src_file_name IN VARCHAR2,
dest_dir IN VARCHAR2,
dest_file_name IN VARCHAR2,
overwrite IN BOOLEAN DEFAULT FALSE)
RETURN VOID;
```

**Table 10-213** DBE\_FILE.RENAME parameters

| Parameter     | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|----------|------------------------|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| src_dir       | VARCHAR2 | IN                     | No           | Directory of the original file (case-sensitive)<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| src_file_name | VARCHAR2 | IN                     | No           | Original file to be renamed                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| dest_dir      | VARCHAR2 | IN                     | No           | Target directory (case-sensitive).<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul>              |

| Parameter      | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                |
|----------------|----------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dest_file_name | VARCHAR2 | IN                     | No           | New file name                                                                                                                                                                                                              |
| overwrite      | BOOLEAN  | IN                     | Yes          | Specifies whether to overwrite the file. If the parameter is left empty or not specified, the parameter is not overwritten. If a file with the same name exists in the target directory, the file will not be overwritten. |

- DBE\_FILE.COPY

Copies data in a continuous area to a new file. If **start\_line** and **end\_line** are omitted, the entire file is copied.

The prototype of the DBE\_FILE.COPY function is as follows:

```
DBE_FILE.COPY (
src_dir IN VARCHAR2,
src_file_name IN VARCHAR2,
dest_dir IN VARCHAR2,
dest_file_name IN VARCHAR2,
start_line IN INTEGER DEFAULT 1,
end_line IN INTEGER DEFAULT NULL)
RETURN VOID;
```

**Table 10-214** DBE\_FILE.COPY parameters

| Parameter      | Type     | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------|----------|------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| src_dir        | VARCHAR2 | IN                     | No           | Directory of the original file<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul>    |
| src_file_name  | VARCHAR2 | IN                     | No           | File to be copied.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| dest_dir       | VARCHAR2 | IN                     | No           | Directory of the destination file<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <b>PG_DIRECTORY</b>. If the input path does not match the path in <b>PG_DIRECTORY</b>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| dest_file_name | VARCHAR2 | IN                     | No           | Target file<br><b>NOTE</b><br>For the write operation, the system checks the file type. If the ELF file is written, an error is reported and the system exits.                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Parameter  | Type                | Input/<br>Output<br>Parameter | Can<br>Be<br>Empty | Description                                                                                                      |
|------------|---------------------|-------------------------------|--------------------|------------------------------------------------------------------------------------------------------------------|
| start_line | IN<br>TE<br>GE<br>R | IN                            | No                 | Number of the line where the copy starts.<br>The default value is <b>1</b> .                                     |
| end_line   | IN<br>TE<br>GE<br>R | IN                            | Yes                | Number of the line where the copy ends.<br>The default value is <b>NULL</b> , indicating the<br>end of the file. |

- DBE\_FILE.GET\_ATTR

Reads and returns the attributes of a disk file.

The prototype of the DBE\_FILE.GET\_ATTR stored procedure is as follows:

```
DBE_FILE.GET_ATTR(
location IN text,
filename IN text,
OUT fexists boolean,
OUT file_length bigint,
OUT block_size integer);
```



**Table 10-215** DBE\_FILE.GET\_ATTR parameters

| Parameter   | Type    | Input/Output Parameter | Can Be Empty | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-------------|---------|------------------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| location    | TEXT    | IN                     | No           | File directory<br><b>NOTE</b> <ul style="list-style-type: none"><li>File directories need to be added to the system catalog <a href="#">PG_DIRECTORY</a>. If the input path does not match the path in <a href="#">PG_DIRECTORY</a>, an error indicating that the path does not exist will be reported. Functions that involve <b>location</b> as parameters also comply with this rule.</li><li>When the GUC parameter <b>safe_data_path</b> is enabled, you can only use an advanced package to read and write files in the file path specified by <b>safe_data_path</b>.</li></ul> |
| filename    | TEXT    | IN                     | No           | File name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| fexists     | BOOLEAN | OUT                    | No           | Specifies whether the file exists.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| file_length | BIGINT  | OUT                    | No           | File length (unit: byte). If the file does not exist, <b>NULL</b> is returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| block_size  | INTEGER | OUT                    | No           | Block size of the file system (unit: byte). If the file does not exist, <b>NULL</b> is returned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- DBE\_FILE.SEEK

Adjusts the position of a file pointer forward or backward based on the specified number of bytes.

The prototype of the DBE\_FILE.SEEK function is as follows:

```
DBE_FILE.SEEK (
file IN INTEGER,
absolute_start IN BIGINT DEFAULT NULL,
```

```
relative_start IN BIGINT DEFAULT NULL)
RETURN VOID;
```

**Table 10-216** DBE\_FILE.SEEK parameters

| Parameter      | Type    | Input/<br>Output<br>Parameter | Can<br>Be<br>Empty | Description                                                                                                                                                                                                                                                            |
|----------------|---------|-------------------------------|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| file           | INTEGER | IN                            | No                 | File handle opened using OPEN.                                                                                                                                                                                                                                         |
| absolute_start | BIGINT  | IN                            | Yes                | Absolute offset of a file. The default value is <b>NULL</b> .                                                                                                                                                                                                          |
| relative_start | BIGINT  | IN                            | Yes                | Relative offset of a file. A positive number indicates forward offset and a negative number indicates backward offset. The default value is <b>NULL</b> . If both <b>absolute_start</b> and this parameter are specified, the <b>absolute_start</b> parameter is used. |

- DBE\_FILE.GET\_POS

Returns the current offset of the file in bytes.

The prototype of the DBE\_FILE.FGETPOS function is as follows:

```
DBE_FILE.GET_POS (
file IN INTEGER)
RETURN BIGINT;
```

**Table 10-217** DBE\_FILE.GET\_POS parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

- DBE\_FILE.IS\_OPEN

Returns the current offset of the file in bytes.

The prototype of the DBE\_FILE.IS\_OPEN function is as follows:

```
DBE_FILE.IS_OPEN (
file IN INTEGER)
RETURN BOOLEAN;
```

**Table 10-218** DBE\_FILE.IS\_OPEN parameters

| Parameter | Type    | Input/Output Parameter | Can Be Empty | Description                    |
|-----------|---------|------------------------|--------------|--------------------------------|
| file      | INTEGER | IN                     | No           | File handle opened using OPEN. |

## Examples

```
-- Add the /temp/ directory to the PG_DIRECTORY system catalog as a SYSADMIN.
CREATE OR REPLACE DIRECTORY dir AS '/tmp/';
-- The expected result is as follows:
CREATE DIRECTORY
-- Open a file and write data into the file.
```

```
DECLARE
 f integer;
 dir text := 'dir';
BEGIN
 f := dbe_file.open(dir, 'sample.txt', 'w');
 PERFORM dbe_file.write_line(f, 'ABC');
 PERFORM dbe_file.write_line(f, '123'::numeric);
 PERFORM dbe_file.write_line(f, '-----');
 PERFORM dbe_file.new_line(f);
 PERFORM dbe_file.write_line(f, '*****');
 PERFORM dbe_file.new_line(f, 0);
 PERFORM dbe_file.write_line(f, '+++++++');
 PERFORM dbe_file.new_line(f, 2);
 PERFORM dbe_file.write_line(f, '#####');
 PERFORM dbe_file.write(f, 'A');
 PERFORM dbe_file.write(f, 'B');
 PERFORM dbe_file.new_line(f);
 PERFORM dbe_file.format_write(f, '[1 -> %s, 2 -> %s, 3 -> %s, 4 -> %s, 5 -> %s]', 'gaussdb', 'dbe', 'file',
'get', 'line');
 PERFORM dbe_file.new_line(f);
 PERFORM dbe_file.write_line(f, '1234567890');
 f := dbe_file.close(f);
END;
/
-- The expected result is as follows:
ANONYMOUS BLOCK EXECUTE

-- Read data from the file mentioned above.
DECLARE
 f integer;
 dir text := 'dir';
BEGIN
 f := dbe_file.open(dir, 'sample.txt', 'r');
 FOR i IN 1..11 LOOP
 RAISE INFO '[%] : %', i, dbe_file.read_line(f);
 END LOOP;
END;
/
-- The expected result is as follows:
INFO: [1] : ABC
INFO: [2] : 123
INFO: [3] : -----
INFO: [4] :
INFO: [5] : *****
INFO: [6] : ++++++++
INFO: [7] :
INFO: [8] :
INFO: [9] : #####
INFO: [10] : AB
INFO: [11] : [1 -> gaussdb, 2 -> dbe, 3 -> file, 4 -> get, 5 -> line]
ANONYMOUS BLOCK EXECUTE

-- Offset the file handle and obtain the current file location.
DECLARE
 l_file integer;
 l_buffer VARCHAR2(32767);
 dir text := 'dir';
 abs_offset number := 100;
 rel_offset number := NULL;
BEGIN
 l_file := dbe_file.open(dir => dir, file_name => 'sample.txt', open_mode => 'R');
 dbe_output.print_line('before seek: current position is ' || dbe_file.get_pos(file => l_file)); -- before seek:
current position is 0
 dbe_file.seek(file => l_file, absolute_start=>abs_offset, relative_start=>rel_offset);
 dbe_output.print_line('fseek: current position is ' || dbe_file.get_pos(file => l_file)); -- seek: current
position is 100
 l_file := dbe_file.close(file => l_file);
END;
/
```

```
-- The expected result is as follows:
before seek: current position is 0
fseek: current position is 100
ANONYMOUS BLOCK EXECUTE
```

### 10.12.2.9 DBE\_UTILITY

#### API Description

**Table 10-219** provides all APIs supported by the **DBE\_UTILITY** package.

**Table 10-219** DBE\_UTILITY

| API                                       | Description                                                               |
|-------------------------------------------|---------------------------------------------------------------------------|
| <b>DBE_UTILITY.FORMAT_ERROR_BACKTRACE</b> | Outputs the call stack of an abnormal stored procedure.                   |
| <b>DBE_UTILITY.FORMAT_ERROR_STACK</b>     | Outputs detailed information about a stored procedure exception.          |
| <b>DBE_UTILITY.FORMAT_CALL_STACK</b>      | Outputs the call stack of a stored procedure.                             |
| <b>DBE_UTILITY.GET_TIME</b>               | Outputs the current time, which is used to obtain the execution duration. |

- **DBE\_UTILITY.FORMAT\_ERROR\_BACKTRACE**

Returns the call stack where an error occurs during execution. The prototype of the **DBE\_UTILITY.FORMAT\_ERROR\_BACKTRACE** function is as follows:

```
DBE_UTILITY.FORMAT_ERROR_BACKTRACE()
RETURN TEXT;
```

- **DBE\_UTILITY.FORMAT\_ERROR\_STACK**

Returns the detailed information about the error location when an error occurs during the execution. The prototype of the **DBE\_UTILITY.FORMAT\_ERROR\_STACK** function is as follows:

```
DBE_UTILITY.FORMAT_ERROR_STACK()
RETURN TEXT;
```

- **DBE\_UTILITY.FORMAT\_CALL\_STACK**

Sets the call stack of the output function. The prototype of the **DBE\_UTILITY.FORMAT\_CALL\_STACK** function is as follows:

```
DBE_UTILITY.FORMAT_CALL_STACK()
RETURN TEXT;
```

- **DBE\_UTILITY.GET\_TIME**

Sets the output time, which is usually used for difference. A separate return value is meaningless. The prototype of the **DBE\_UTILITY.GET\_TIME** function is as follows:

```
DBE_UTILITY.GET_TIME()
RETURN BIGINT;
```

## Examples

```
CREATE OR REPLACE PROCEDURE test_get_time1()
AS
declare
 start_time bigint;
 end_time bigint;
BEGIN
 start_time:= dbe_utility.get_time ();
 pg_sleep(1);
 end_time:=dbe_utility.get_time ();
 dbe_output.print_line(end_time - start_time);
END;
/
```

### 10.12.2.10 DBE\_SESSION

#### API Description

**Table 10-220** provides all APIs supported by the **DBE\_SESSION** package. **DBE\_SESSION** takes effect at the session level.

**Table 10-220** DBE\_SESSION

| API                                        | Description                                               |
|--------------------------------------------|-----------------------------------------------------------|
| <a href="#">DBE_SESSION.SET_CONTEXT</a>    | Sets the value of an attribute in a specified context.    |
| <a href="#">DBE_SESSION.CLEAR_CONTEXT</a>  | Clears the value of an attribute in a specified context.  |
| <a href="#">DBE_SESSION.SEARCH_CONTEXT</a> | Queries the value of an attribute in a specified context. |

- **DBE\_SESSION.SET\_CONTEXT**

Sets the value of an attribute in a specified namespace (context). The prototype of the **DBE\_SESSION.SET\_CONTEXT** function is as follows:

```
DBE_SESSION.SET_CONTEXT(
 namespace text,
 attribute text,
 value text
)returns void;
```

**Table 10-221** DBE\_SESSION.SET\_CONTEXT API parameters

| Parameter | Description                                                                                                                     |
|-----------|---------------------------------------------------------------------------------------------------------------------------------|
| namespace | Name of the context to be set. If the context does not exist, create a context. The value contains a maximum of 128 characters. |
| attribute | Attribute name. The value contains a maximum of 128 characters.                                                                 |

| Parameter | Description                                                                  |
|-----------|------------------------------------------------------------------------------|
| value     | Name of the value to be set. The value contains a maximum of 128 characters. |

- DBE\_SESSION.CLEAR\_CONTEXT

Clears the value of an attribute in a specified namespace (context). The prototype of the DBE\_SESSION.CLEAR\_CONTEXT function is as follows:

```
DBE_SESSION.CLEAR_CONTEXT (
 namespace text,
 client_identifier text default 'null',
 attribute text
)returns void ;
```

**Table 10-222** DBE\_SESSION.CLEAR\_CONTEXT API parameters

| Parameter         | Description                                                                                                             |
|-------------------|-------------------------------------------------------------------------------------------------------------------------|
| namespace         | User-specified context.                                                                                                 |
| client_identifier | Client authentication. The default value is ' <b>null</b> '. Generally, you do not need to manually set this parameter. |
| attribute         | Attribute to be cleared.                                                                                                |

- DBE\_SESSION.SEARCH\_CONTEXT

Queries the value of an attribute in a specified namespace (context). The prototype of the DBE\_SESSION.SEARCH\_CONTEXT function is as follows:

```
DBE_SESSION.SEARCH_CONTEXT (
 namespace text,
 attribute text
)returns text;
```

**Table 10-223** DBE\_SESSION.SEARCH\_CONTEXT API parameters

| Parameter | Description              |
|-----------|--------------------------|
| namespace | User-specified context.  |
| attribute | Attribute to be queried. |

## Example

```
BEGIN
 select DBE_SESSION.set_context('test', 'gaussdb', 'one'); -- Set the gaussdb attribute of the test context to one.
 select DBE_SESSION.search_context('test', 'gaussdb');
 select DBE_SESSION.clear_context('test', 'test','gaussdb');
END;
```

### 10.12.2.11 DBE\_MATCH

#### API Description

[Table 10-224](#) provides all APIs supported by the **DBE\_MATCH** package.

**Table 10-224** DBE\_MATCH

| API                                | Description                                                                                                                                                                                                                                                                                                                            |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DBE_MATCH.EDIT_DISTANCE_SIMILARITY | Compares the difference between two character strings (minimum steps of deletion, addition, and conversion) and normalizes the difference to a value ranging from 0 to 100. The value <b>100</b> indicates that the two character strings are the same, and the value <b>0</b> indicates that the two character strings are different. |

- DBE\_MATCH.EDIT\_DISTANCE\_SIMILARITY

Compares the difference between two character strings (minimum steps of deletion, addition, and conversion) and normalizes the difference to a value ranging from 0 to 100. The value **100** indicates that the two character strings are the same, and the value **0** indicates that the two character strings are different. The prototype of the **DBE\_MATCH.EDIT\_DISTANCE\_SIMILARITY** function is as follows:

```
DBE_MATCH.EDIT_DISTANCE_SIMILARITY(
 str1 IN text,
 str2 IN text
)returns integer ;
```

**Table 10-225** DBE\_MATCH.EDIT\_DISTANCE\_SIMILARITY API parameters

| Parameter | Description                                                                  |
|-----------|------------------------------------------------------------------------------|
| str1      | First character string. If the value is <b>null</b> , <b>0</b> is returned.  |
| str2      | Second character string. If the value is <b>null</b> , <b>0</b> is returned. |

### 10.12.2.12 DBE\_APPLICATION\_INFO

#### API Description

[Table 10-226](#) provides all APIs supported by the **DBE\_APPLICATION\_INFO** package. **DBE\_APPLICATION\_INFO** applies to the current session.



**Table 10-226** DBE\_APPLICATION\_INFO

| API                                   | Description                |
|---------------------------------------|----------------------------|
| DBE_APPLICATION_INFO.SET_CLIENT_INFO  | Writes client information. |
| DBE_APPLICATION_INFO.READ_CLIENT_INFO | Reads client information.  |

- DBE\_APPLICATION\_INFO.SET\_CLIENT\_INFO  
Writes client information. The prototype of the **DBE\_APPLICATION\_INFO.SET\_CLIENT\_INFO** function is as follows:

```
DBE_APPLICATION_INFO.SET_CLIENT_INFO(
 str text
)returns void;
```

**Table 10-227** DBE\_APPLICATION\_INFO.SET\_CLIENT\_INFO API parameters

| Parameter | Description                |
|-----------|----------------------------|
| str       | Writes client information. |

- DBE\_APPLICATION\_INFO.READ\_CLIENT\_INFO  
The prototype of the **DBE\_APPLICATION\_INFO.READ\_CLIENT\_INFO** function is as follows:

```
DBE_APPLICATION_INFO.READ_CLIENT_INFO(
 OUT client_info text);
```

**Table 10-228** DBE\_APPLICATION\_INFO.READ\_CLIENT\_INFO API parameters

| Parameter   | Description        |
|-------------|--------------------|
| client_info | Client information |

## 10.13 Retry Management

Retry is a process in which the database executes an SQL statement or stored procedure (including anonymous block) again in the case of execution failure, improving the execution success rate and user experience. The database checks the error code and retry configuration to determine whether to retry.

- If the execution fails, the system rolls back the executed statements and executes the stored procedure again.

Example:

```
openGauss=# CREATE OR REPLACE PROCEDURE retry_basic (IN x INT)
AS
BEGIN
 INSERT INTO t1 (a) VALUES (x);
 INSERT INTO t1 (a) VALUES (x+1);
```

```
END;
/
openGauss=# CALL retry_basic(1);
```

## 10.14 Debugging

### Syntax

#### RAISE

The syntax of RAISE is as follows:

**Figure 10-35** raise\_format::=




**Figure 10-36** raise\_condition::=




**Figure 10-37** raise\_sqlstate::=




**Figure 10-38** raise\_option::=





**Figure 10-39** raise::=**Parameter description:**

- The level option is used to specify the error level, that is, **DEBUG**, **LOG**, **INFO**, **NOTICE**, **WARNING**, or **EXCEPTION** (default). **EXCEPTION** throws an error that normally terminates the current transaction and the others only generate information at their levels. The parameters **log\_min\_messages** and **client\_min\_messages** determine whether the error messages of specific levels are reported to the client and are written to the server log.
- **format**: specifies the error message text to be reported, a format string. The format string can be appended with an expression for insertion to the message text. In a format string, **%** is replaced by the parameter value attached to format and **%%** is used to print **%**. For example:  
--v\_job\_id replaces % in the string.  
`RAISE NOTICE 'Calling cs_create_job(%)',v_job_id;`
- **option = expression**: inserts additional information to an error report. The keyword option can be **MESSAGE**, **DETAIL**, **HINT**, or **ERRCODE**, and each expression can be any string.
  - **MESSAGE**: specifies the error message text. This option cannot be used in a **RAISE** statement that contains a format character string in front of **USING**.
  - **DETAIL**: specifies detailed information of an error.
  - **HINT**: prints hint information.
  - **ERRCODE**: designates an error code (SQLSTATE) to a report. A condition name or a five-character SQLSTATE error code can be used.
- **condition\_name**: specifies the condition name corresponding to the error code.
- **sqlstate**: specifies the error code.

If neither a condition name nor an **SQLSTATE** is designated in a **RAISE EXCEPTION** command, the **RAISE EXCEPTION (P0001)** is used by default. If no message text is designated, the condition name or SQLSTATE is used as the message text by default.

**NOTICE**


- If the **SQLSTATE** designates an error code, the error code is not limited to a defined error code. It can be any error code containing five digits or ASCII uppercase rather than **00000**. Do not use an error code ended with three zeros because such error codes are category codes and can be captured by the whole category.
- In O-compatible mode, **SQLCODE** is equivalent to **SQLSTATE**.

**NOTE**

The syntax described in [Figure 10-39](#) does not append any parameter. This form is used only for the **EXCEPTION** statement in a **BEGIN** block so that the error can be re-processed.

**EXCEPTION\_INIT**

In O-compatible mode, **EXCEPTION\_INIT** can be used to define the **SQLCODE** error code. The syntax is as follows:

**Figure 10-40** exception\_init::=**Parameter description:**

- **exception\_name**: indicates the name of the exception declared by the user. The **EXCEPTION\_INIT** syntax must follow the declared exception.
- **sqlcode** is a customized SQL code, which must be a negative integer ranging from -2147483647 to -1.

**NOTICE**

When **EXCEPTION\_INIT** is used to customize an SQL code, **SQLSTATE** is equivalent to **SQLCODE**, and **SQLERRM** is in the format of **xxx: non-GaussDB Exception**. For example, if the customized SQL code is -1, **SQLSTATE** is -1 and **SQLERRM** is 1: **non-GaussDB Exception**.

**Example**

Display error and hint information when a transaction terminates:

```
CREATE OR REPLACE PROCEDURE proc_raise1(user_id in integer)
AS
BEGIN
RAISE EXCEPTION 'Noexistence ID --> %',user_id USING HINT = 'Please check your user ID';
END;
/

call proc_raise1(300011);

-- Execution result:
ERROR: Noexistence ID --> 300011
HINT: Please check your user ID
```

Two methods are available for setting **SQLSTATE**:

```
CREATE OR REPLACE PROCEDURE proc_raise2(user_id in integer)
AS
BEGIN
RAISE 'Duplicate user ID: %',user_id USING ERRCODE = 'unique_violation';
END;
/

\set VERBOSITY verbose
call proc_raise2(300011);

-- Execution result:
```

```
ERROR: Duplicate user ID: 300011
SQLSTATE: 23505
LOCATION: exec_stmt_raise, pl_exec.cpp:3482
```

If the main parameter is a condition name or **SQLSTATE**, the following applies:

```
RAISE division_by_zero;
```

```
RAISE SQLSTATE '22012';
```

For example:

```
CREATE OR REPLACE PROCEDURE division(div in integer, dividend in integer)
AS
DECLARE
res int;
BEGIN
IF dividend=0 THEN
RAISE division_by_zero;
RETURN;
ELSE
res := div/dividend;
RAISE INFO 'division result: %', res;
RETURN;
END IF;
END;
/
call division(3,0);

-- Execution result:
ERROR: division_by_zero
```

Alternatively:

```
RAISE unique_violation USING MESSAGE = 'Duplicate user ID: ' || user_id;
```

In O-compatible mode, EXCEPTION\_INIT can be used to customize error codes SQLCODE.

```
declare
deadlock_detected exception;
pragma exception_init(deadlock_detected, -1);
begin
if 1 > 0 then
raise deadlock_detected;
end if;
exception
when deadlock_detected then
raise notice 'sqlcode:%,sqlstate:%,sqlerrm:%',sqlcode,sqlstate,sqlerrm;
end;
/
-- Execution result:
NOTICE: sqlcode:-1,sqlstate:-1,sqlerrm: 1: non-GaussDB Exception
```

## 10.15 Package

A package is a combination of PL/SQL programs, such as stored procedures, functions, variables, constants, and cursors. It is object-oriented and can encapsulate PL/SQL program design elements. Functions in a package are created, deleted, and modified in a unified manner.

A package contains two parts: package specifications and package body. The declaration contained in the package specifications can be accessed by external functions and anonymous blocks. The declaration contained in the package body cannot be accessed by external functions or anonymous blocks, but can be accessed only by functions and stored procedures in the package body.

For details about how to create a package, see [CREATE PACKAGE](#).

---

**NOTICE**

- Cross-package variables cannot be used as control variables in the for loops.
  - Types defined in a package cannot be deleted or modified, and cannot be used to define tables.
  - Cursor variables cannot be referenced in SCHEMA.PACKAGE.CUROS mode.
  - A cursor with parameters can be opened only in the current package.
-

# 11 Autonomous Transaction

---

An autonomous transaction is an independent transaction that is started during the execution of a main transaction. Committing and rolling back an autonomous transaction does not affect the data that has been committed by the main transaction. In addition, an autonomous transaction is not affected by the main transaction.

Autonomous transactions are defined in stored procedures, functions, anonymous blocks, and packages, and are declared using the **PRAGMA AUTONOMOUS\_TRANSACTION** keyword.

## 11.1 Stored Procedure Supporting Autonomous Transaction

An autonomous transaction can be defined in a stored procedure. The identifier of an autonomous transaction is **PRAGMA AUTONOMOUS\_TRANSACTION**. The syntax of an autonomous transaction is the same as that of creating a stored procedure. The following is an example.

```
-- Create a table.
create table t2(a int, b int);
insert into t2 values(1,2);
select * from t2;

-- Create a stored procedure that contains an autonomous transaction.
CREATE OR REPLACE PROCEDURE autonomous_4(a int, b int) AS
DECLARE
 num3 int := a;
 num4 int := b;
 PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 insert into t2 values(num3, num4);
 db_output.print_line('just use call. ');
END;
/

-- Create a common stored procedure that calls an autonomous transaction stored procedure.
CREATE OR REPLACE PROCEDURE autonomous_5(a int, b int) AS
DECLARE
BEGIN
 db_output.print_line('just no use call. ');
 insert into t2 values(666, 666);
 autonomous_4(a,b);
 rollback;
END;
```

```
/
-- Call a common stored procedure.
select autonomous_5(11,22);
-- View the table result.
select * from t2 order by a;
```

In the preceding example, a stored procedure containing an autonomous transaction is finally executed in a transaction block to be rolled back, which directly illustrates a characteristic of the autonomous transaction, that is, rollback of the main transaction does not affect content that has been committed by the autonomous transaction.

## 11.2 Anonymous Block Supporting Autonomous Transaction

An autonomous transaction can be defined in an anonymous block. The identifier of an autonomous transaction is **PRAGMA AUTONOMOUS\_TRANSACTION**. The syntax of an autonomous transaction is the same as that of creating an anonymous block. The following is an example.

```
create table t1(a int ,b text);

START TRANSACTION;
DECLARE
 PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 db_output.print_line('just use call. ');
 insert into t1 values(1,'you are so cute,will commit!');
END;
/
insert into t1 values(1,'you will rollback!');
rollback;

select * from t1;
```

In the preceding example, an anonymous block containing an autonomous transaction is finally executed before a transaction block to be rolled back, which directly illustrates a characteristic of the autonomous transaction, that is, rollback of the main transaction does not affect content that has been committed by the autonomous transaction.

## 11.3 Function Supporting Autonomous Transaction

An autonomous transaction can be defined in a function. The identifier of an autonomous transaction is **PRAGMA AUTONOMOUS\_TRANSACTION**. The syntax of an autonomous transaction is the same as that of creating a function. The following is an example.

```
create table t4(a int, b int, c text);

CREATE OR REPLACE function autonomous_32(a int ,b int ,c text) RETURN int AS
DECLARE
 PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 insert into t4 values(a, b, c);
 return 1;
END;
/
CREATE OR REPLACE function autonomous_33(num1 int) RETURN int AS
```



```
DECLARE
 num3 int := 220;
 tmp int;
 PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 num3 := num3/num1;
 return num3;
EXCEPTION
 WHEN division_by_zero THEN
 select autonomous_32(num3, num1, sqlerrm) into tmp;
 return 0;
END;
/

select autonomous_33(0);

select * from t4;
```

## 11.4 Package Supporting Autonomous Transaction

An autonomous transaction can be defined in a stored procedure or function in a package. The identifier of an autonomous transaction is **PRAGMA AUTONOMOUS\_TRANSACTION**. The syntax of an autonomous transaction is the same as that of creating a stored procedure or function in a package. The following is an example.

```
-- Create a table.
drop table t2;
create table t2(a int, b int);
insert into t2 values(1,2);
select * from t2;

-- Create a stored procedure or function in a package that contains autonomous transactions.
CREATE OR REPLACE PACKAGE autonomous_pkg AS
 PROCEDURE autonomous_4(a int, b int);
 FUNCTION autonomous_32(a int ,b int) RETURN int;
END autonomous_pkg;
/
CREATE OR REPLACE PACKAGE body autonomous_pkg AS
 PROCEDURE autonomous_4(a int, b int) AS
 DECLARE
 num3 int := a;
 num4 int := b;
 PRAGMA AUTONOMOUS_TRANSACTION;
 BEGIN
 insert into t2 values(num3, num4);
 END;
 FUNCTION autonomous_32(a int ,b int) RETURN int AS
 DECLARE
 PRAGMA AUTONOMOUS_TRANSACTION;
 BEGIN
 insert into t2 values(a, b);
 return 1;
 END;
END autonomous_pkg;
/

-- Create a common stored procedure that calls a stored procedure or function from a package that
contains autonomous transactions.
CREATE OR REPLACE PROCEDURE autonomous_5(a int, b int) AS
 DECLARE
 va int;
 BEGIN
 insert into t2 values(666, 666);
 autonomous_pkg.autonomous_4(a,b);
 va := autonomous_pkg.autonomous_32(a + 1, b + 1);
 rollback;
```

```
END;
/
-- Call a common stored procedure.
select autonomous_5(11,22);
-- View the table result.
select * from t2 order by a;
```

In the preceding example, a stored procedure or function in a package containing autonomous transactions is finally executed in a transaction block that is rolled back, which directly illustrates a characteristic of the autonomous transaction, that is, rollback of the main transaction does not affect content that has been committed by an autonomous transaction.

## 11.5 Restrictions

### CAUTION

- When an autonomous transaction is executed, an autonomous transaction session is started in the background. You can use **max\_concurrent\_autonomous\_transactions** to set the maximum number of concurrent autonomous transactions. The value range is 0 to 1024, and the default value is **10**.
- When **max\_concurrent\_autonomous\_transactions** is set to **0**, autonomous transactions cannot be executed.
- After a new session is started for an autonomous transaction, the default session parameters are used and objects (including session-level variables, local temporary variables, and global temporary table data) of the primary session are not shared.
- Autonomous transactions are affected by the communication buffer. The size of the information returned to the client is limited by the length of the communication buffer. If the size exceeds the length of the communication buffer, an error is reported.

- A trigger function does not support autonomous transactions.

```
CREATE TABLE test_trigger_des_tbl(id1 INT, id2 INT, id3 INT);

CREATE OR REPLACE FUNCTION tri_insert_func() RETURNS TRIGGER AS
$$
DECLARE
 PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 INSERT INTO test_trigger_des_tbl VALUES(NEW.id1, NEW.id2, NEW.id3);
 RETURN NEW;
END
$$ LANGUAGE plpgsql;
```

- Autonomous transactions cannot be called by non-top-layer anonymous blocks (but can only be called by top-layer autonomous transactions, including stored procedures, functions, and anonymous blocks).

```
create table t1(a int ,b text);

DECLARE
 --PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
 DECLARE
 PRAGMA AUTONOMOUS_TRANSACTION;
 BEGIN
```

```
dbe_output.print_line('just use call.');
```

```
insert into t1 values(1,'can you rollback!');
```

```
END;
```

```
insert into t1 values(2,'I will rollback!');
```

```
rollback;
```

```
END;
```

```
/
```

```
select * from t1;
```

- In an autonomous transaction, the **ref cursor** parameter can be passed only through the **PROCEDURE OUT** parameter. The **ref cursor** parameter cannot be passed through the **IN**, **INOUT**, or **FUNCTION** parameter.

```
create table sections(section_ID int);
```

```
insert into sections values(1);
```

```
insert into sections values(1);
```

```
insert into sections values(1);
```

```
insert into sections values(1);
```

1. The **PROCEDURE OUT** output parameter passes the **ref cursor** parameter (supported).

```
CREATE OR REPLACE PROCEDURE proc_sys_ref(OUT c1 refcursor)
```

```
IS
```

```
declare
```

```
 PRAGMA AUTONOMOUS_TRANSACTION;
```

```
BEGIN
```

```
 OPEN c1 FOR SELECT section_ID FROM sections ORDER BY section_ID;
```

```
END;
```

```
/
```

```
CREATE OR REPLACE PROCEDURE proc_sys_call() AS
```

```
DECLARE
```

```
 c1 SYS_REFCURSOR;
```

```
 TEMP NUMBER(4);
```

```
BEGIN
```

```
 proc_sys_ref(c1);
```

```
 if c1%isopen then
```

```
 raise notice '%','ok';
```

```
 end if;
```

```
LOOP
```

```
 FETCH C1 INTO TEMP;
```

```
 raise notice '%',C1%ROWCOUNT;
```

```
 EXIT WHEN C1%NOTFOUND;
```

```
END LOOP;
```

```
END;
```

```
/
```

2. The **PROCEDURE IN** or **INOUT** output parameter passes the **ref cursor** parameter (not supported).

```
CREATE OR REPLACE PROCEDURE proc_sys_ref(IN c1 refcursor)
```

```
IS
```

```
declare
```

```
 PRAGMA AUTONOMOUS_TRANSACTION;
```

```
 TEMP NUMBER(4);
```

```
BEGIN
```

```
 if c1%isopen then
```

```
 raise notice '%','ok';
```

```
 end if;
```

```
LOOP
```

```
 FETCH C1 INTO TEMP;
```

```
 raise notice '%',C1%ROWCOUNT;
```

```
 EXIT WHEN C1%NOTFOUND;
```

```
END LOOP;
```

```
END;
```

```
/
```

```
CREATE OR REPLACE PROCEDURE proc_sys_call() AS
```

```
DECLARE
```

```
c1 SYS_REFCURSOR;
TEMP NUMBER(4);
BEGIN
OPEN c1 FOR SELECT section_ID FROM sections ORDER BY section_ID;
proc_sys_ref(c1);
END;
/
```

3. **FUNCTION RETURN** passes the **ref cursor** parameter (not supported)  
CREATE OR REPLACE function proc\_sys\_ref()  
return SYS\_REFCURSOR

```
IS
declare
PRAGMA AUTONOMOUS_TRANSACTION;
C1 SYS_REFCURSOR;
BEGIN
OPEN C1 FOR SELECT section_ID FROM sections ORDER BY section_ID;
return C1;
END;
/
```

4. The **FUNCTION OUT** output parameter passes the **ref cursor** parameter (not supported).  
CREATE OR REPLACE function proc\_sys\_ref(C1 out SYS\_REFCURSOR)  
return SYS\_REFCURSOR

```
IS
declare
PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
OPEN C1 FOR SELECT section_ID FROM sections ORDER BY section_ID;
return 1;
END;
/
```

- Autonomous transaction functions only return records in the out format.

```
create table test_in (id int,a date);
CREATE OR REPLACE FUNCTION autonomous_out()
RETURNS record
LANGUAGE plpgsql AS $$
DECLARE PRAGMA AUTONOMOUS_TRANSACTION;
r1 record;
BEGIN
DBE_OUTPUT.PRINT_LINE('this is in autonomous_f_139_7');
truncate test_in;
insert into test_in values (1,'1909-01-01');
select * into r1 from test_in;
RETURN r1;
END;
$$;
select ok.id,ok.a from autonomous_f_139_7() as ok(id int,a date);
```

- The isolation level of an autonomous transaction cannot be changed.
- Autonomous transactions do not support the **setof** return type.

```
create table test_in (id int,a date);
create table test_main (id int,a date);
insert into test_main values (1111,'2021-01-01'),(2222,'2021-02-02');
truncate test_in,test_main;
CREATE OR REPLACE FUNCTION autonomous_f_022(num1 int) RETURNS SETOF test_in
LANGUAGE plpgsql AS $$
DECLARE
count int :=3;
test_row test_in%ROWTYPE;
PRAGMA AUTONOMOUS_TRANSACTION;
BEGIN
while true
loop
if count=3 then
null;
else
if count=2 then
insert into test_main values (count,'2021-03-03');
```

```
 goto pos1;
 end if;
end if;
count=count-1;
end loop;
insert into test_main values (1000,'2021-04-04');
<<pos1>>
for test_row in select * from test_main
loop
 return next test_row;
end loop;
return;
END;
$$
;
```

# 12 System Catalogs and System Views

---

## 12.1 Overview of System Catalogs and System Views

System catalogs store structured metadata of GaussDB. They are the source of information used by GaussDB to control system running and are a core component of the database system.

System views provide ways to query the system catalogs and internal database status.

System catalogs and system views are visible to either users with the SYSADMIN permission or all users. Some system catalogs and views have marked the need of administrator permissions, so they are accessible only to administrators.

You can delete and re-create system catalogs, add columns to them, and insert and update values in them, but doing so may make system information inconsistent and cause system faults. Generally, users should not modify system catalogs or system views, or rename their schemas. They are automatically maintained by the system.

---

### NOTICE

- You are advised not to modify the permissions on system catalogs or system views.
  - Do not add, delete, or modify system catalogs because doing so will result in exceptions or even database unavailability.
  - For details about column types in system catalogs and system views, see [Data Type](#).
- 

## 12.2 System Catalogs

## 12.2.1 GS\_ASP

GS\_ASP displays the persistent ACTIVE SESSION PROFILE samples. This catalog can be queried only in the system database but cannot be queried in the user database.

**Table 12-1** GS\_ASP columns

| Name              | Type                     | Description                                                                                                                                                      |
|-------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sampleid          | bigint                   | Sample ID                                                                                                                                                        |
| sample_time       | timestamp with time zone | Sampling time                                                                                                                                                    |
| need_flush_sample | Boolean                  | Specifies whether the sample needs to be flushed to disks. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul> |
| databaseid        | oid                      | Database ID                                                                                                                                                      |
| thread_id         | bigint                   | Thread ID                                                                                                                                                        |
| sessionid         | bigint                   | Session ID                                                                                                                                                       |
| start_time        | timestamp with time zone | Start time of a session                                                                                                                                          |
| event             | text                     | Event name                                                                                                                                                       |
| lwtid             | integer                  | Lightweight thread ID of the current thread                                                                                                                      |
| psessionid        | bigint                   | Parent thread of the streaming thread                                                                                                                            |
| tlevel            | integer                  | Level of the streaming thread, which corresponds to the level ( <b>id</b> ) of the execution plan                                                                |
| smpid             | integer                  | Concurrent thread ID in SMP execution mode                                                                                                                       |
| userid            | oid                      | ID of a session user                                                                                                                                             |
| application_name  | text                     | Name of an application                                                                                                                                           |
| client_addr       | inet                     | IP address of a client                                                                                                                                           |
| client_hostname   | text                     | Name of a client                                                                                                                                                 |
| client_port       | integer                  | TCP port number used by a client to communicate with the backend                                                                                                 |
| query_id          | bigint                   | Debug query ID                                                                                                                                                   |

| Name             | Type                     | Description                                                                                                                                                                                                                                                    |
|------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| unique_query_id  | bigint                   | Unique query ID                                                                                                                                                                                                                                                |
| user_id          | oid                      | User ID in the key of the unique query                                                                                                                                                                                                                         |
| cn_id            | integer                  | ID of the node that delivers the unique SQL statement. <b>cn_id</b> is in the key of the unique query.                                                                                                                                                         |
| unique_query     | text                     | Standardized unique SQL text string.                                                                                                                                                                                                                           |
| locktag          | text                     | Information of a lock that the session waits for, which can be parsed using <b>locktag_decode</b>                                                                                                                                                              |
| lockmode         | text                     | Mode of a lock that the session waits for <ul style="list-style-type: none"><li>• <b>LW_EXCLUSIVE</b>: exclusive lock</li><li>• <b>LW_SHARED</b>: shared lock</li><li>• <b>LW_WAIT_UNTIL_FREE</b>: waits for the <b>LW_EXCLUSIVE</b> to be available</li></ul> |
| block_sessionid  | bigint                   | Blocks a session from obtaining the session ID of a lock if the session is waiting for the lock.                                                                                                                                                               |
| wait_status      | text                     | Provides more details about an event column.                                                                                                                                                                                                                   |
| global_sessionid | text                     | Global session ID                                                                                                                                                                                                                                              |
| xact_start_time  | timestamp with time zone | Start time of the transaction                                                                                                                                                                                                                                  |
| query_start_time | timestamp with time zone | Time when the statement starts to be executed                                                                                                                                                                                                                  |
| state            | text                     | Current transaction state.<br>The value can be <b>active</b> , <b>idle in transaction</b> , <b>fastpath function call</b> , <b>idle in transaction (aborted)</b> , <b>disabled</b> , or <b>retrying</b> .                                                      |

### 12.2.2 GS\_AUDITING\_POLICY

GS\_AUDITING\_POLICY records the main information about the unified audit. Each record corresponds to a design policy. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.



**Table 12-2** GS\_AUDITING\_POLICY columns

| Name        | Type                        | Description                                                                                                                                            |
|-------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid         | oid                         | Row identifier (hidden attribute, which must be specified)                                                                                             |
| polname     | name                        | Policy name, which must be unique                                                                                                                      |
| polcomments | name                        | Policy description field, which records policy-related description information and is represented by the <b>COMMENTS</b> keyword                       |
| modifydate  | timestamp without time zone | Latest timestamp when a policy is created or modified                                                                                                  |
| polenabled  | boolean                     | Specifies whether to enable a policy. <ul style="list-style-type: none"><li>• <b>t</b> (true): enabled.</li><li>• <b>f</b> (false): disabled</li></ul> |

### 12.2.3 GS\_AUDITING\_POLICY\_ACCESS

GS\_AUDITING\_POLICY\_ACCESS records the unified audit information about DML database operations. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-3** GS\_AUDITING\_POLICY\_ACCESS columns

| Name       | Type                        | Description                                                                                                                             |
|------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| oid        | oid                         | Row identifier (hidden attribute, which must be specified).                                                                             |
| accesstype | name                        | DML database operation type. For example, SELECT, INSERT, and DELETE.                                                                   |
| labelname  | name                        | Resource label name. This column corresponds to the <b>polname</b> column in the system catalog in <a href="#">GS_AUDITING_POLICY</a> . |
| policyoid  | oid                         | OID of the audit policy, corresponding to the OID in the GS_AUDITING_POLICY system catalog.                                             |
| modifydate | timestamp without time zone | Latest creation or modification timestamp.                                                                                              |

## 12.2.4 GS\_AUDITING\_POLICY\_FILTERS

GS\_AUDITING\_POLICY\_FILTERS records the filtering policies about the unified audit. Each record corresponds to a design policy. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-4** GS\_AUDITING\_POLICY\_FILTERS columns

| Name            | Type                        | Description                                                                                             |
|-----------------|-----------------------------|---------------------------------------------------------------------------------------------------------|
| oid             | oid                         | Row identifier (hidden attribute, which must be specified).                                             |
| filtertype      | name                        | Filter type. Currently, the value is <b>logical_expr</b> .                                              |
| labelname       | name                        | Name. Currently, the value is <b>logical_expr</b> .                                                     |
| policyoid       | oid                         | OID of the audit policy, corresponding to OID in the <a href="#">GS_AUDITING_POLICY</a> system catalog. |
| modifydate      | timestamp without time zone | Latest creation or modification timestamp.                                                              |
| logicaloperator | text                        | Logical character string of a filter criterion.                                                         |

## 12.2.5 GS\_AUDITING\_POLICY\_PRIVILEGES

GS\_AUDITING\_POLICY\_PRIVILEGES records the DDL database operations about the unified audit. Each record corresponds to a design policy. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-5** GS\_AUDITING\_POLICY\_PRIVILEGES columns

| Name            | Type | Description                                                                                                                             |
|-----------------|------|-----------------------------------------------------------------------------------------------------------------------------------------|
| oid             | oid  | Row identifier (hidden attribute, which must be specified).                                                                             |
| privilegegetype | name | DDL database operation type. For example, CREATE, ALTER, and DROP.                                                                      |
| labelname       | name | Resource label name. This column corresponds to the <b>polname</b> column in the system catalog in <a href="#">GS_AUDITING_POLICY</a> . |

| Name       | Type                        | Description                                                                                  |
|------------|-----------------------------|----------------------------------------------------------------------------------------------|
| policyoid  | oid                         | This column corresponds to the OID in the <a href="#">GS_AUDITING_POLICY</a> system catalog. |
| modifydate | timestamp without time zone | Latest creation or modification timestamp.                                                   |

## 12.2.6 GS\_DB\_PRIVILEGE

**GS\_DB\_PRIVILEGE** records the granting of ANY permissions. Each record corresponds to a piece of authorization information.

**Table 12-6** GS\_DB\_PRIVILEGE columns

| Name           | Type    | Description                                                                                                                                                                                   |
|----------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid            | oid     | Row identifier (hidden attribute, which must be specified).                                                                                                                                   |
| roleid         | oid     | User ID.                                                                                                                                                                                      |
| privilege_type | text    | ANY permission of a user. For details about the value, see <a href="#">Table 7-108</a> .                                                                                                      |
| admin_option   | boolean | Specifies whether the ANY permission recorded in the <b>privilege_type</b> column can be re-granted. <ul style="list-style-type: none"><li>• <b>t</b>: yes.</li><li>• <b>f</b>: no.</li></ul> |

## 12.2.7 GS\_GLOBAL\_CONFIG

**GS\_GLOBAL\_CONFIG** records the parameter values specified by users during database instance initialization. In addition, it also stores weak passwords set by users. Initial database users can write, modify, and delete parameters in system catalogs by using **ALTER** and **DROP**.

**Table 12-7** GS\_GLOBAL\_CONFIG columns

| Name | Type | Description                                                                                                                      |
|------|------|----------------------------------------------------------------------------------------------------------------------------------|
| name | name | Specifies the preset parameter name, weak password name, or parameter required by users during database instance initialization. |

| Name  | Type | Description                                                                                                                              |
|-------|------|------------------------------------------------------------------------------------------------------------------------------------------|
| value | text | Specifies the preset parameter value, weak password value, or parameter value required by users during database instance initialization. |

## 12.2.8 GS\_JOB\_ARGUMENT

GS\_JOB\_ARGUMENT provides the parameter attributes of DBE\_SCHEDULER scheduled tasks and programs.

**Table 12-8** GS\_JOB\_ARGUMENT columns

| Name              | Type    | Description                                                                                                                                      |
|-------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| oid               | oid     | Row identifier (hidden column)                                                                                                                   |
| argument_position | integer | Location of a parameter of a scheduled task or program.                                                                                          |
| argument_type     | name    | Parameter type of a scheduled task or program.                                                                                                   |
| job_name          | text    | Name of a scheduled task or program.                                                                                                             |
| argument_name     | text    | Parameter name of a scheduled task or program. The scheduled task inherits the parameter name of the program. Therefore, this parameter is null. |
| argument_value    | text    | Parameter value of a scheduled task. (The program cannot bind a value.)                                                                          |
| default_value     | text    | Default parameter value of a program.                                                                                                            |

## 12.2.9 GS\_JOB\_ATTRIBUTE

GS\_JOB\_ATTRIBUTE records attributes of DBE\_SCHEDULER scheduled tasks, including basic attributes of scheduled tasks, scheduled task classes, certificates, authorization, programs, and schedules.

**Table 12-9** GS\_JOB\_ATTRIBUTE columns

| Name | Type | Description                    |
|------|------|--------------------------------|
| oid  | oid  | Row identifier (hidden column) |

| Name            | Type | Description                                                                                                                |
|-----------------|------|----------------------------------------------------------------------------------------------------------------------------|
| job_name        | text | Names of scheduled tasks, scheduled task classes, certificates, programs, and schedules, and authorized user names.        |
| attribute_name  | text | Attribute names of scheduled tasks, scheduled task classes, certificates, programs, and schedules, and authorized content. |
| attribute_value | text | Attribute values of scheduled tasks, scheduled task classes, certificates, programs, and schedules.                        |

## 12.2.10 GS\_MASKING\_POLICY

GS\_MASKING\_POLICY records the main information about dynamic data masking policies. Each record corresponds to a masking policy. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-10** GS\_MASKING\_POLICY columns

| Name        | Type                        | Description                                                                                                                                            |
|-------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid         | oid                         | Row identifier (hidden attribute, which must be specified).                                                                                            |
| polname     | name                        | Policy name, which must be unique                                                                                                                      |
| polcomments | name                        | Policy description field, which records policy-related description information and is represented by the <b>COMMENTS</b> keyword                       |
| modifydate  | timestamp without time zone | Latest timestamp when a policy is created or modified                                                                                                  |
| polenabled  | boolean                     | Specifies whether to enable a policy. <ul style="list-style-type: none"><li>• <b>t</b> (true): enabled.</li><li>• <b>f</b> (false): disabled</li></ul> |

## 12.2.11 GS\_MASKING\_POLICY\_ACTIONS

GS\_MASKING\_POLICY\_ACTIONS records the masking actions of a masking policy in the dynamic data masking policies. One masking policy corresponds to one or more rows of records in the catalog. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-11** GS\_MASKING\_POLICY\_ACTIONS columns

| Name          | Type                        | Description                                                                                                      |
|---------------|-----------------------------|------------------------------------------------------------------------------------------------------------------|
| oid           | oid                         | Row identifier (hidden attribute, which must be specified).                                                      |
| actiontype    | name                        | Name of a masking function used by a masking policy                                                              |
| actparams     | name                        | Parameter information specified in a masking function                                                            |
| actlabelname  | name                        | Name of a masked label                                                                                           |
| policyoid     | oid                         | OID of a masking policy to which a record belongs, corresponding to the OID in <a href="#">GS_MASKING_POLICY</a> |
| actmodifydate | timestamp without time zone | Latest timestamp when a record is created or modified                                                            |

## 12.2.12 GS\_MASKING\_POLICY\_FILTERS

**GS\_MASKING\_POLICY\_FILTERS** records the user filter criteria corresponding to the dynamic data masking policies. The corresponding masking policy takes effect only when the user information meets the filter criteria. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

**Table 12-12** GS\_MASKING\_POLICY\_FILTERS columns

| Name            | Type                        | Description                                                                                                                           |
|-----------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| oid             | oid                         | Row identifier (hidden attribute, which must be specified).                                                                           |
| filtertype      | name                        | Filter type. Currently, the value is <b>logical_expr</b> .                                                                            |
| filterlabelname | name                        | Filter range. Currently, the value is <b>logical_expr</b> .                                                                           |
| policyoid       | oid                         | OID of the masking policy to which a user filter criterion belongs, which corresponds to the OID in <a href="#">GS_MASKING_POLICY</a> |
| modifydate      | timestamp without time zone | Latest timestamp when a user filter criterion is created or modified                                                                  |
| logicaloperator | text                        | Polish notation of a filter criterion                                                                                                 |

## 12.2.13 GS\_MATVIEW

**GS\_MATVIEW** provides information about each materialized view in the database.

**Table 12-13** GS\_MATVIEW columns

| Name        | Type      | Description                                                                                                                                                                                                                                     |
|-------------|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid         | oid       | Row identifier (hidden attribute, which must be specified)                                                                                                                                                                                      |
| matviewid   | oid       | OID of a materialized view                                                                                                                                                                                                                      |
| mapid       | oid       | OID of a map table associated with a materialized view. Each map table corresponds to one materialized view. If a complete-refresh materialized view does not correspond to a map table, the value of this column is <b>0</b> .                 |
| ivm         | boolean   | Type of a materialized view. The value <b>t</b> indicates a fast-refresh materialized view, and the value <b>f</b> indicates a complete-refresh materialized view.                                                                              |
| needrefresh | boolean   | Reserved column                                                                                                                                                                                                                                 |
| refreshtime | timestamp | Last time when a materialized view was refreshed. If the materialized view is not refreshed, the value is null. This column is maintained only for fast-refresh materialized views. For complete-refresh materialized views, the value is null. |

## 12.2.14 GS\_MATVIEW\_DEPENDENCY

**GS\_MATVIEW\_DEPENDENCY** provides association information about each fast-refresh materialized view, base table, and Mlog table in the database. The Mlog table corresponding to the base table does not exist in the complete-refresh materialized view. Therefore, no record is written into the Mlog table.

**Table 12-14** GS\_MATVIEW\_DEPENDENCY columns

| Name      | Type | Description                                                |
|-----------|------|------------------------------------------------------------|
| oid       | oid  | Row identifier (hidden attribute, which must be specified) |
| matviewid | oid  | OID of a materialized view                                 |
| relid     | oid  | OID of a base table of a materialized view                 |

| Name   | Type | Description                                                                                                       |
|--------|------|-------------------------------------------------------------------------------------------------------------------|
| mlogid | oid  | OID of a Mlog table which is the log table of a materialized view. Each Mlog table corresponds to one base table. |
| mxmin  | int4 | Reserved column                                                                                                   |

## 12.2.15 GS\_OPT\_MODEL

GS\_OPT\_MODEL is a data table used when the AI engine is enabled to predict the planned time. It records the configurations, training results, features, corresponding system functions, and training history of machine learning models.

**Table 12-15** GS\_OPT\_MODEL columns

| Name          | Type    | Description                                                                                                                                                         |
|---------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid     | Database object ID.                                                                                                                                                 |
| template_name | name    | Template name of the machine learning model, which determines the APIs called for training and prediction. Currently, only rlstm is implemented.                    |
| model_name    | name    | Model name. Each model corresponds to a set of parameters, training logs, and model coefficients in the AI engine online learning process. The name must be unique. |
| datname       | name    | Name of the database served by the model. Each model is specific to a single database. This parameter determines data used for training.                            |
| ip            | name    | IP address of the host where the AI engine is deployed                                                                                                              |
| port          | integer | Listening port number of the AI engine                                                                                                                              |



| Name          | Type    | Description                                                                                                                                                       |
|---------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| max_epoch     | integer | Maximum number of iterations in an epoch                                                                                                                          |
| learning_rate | real    | Learning rate of model training. The default value <b>1</b> is recommended.                                                                                       |
| dim_red       | real    | Number of model feature dimensions whose retention is reduced                                                                                                     |
| hidden_units  | integer | Number of neurons in the model's hidden layer. If the model cannot be converged for a long time, increase the value of this parameter.                            |
| batch_size    | integer | Size of a batch in each iteration. It is recommended that the size be greater than or equal to the total training data volume to accelerate model convergence.    |
| feature_size  | integer | Length of the model feature, which is used to trigger retraining. This parameter is automatically updated after model training and does not need to be specified. |
| available     | Boolean | Specifies whether the model is converged. This parameter does not need to be specified.                                                                           |
| Is_training   | Boolean | Specifies whether the model is being trained. This parameter does not need to be specified.                                                                       |

| Name        | Type     | Description                                                                                                                                                                                                                                   |
|-------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| label       | "char"[] | Target task of the model. <ul style="list-style-type: none"><li>• S: startup time</li><li>• T: total time</li><li>• R: rows</li><li>• M: peak memory</li></ul> Currently, {S, T} or {R} is recommended due to model performance restrictions. |
| max         | bigint[] | Maximum value of each task label of the model, which is used to trigger retraining. This parameter does not need to be specified.                                                                                                             |
| acc         | real[]   | Accuracy of each model task. This parameter does not need to be specified.                                                                                                                                                                    |
| description | text     | Model comment                                                                                                                                                                                                                                 |

## 12.2.16 GS\_PACKAGE

**GS\_PACKAGE** records package information.

**Table 12-16** GS\_PACKAGE columns

| Name           | Type      | Description                                                |
|----------------|-----------|------------------------------------------------------------|
| oid            | oid       | Row identifier (hidden attribute, which must be specified) |
| pkgnamespace   | oid       | Schema to which a package belongs                          |
| pkgowner       | oid       | Owner of a package                                         |
| pkgname        | name      | Name of a package                                          |
| pkgspecsrc     | text      | Package specification                                      |
| pkgbodydeclsrc | text      | Package body                                               |
| pkgbodyinitsrc | text      | Package initialization source                              |
| pkgacl         | aclitem[] | Access permission                                          |

| Name      | Type    | Description                                                         |
|-----------|---------|---------------------------------------------------------------------|
| pkgsecdef | boolean | Specifies whether a user has the definer permission on the package. |

## 12.2.17 GS\_POLICY\_LABEL

**GS\_POLICY\_LABEL** records the resource label configuration information. One resource label corresponds to one or more records, and each record identifies the resource label to which a database resource belongs. Only users with the SYSADMIN or POLADMIN permission can access this system catalog.

Fully Qualified Domain Name (FQDN) identifies an absolute path of a database resource.

**Table 12-17** GS\_POLICY\_LABEL columns

| Name          | Type | Description                                                                                                                           |
|---------------|------|---------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid  | Row identifier (hidden attribute, which must be specified)                                                                            |
| labelname     | name | Resource label name                                                                                                                   |
| labeltype     | name | Resource tag type. Currently, the value is <b>RESOURCE</b> .                                                                          |
| fqdnnamespace | oid  | OID of a namespace to which an identified database resource belongs                                                                   |
| fqdnid        | oid  | OID of an identified database resource. If the database resource is a column, this column is the OID of the catalog.                  |
| relcolumn     | name | Column name. If the identified database resource is a column, this column indicates the column name. Otherwise, this column is empty. |
| fqdtype       | name | Type of the identified database resource, for example, schema, table, column, or view                                                 |

## 12.2.18 GS\_RECYCLEBIN

**GS\_RECYCLEBIN** describes details about objects in the recycle bin.

**Table 12-18** GS\_RECYCLEBIN columns

| Name           | Type                     | Description                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid            | oid                      | System column.                                                                                                                                                                                                                                                                                                                                                                                          |
| rcybaseid      | oid                      | Base table object ID, which references <b>gs_recyclebin.oid</b> .                                                                                                                                                                                                                                                                                                                                       |
| rcydbid        | oid                      | OID of the database to which the current object belongs.                                                                                                                                                                                                                                                                                                                                                |
| rcyrelid       | oid                      | OID of the current object.                                                                                                                                                                                                                                                                                                                                                                              |
| rcyname        | name                     | Name of the object in the recycle bin. The format is BIN\$ <i>unique_id</i> \$ <i>oid</i> \$. <i>unique_id</i> indicates the unique identifier with a maximum of 16 characters, and <i>oid</i> indicates the OID.                                                                                                                                                                                       |
| rcyoriginname  | name                     | Original object name.                                                                                                                                                                                                                                                                                                                                                                                   |
| rcyoperation   | char                     | Operation type. <ul style="list-style-type: none"><li>• <b>d</b>: drop</li><li>• <b>t</b>: truncate</li></ul>                                                                                                                                                                                                                                                                                           |
| rcytype        | integer                  | Object type. <ul style="list-style-type: none"><li>• <b>0</b>: table</li><li>• <b>1</b>: index</li><li>• <b>2</b>: TOAST table</li><li>• <b>3</b>: TOAST index</li><li>• <b>4</b>: sequence, indicating the sequence object that is automatically associated with the serial, bigserial, smallserial, and largeserial types.</li><li>• <b>5</b>: partition.</li><li>• <b>6</b>: global index.</li></ul> |
| rcyrecyclecsn  | bigint                   | CSN when an object is dropped or truncated.                                                                                                                                                                                                                                                                                                                                                             |
| rcyrecycletime | timestamp with time zone | Time when an object is dropped or truncated.                                                                                                                                                                                                                                                                                                                                                            |
| rcycreatecsn   | bigint                   | CSN when an object is created.                                                                                                                                                                                                                                                                                                                                                                          |
| rcychangecsn   | bigint                   | CSN when an object definition is modified.                                                                                                                                                                                                                                                                                                                                                              |
| rcynamespace   | oid                      | OID of the namespace that contains this relationship                                                                                                                                                                                                                                                                                                                                                    |
| rcyowner       | oid                      | Relationship owner                                                                                                                                                                                                                                                                                                                                                                                      |

| Name           | Type  | Description                                                                                                                                                                                 |
|----------------|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rcytablespace  | oid   | Tablespace where the relationship is stored. If the value is <b>0</b> , the default tablespace of the database is used. This column is meaningless if the relationship has no on-disk file. |
| rcyrelfilenode | oid   | File name of the recycle bin object on a disk, or <b>0</b> if none, used to restore the texture file when the <b>TRUNCATE</b> object is restored.                                           |
| rcycanrestore  | bool  | Specifies whether flashback can be performed separately.                                                                                                                                    |
| rcycanpurge    | bool  | Specifies whether the purge operation can be performed independently.                                                                                                                       |
| rcyfrozenxid   | xid32 | All transaction IDs before this one have been replaced by a permanent (frozen) transaction ID in the table.                                                                                 |
| rcyfrozenxid64 | xid   | All transaction IDs before this one have been replaced by a permanent (frozen) transaction ID in the table.                                                                                 |

## 12.2.19 GS\_SQL\_PATCH

**GS\_SQL\_PATCH** records the status information about all SQL patches.

**Table 12-19** GS\_SQL\_PATCH columns

| Name           | Type         | Description                                   |
|----------------|--------------|-----------------------------------------------|
| patch_name     | name         | Patch name.                                   |
| unique_sql_id  | bigint       | Global unique ID.                             |
| owner          | oid          | ID of the user who creates the patch.         |
| enable         | boolean      | Specifies whether the patch takes effect.     |
| status         | "char"       | Patch status (reserved column).               |
| abort          | boolean      | Specifies whether the patch is an abort hint. |
| hint_string    | text         | Hint text.                                    |
| hint_node      | pg_node_tree | Hint parsing and serialization result.        |
| original_query | text         | Original statement (reserved column).         |

| Name                | Type         | Description                                          |
|---------------------|--------------|------------------------------------------------------|
| patched_query       | text         | Patched statement (reserved column).                 |
| original_query_tree | pg_node_tree | Original statement parsing result (reserved column). |
| patched_query_tree  | pg_node_tree | Patched statement parsing result (reserved column).  |
| description         | text         | Patch description.                                   |

## 12.2.20 GS\_TXN\_SNAPSHOT

GS\_TXN\_SNAPSHOT is a timestamp-CSN mapping table. It periodically samples and maintains an appropriate time range to estimate the CSN value corresponding to the timestamp in the range.

**Table 12-20** GS\_TXN\_SNAPSHOT columns

| Name        | Data Type                | Description              |
|-------------|--------------------------|--------------------------|
| snptime     | timestamp with time zone | Snapshot time            |
| snpxmin     | bigint                   | Minimum snapshot ID      |
| snpcsn      | bigint                   | Snapshot CSN             |
| snpsnapshot | text                     | Serialized snapshot text |

## 12.2.21 GS\_UID

GS\_UID records the unique identification meta information of the hasuids attribute table in the database.

**Table 12-21** GS\_UID columns

| Name       | Type   | Description                                                |
|------------|--------|------------------------------------------------------------|
| relid      | oid    | OID of a table.                                            |
| uid_backup | bigint | Largest unique identifier that can be assigned to a table. |

## 12.2.22 PG\_AGGREGATE

PG\_AGGREGATE records information about aggregate functions. Each entry in PG\_AGGREGATE is an extension of an entry in PG\_PROC. The PG\_PROC entry carries the aggregate's name, input and output data types, and other information that is similar to ordinary functions.

**Table 12-22** PG\_AGGREGATE columns

| Name           | Type    | Reference                                 | Description                                                                                                                                                                                                                                              |
|----------------|---------|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| aggfnoid       | regproc | <b>prname</b> in <a href="#">PG_PROC</a>  | <a href="#">PG_PROC</a> prname of the aggregate function.                                                                                                                                                                                                |
| aggtransfn     | regproc | <b>prname</b> in <a href="#">PG_PROC</a>  | Transition function                                                                                                                                                                                                                                      |
| aggcollectfn   | regproc | <b>prname</b> in <a href="#">PG_PROC</a>  | Collect function                                                                                                                                                                                                                                         |
| aggfinalfn     | regproc | <b>prname</b> in <a href="#">PG_PROC</a>  | Final function ( <b>0</b> if none)                                                                                                                                                                                                                       |
| aggsortop      | oid     | <b>oid</b> in <a href="#">PG_OPERATOR</a> | Associated sort operator ( <b>0</b> if none)                                                                                                                                                                                                             |
| aggtranstype   | oid     | <b>oid</b> in <a href="#">PG_TYPE</a>     | Data type of the aggregate function's internal transition (state) data<br><br>The possible values and their meanings are defined by the types in <a href="#">pg_type.h</a> . The main two types are polymorphic (isPolymorphicType) and non-polymorphic. |
| agginitval     | text    | -                                         | Initial value of the transition state. This is a text column containing the initial value in its external string representation. If this column is null, the transition state value starts from null.                                                    |
| agginitcollect | text    | -                                         | Initial value of the collection state. This is a text column containing the initial value in its external string representation. If this column is null, the collection state value starts from null.                                                    |

| Name              | Type     | Reference | Description                                                                                                                                                                                                   |
|-------------------|----------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| aggkind           | "char"   | -         | Type of the aggregate function: <ul style="list-style-type: none"><li>• <b>n</b>: normal aggregate</li><li>• <b>o</b>: ordered set aggregate</li></ul>                                                        |
| aggnumdirect args | smallint | -         | Number of direct parameters (non-aggregation-related parameters) of the aggregate function of the ordered set aggregate type. For an aggregate function of the normal aggregate type, the value is <b>0</b> . |

### 12.2.23 PG\_AM

PG\_AM records information about index access methods. There is one row for each index access method supported by the system.

**Table 12-23** PG\_AM columns

| Name         | Type     | Reference | Description                                                                                                                                                                                                               |
|--------------|----------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid          | oid      | -         | Row identifier (hidden attribute, which must be specified)                                                                                                                                                                |
| amname       | name     | -         | Name of the access method                                                                                                                                                                                                 |
| amstrategies | smallint | -         | Number of operator strategies for the access method ( <b>0</b> if the access method does not have a fixed set of operator strategies)                                                                                     |
| amsupport    | smallint | -         | Number of support routines for the access method                                                                                                                                                                          |
| amcanorder   | Boolean  | -         | Specifies whether the access method supports ordered scans sorted by the indexed column's value. <ul style="list-style-type: none"><li>• <b>t</b> (true): supported.</li><li>• <b>f</b> (false): not supported.</li></ul> |



| Name           | Type    | Reference | Description                                                                                                                                                                                                                                          |
|----------------|---------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| amcanorderbyop | Boolean | -         | Specifies whether the access method supports ordered scans sorted by the result of an operator on the indexed column.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul> |
| amcanbackward  | Boolean | -         | Specifies whether the access method supports backward scanning.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>                                                       |
| amcanunique    | Boolean | -         | Specifies whether the access method supports unique indexes.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>                                                          |
| amcanmulticol  | Boolean | -         | Specifies whether the access method supports multi-column indexes.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>                                                    |
| amoptionalkey  | Boolean | -         | Specifies whether the access method supports scanning without any constraint for the first index column.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>              |
| amsearcharray  | Boolean | -         | Specifies whether the access method supports <b>ScalarArrayOpExpr</b> searches.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>                                       |
| amsearchnulls  | Boolean | -         | Specifies whether the access method supports <b>IS NULL/NOT NULL</b> searches.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): supported.</li> <li>• <b>f</b> (false): not supported.</li> </ul>                                        |

| Name             | Type    | Reference                                | Description                                                                                                                                                                                               |
|------------------|---------|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| amstorage        | Boolean | -                                        | Specifies whether the index storage data type can differ from the column data type. <ul style="list-style-type: none"><li>• <b>t</b> (true): allowed.</li><li>• <b>f</b> (false): no.t allowed.</li></ul> |
| amclusterable    | Boolean | -                                        | Specifies whether an index of this type can be clustered on. <ul style="list-style-type: none"><li>• <b>t</b> (true): allowed.</li><li>• <b>f</b> (false): no.t allowed.</li></ul>                        |
| ampredlocks      | Boolean | -                                        | Specifies whether an index of this type manages fine-grained predicate locks. <ul style="list-style-type: none"><li>• <b>t</b> (true): allowed.</li><li>• <b>f</b> (false): no.t allowed.</li></ul>       |
| amkeytype        | oid     | <b>oid</b> in <a href="#">PG_TYPE</a>    | Type of data stored in index ( <b>0</b> if it is not a fixed type).                                                                                                                                       |
| aminsert         | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Insert this tuple" function                                                                                                                                                                              |
| ambeginscan      | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Prepare for index scan" function                                                                                                                                                                         |
| amgettupl<br>e   | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Next valid tuple" function ( <b>0</b> if none)                                                                                                                                                           |
| amgetbitm<br>ap  | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Fetch all valid tuples" function ( <b>0</b> if none)                                                                                                                                                     |
| amrescan         | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "(Re)start index scan" function                                                                                                                                                                           |
| amendsca<br>n    | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Clean up after index scan" function                                                                                                                                                                      |
| ammarkpos        | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Mark current scan position" function                                                                                                                                                                     |
| amrestrpos       | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Restore marked scan position" function                                                                                                                                                                   |
| ammerge          | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Merge multiple indexes" function                                                                                                                                                                         |
| ambuild          | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Build new index" function                                                                                                                                                                                |
| ambuilde<br>mpty | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | "Build empty index" function                                                                                                                                                                              |

| Name            | Type    | Reference                                | Description                                                                       |
|-----------------|---------|------------------------------------------|-----------------------------------------------------------------------------------|
| ambulkdelete    | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Bulk-delete function                                                              |
| amvacuumcleanup | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Post-VACUUM cleanup function.                                                     |
| amcanreturn     | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Function to check whether the index supports index-only scans ( <b>0</b> if none) |
| amcostestimate  | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Function to estimate cost of an index scan                                        |
| amoptions       | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Function to parse and validate <b>reloptions</b> for an index                     |

## 12.2.24 PG\_AMOP

PG\_AMOP records information about operators associated with access method operator families. There is one row for each operator that is a member of an operator family. A family member can be either a search operator or an ordering operator. An operator can appear in more than one family, but cannot appear in more than one search position nor more than one ordering position within a family.

**Table 12-24** PG\_AMOP columns

| Name           | Type | Reference                                 | Description                                                                                                                              |
|----------------|------|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| oid            | oid  | -                                         | Row identifier (hidden attribute, which must be specified).                                                                              |
| amopfamily     | oid  | <b>oid</b> in <a href="#">PG_OPFAMILY</a> | Operator family of this entry.                                                                                                           |
| amoplefttype   | oid  | <b>oid</b> in <a href="#">PG_TYPE</a>     | Left-hand input data type of the operator. For details about the possible values and their description, see <a href="#">Data Type</a> .  |
| amoprightright | oid  | <b>oid</b> in <a href="#">PG_TYPE</a>     | Right-hand input data type of the operator. For details about the possible values and their description, see <a href="#">Data Type</a> . |

| Name           | Type     | Reference                          | Description                                                                                                               |
|----------------|----------|------------------------------------|---------------------------------------------------------------------------------------------------------------------------|
| amopstrategy   | smallint | -                                  | Number of operator strategies.                                                                                            |
| amoppurpose    | "char"   | -                                  | Purpose of the operator. <ul style="list-style-type: none"><li>• s: search</li><li>• o: order</li></ul>                   |
| amopopr        | oid      | oid in <a href="#">PG_OPERATOR</a> | OID of the operator.                                                                                                      |
| amopmethod     | oid      | oid in <a href="#">PG_AM</a>       | Operator family of the index access method.                                                                               |
| amopsortfamily | oid      | oid in <a href="#">PG_OPFAMILY</a> | The B-tree operator family according to which this entry sorts for an ordering operator ( <b>0</b> for a search operator) |

A search operator entry indicates that an index of this operator family can be searched to find all rows satisfying **WHERE indexed\_column operator constant**. Obviously, such an operator must return a Boolean value, and its left-hand input type must match the index's column data type.

An ordering operator entry indicates that an index of this operator family can be scanned to return rows in the order represented by **ORDER BY indexed\_column operator constant**. Such an operator could return any sortable data type, though again its left-hand input type must match the index's column data type. The exact semantics of **ORDER BY** are specified by the **amopsortfamily** column, which must reference the B-tree operator family for the operator's result type.

## 12.2.25 PG\_AMPROC

PG\_AMPROC records information about the support procedures associated with the access method operator families. There is one row for each support procedure that belongs to an operator family.

**Table 12-25** PG\_AMPROC columns

| Name         | Type | Reference                          | Description                                                 |
|--------------|------|------------------------------------|-------------------------------------------------------------|
| oid          | oid  | -                                  | Row identifier (hidden attribute, which must be specified). |
| amprocfamily | oid  | oid in <a href="#">PG_OPFAMILY</a> | Operator family of this entry.                              |

| Name            | Type     | Reference                         | Description                                                                                                                                         |
|-----------------|----------|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| amproclefttype  | oid      | oid in <a href="#">PG_TYPE</a>    | Left-hand input data type of the associated operator. For details about the possible values and their description, see <a href="#">Data Type</a> .  |
| amprocrighttype | oid      | oid in <a href="#">PG_TYPE</a>    | Right-hand input data type of the associated operator. For details about the possible values and their description, see <a href="#">Data Type</a> . |
| amprocnum       | smallint | -                                 | Support procedure number.                                                                                                                           |
| amproc          | regproc  | prname in <a href="#">PG_PROC</a> | OID of the procedure.                                                                                                                               |

The usual interpretation of the **amproclefttype** and **amprocrighttype** columns is that they identify the left and right input types of the operator(s) that a particular support procedure supports. For some access methods, these match the input data type(s) of the support procedure itself; for others not. There is a notion of "default" support procedures for an index, which are those with **amproclefttype** and **amprocrighttype** both equal to the index opclass's **opcintype**.

## 12.2.26 PG\_APP\_WORKLOADGROUP\_MAPPING

**PG\_APP\_WORKLOADGROUP\_MAPPING** provides load mapping group information in the database.

**Table 12-26** PG\_APP\_WORKLOADGROUP\_MAPPING columns

| Name            | Type | Description                                                |
|-----------------|------|------------------------------------------------------------|
| oid             | oid  | Row identifier (hidden attribute, which must be specified) |
| appname         | name | Application name                                           |
| workload_gpname | name | Mapped workload group name                                 |

## 12.2.27 PG\_ATTRDEF

**PG\_ATTRDEF** records default values of columns.

**Table 12-27** PG\_ATTRDEF columns

| Name     | Type         | Description                                                                                                                                                                                                                     |
|----------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid      | oid          | Row identifier (hidden attribute, which must be specified)                                                                                                                                                                      |
| adrelid  | oid          | Table to which a column belongs                                                                                                                                                                                                 |
| adnum    | smallint     | Number of columns                                                                                                                                                                                                               |
| adbin    | pg_node_tree | Internal representation of the default value of a column or of a generated expression                                                                                                                                           |
| adsrc    | text         | Internal representation of a readable default value or of a generated expression                                                                                                                                                |
| adgencol | "char"       | Specifies whether a column is a generated column. The value <b>s</b> indicates that the column is a generated column, and the value <b>\0</b> indicates that the column is an ordinary column. The default value is <b>\0</b> . |

## 12.2.28 PG\_ATTRIBUTE

PG\_ATTRIBUTE records information about table columns.

**Table 12-28** PG\_ATTRIBUTE columns

| Name          | Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| attrelid      | oid      | Table to which a column belongs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| attname       | name     | Column name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| atttypid      | oid      | Column type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| attstattarget | integer  | Level of details of statistics collected for this column by <b>ANALYZE</b> . <ul style="list-style-type: none"><li>• The value <b>0</b> indicates that no statistics should be collected.</li><li>• A negative value indicates that the system default statistic object is used.</li><li>• The exact meaning of positive values is data type-dependent.</li></ul> For scalar data types, <b>attstattarget</b> is both the target number of "most common values" to collect, and the target number of histogram bins to create. |
| attlen        | smallint | Copy of <b>PG_TYPE</b> in the column's type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| attnum        | smallint | Number of the column                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| attn_dims     | integer  | Number of dimensions if the column is an array ( <b>0</b> in other cases)                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Name          | Type      | Description                                                                                                                                                                                                                                                                                                                      |
|---------------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| attcacheoff   | integer   | This column is always set to <b>-1</b> on disks. When it is loaded into a row descriptor in the memory, it may be updated to cache the offset of the columns in the row.                                                                                                                                                         |
| atttypmod     | integer   | Type-specific data supplied at the table creation time (for example, the maximum length of a <b>varchar</b> column). This column is used as the third parameter when passing to type-specific input functions and length coercion functions. The value will generally be <b>-1</b> for types that do not need <b>atttypmod</b> . |
| attbyval      | Boolean   | Copy of <b>PG_TYPE</b> of this column's type.                                                                                                                                                                                                                                                                                    |
| attstorage    | "char"    | Copy of <b>PG_TYPE</b> of this column's type.                                                                                                                                                                                                                                                                                    |
| attalign      | "char"    | Copy of <b>PG_TYPE</b> of this column's type.                                                                                                                                                                                                                                                                                    |
| attnotnull    | Boolean   | A NOT NULL constraint. It is possible to change this column to enable or disable the constraint.                                                                                                                                                                                                                                 |
| atthasdef     | Boolean   | This column has a default value, in which case there will be a corresponding entry in the <b>PG_ATTRDEF</b> table that actually defines the value.                                                                                                                                                                               |
| attisdropped  | Boolean   | Indicates that this column has been deleted and is no longer valid. A deleted column is still physically present in the table but is ignored by the analyzer, so it cannot be accessed through SQL.                                                                                                                              |
| attislocal    | Boolean   | Indicates that this column is locally defined in the relationship. Note that a column can be locally defined and inherited simultaneously.                                                                                                                                                                                       |
| attinhcount   | integer   | Number of direct ancestors that this column has. A column with an ancestor cannot be dropped nor renamed.                                                                                                                                                                                                                        |
| attcollation  | oid       | Defined collation of a column.                                                                                                                                                                                                                                                                                                   |
| attacl        | aclitem[] | Permissions for column-level access.                                                                                                                                                                                                                                                                                             |
| attoptions    | text[]    | Column attribute. Currently, the following attributes are supported:<br><b>n_distinct</b> : number of <b>distinct</b> values of a column (excluding subtables).<br><b>n_distinct_inherited</b> : number of <b>distinct</b> values of a column (including subtables).                                                             |
| attfdwoptions | text[]    | Column attribute of a foreign table. Currently, <b>dist_fdw</b> , <b>file_fdw</b> , and <b>log_fdw</b> do not use foreign table column attributes.                                                                                                                                                                               |

| Name          | Type    | Description                                                                                                                                                                                                             |
|---------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| attinitdefval | bytea   | <b>attinitdefval</b> stores the default value expression. <b>ADD COLUMN</b> in the row-store table must use this column.                                                                                                |
| attkvtype     | tinyint | Specifies the key value type for a column. Value:<br>0. <b>ATT_KV_UNDEFINED</b> : default value.<br>1. <b>ATT_KV_TAG</b> : dimension.<br>2. <b>ATT_KV_FIELD</b> : indicator.<br>3. <b>ATT_KV_TIMETAG</b> : time column. |

## 12.2.29 PG\_AUTHID

PG\_AUTHID records information about database authentication identifiers (roles). The concept of users is contained in that of roles. A user is actually a role whose **rolcanlogin** has been set. Any role, whether its **rolcanlogin** is set or not, can use other roles as members.

For GaussDB, only one PG\_AUTHID exists, which is not available for every database. This system catalog is accessible only to users with the SYSADMIN permission.

**Table 12-29** PG\_AUTHID columns

| Name          | Type    | Description                                                                                                                                                                                                                                      |
|---------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid     | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                      |
| rolname       | name    | Role name.                                                                                                                                                                                                                                       |
| rolsuper      | Boolean | Specifies whether a role is the initial SYSADMIN with the highest permission. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                            |
| rolinherit    | Boolean | Specifies whether a role automatically inherits permissions of roles of which it is a member. <ul style="list-style-type: none"><li>• <b>t</b> (true): automatically inherited</li><li>• <b>f</b> (false): not automatically inherited</li></ul> |
| rolcreatorole | Boolean | Specifies whether a role can create more roles. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                          |



| Name           | Type                     | Description                                                                                                                                                                                                                                                                                                          |
|----------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rolcreatedb    | Boolean                  | Specifies whether a role can create databases.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                                         |
| rolcatupdate   | Boolean                  | Specifies whether a role can directly update system catalogs. Only the initial SYSADMIN whose <b>usesysid</b> is set to <b>10</b> has this permission. This permission is unavailable for other users.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul> |
| rolcanlogin    | Boolean                  | Specifies whether a role can log in, that is, whether a role can be given as the initial session authorization identifier.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                             |
| rolreplication | Boolean                  | Specifies whether a role has the replication permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                               |
| rolauditadmin  | Boolean                  | Specifies whether a role has the AUDITADMIN permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                                |
| rolsystemadmin | Boolean                  | Specifies whether a role has the SYSADMIN permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                                  |
| rolconnlimit   | integer                  | Maximum number of concurrent connections that the role can make (valid for roles that can log in).<br>The value <b>-1</b> indicates there is no limit.                                                                                                                                                               |
| rolpassword    | text                     | Password ciphertext ( <b>NULL</b> if no password).                                                                                                                                                                                                                                                                   |
| rolvalidbegin  | timestamp with time zone | Account validity start time ( <b>NULL</b> if no start time).                                                                                                                                                                                                                                                         |
| rolvaliduntil  | timestamp with time zone | Password expiry time ( <b>NULL</b> if no expiration).                                                                                                                                                                                                                                                                |

| Name              | Type    | Description                                                                                                                                                                      |
|-------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| roluseft          | Boolean | Specifies whether a role can perform operations on foreign tables.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul> |
| rolparentid       | oid     | OID of a group user to which the user belongs.                                                                                                                                   |
| roltabspace       | text    | Maximum size of a user data table.                                                                                                                                               |
| rolkind           | "char"  | Special user types, including private users and common users.                                                                                                                    |
| rolnodegroup      | oid     | Unsupported currently.                                                                                                                                                           |
| roltemp space     | text    | Maximum size of a user's temporary table, in KB.                                                                                                                                 |
| rolspillspace     | text    | Maximum size of data that can be written to disks when a user executes a job, in KB.                                                                                             |
| rolexcpdata       | text    | Query rules that can be set by users (reserved).                                                                                                                                 |
| rolmonitorad min  | Boolean | Specifies whether a role has the MONADMIN permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>              |
| roloperatorad min | Boolean | Specifies whether a role has the OPRADMIN permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>              |
| rolpolicyadmin    | Boolean | Specifies whether a role has the POLADMIN permission.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>              |

## 12.2.30 PG\_AUTH\_HISTORY

PG\_AUTH\_HISTORY records the authentication history of a role. This system catalog is accessible only to users with the SYSADMIN permission.

**Table 12-30** PG\_AUTH\_HISTORY columns

| Name         | Type                     | Description                                                                                                               |
|--------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------|
| oid          | oid                      | Row identifier (hidden attribute, which must be specified).                                                               |
| roloid       | oid                      | ID of a role.                                                                                                             |
| passwordtime | timestamp with time zone | Time of password creation and change.                                                                                     |
| rolpassword  | text                     | Ciphertext of the role password. The encryption mode is determined by the GUC parameter <b>password_encryption_type</b> . |

## 12.2.31 PG\_AUTH\_MEMBERS

PG\_AUTH\_MEMBERS records the membership between roles.

**Table 12-31** PG\_AUTH\_MEMBERS columns

| Name         | Type    | Description                                                                                                                    |
|--------------|---------|--------------------------------------------------------------------------------------------------------------------------------|
| roleid       | oid     | ID of a role that has a member.                                                                                                |
| member       | oid     | ID of a role that is a member of ROLEID.                                                                                       |
| grantor      | oid     | ID of a role that grants this membership.                                                                                      |
| admin_option | Boolean | Specifies whether a member can grant membership in ROLEID to others. The value can be <b>true</b> (yes) and <b>false</b> (no). |

## 12.2.32 PG\_CAST

PG\_CAST records the conversion relationship between data types.

**Table 12-32** PG\_CAST columns

| Name       | Type | Description                                                                      |
|------------|------|----------------------------------------------------------------------------------|
| oid        | oid  | Row identifier (hidden attribute, which must be specified).                      |
| castsource | oid  | OID of the source data type                                                      |
| casttarget | oid  | OID of the target data type                                                      |
| castfunc   | oid  | OID of the conversion function ( <b>0</b> if no conversion function is required) |

| Name        | Type   | Description                                                                                                                                                                                                                                                                                                                                             |
|-------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| castcontext | "char" | Conversion mode between the source and target data types. <ul style="list-style-type: none"><li>• <b>e</b>: Only explicit conversion can be performed (using the CAST or :: syntax).</li><li>• <b>i</b>: Implicit conversion can be performed.</li><li>• <b>a</b>: Both explicit and implicit conversion can be performed between data types.</li></ul> |
| castmethod  | "char" | Conversion method. <ul style="list-style-type: none"><li>• <b>f</b>: Conversion is performed using the specified function in the <b>castfunc</b> column.</li><li>• <b>b</b>: Binary forcible conversion rather than the specified function in the <b>castfunc</b> column is performed between data types.</li></ul>                                     |

## 12.2.33 PG\_CLASS

PG\_CLASS records database objects and their relationship.

**Table 12-33** PG\_CLASS columns

| Name          | Type | Description                                                                                                                                                                                      |
|---------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid  | Row identifier (hidden attribute, which must be specified).                                                                                                                                      |
| relname       | name | Name of an object, such as a table, index, or view.                                                                                                                                              |
| relnamespace  | oid  | OID of the namespace that contains the relationship.                                                                                                                                             |
| reltype       | oid  | Data type that corresponds to the table's row type. The index is 0 because the index does not have <b>PG_TYPE</b> records.                                                                       |
| reloftype     | oid  | OID of the composite type ( <b>0</b> for other types).                                                                                                                                           |
| relowner      | oid  | Owner of the relationship.                                                                                                                                                                       |
| relam         | oid  | Access method used, such as B-tree and hash, if this is an index.                                                                                                                                |
| relfilenode   | oid  | Name of the on-disk file of this relationship ( <b>0</b> if such file does not exist).                                                                                                           |
| reltablespace | oid  | Tablespace in which this relationship is stored. If the value is <b>0</b> , the default tablespace in this database is used. This column is meaningless if the relationship has no on-disk file. |

| Name           | Type             | Description                                                                                                                                                                                                                                                                                                                             |
|----------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| relpages       | double precision | Size of the on-disk representation of the table in pages (of size <b>BLCKSZ</b> ). This is only an estimate used by the optimizer.                                                                                                                                                                                                      |
| reltuples      | double precision | Number of rows in the table. This is only an estimate used by the optimizer.                                                                                                                                                                                                                                                            |
| relallvisible  | integer          | Number of pages marked as all visible in the table. This column is used by the optimizer for optimizing SQL execution. It is updated by <b>VACUUM</b> , <b>ANALYZE</b> , and a few DDL statements such as <b>CREATE INDEX</b> .                                                                                                         |
| reltoastrelid  | oid              | OID of the TOAST table associated with the table ( <b>0</b> if no TOAST table exists).<br>The TOAST table stores large columns "offline" in a secondary table.                                                                                                                                                                          |
| reltoastidxid  | oid              | OID of the index for a TOAST table ( <b>0</b> for a table other than a TOAST table).                                                                                                                                                                                                                                                    |
| relhasindex    | Boolean          | Its value is <b>true</b> if this column is a table and has (or recently had) at least one index.<br>It is set by <b>CREATE INDEX</b> but is not immediately cleared by <b>DROP INDEX</b> . If the <b>VACUUM</b> process detects that a table has no index, it clears the <b>relhasindex</b> column and sets the value to <b>false</b> . |
| relisshared    | Boolean          | Its value is <b>true</b> if the table is shared across all database nodes in the database. Only certain system catalogs (such as PG_DATABASE) are shared.                                                                                                                                                                               |
| relpersistence | "char"           | <ul style="list-style-type: none"><li>• <b>p</b>: permanent table</li><li>• <b>u</b>: non-log table</li><li>• <b>t</b>: temporary table</li><li>• <b>g</b>: global temporary table.</li></ul>                                                                                                                                           |

| Name             | Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| relkind          | "char"   | <ul style="list-style-type: none"> <li>• <b>r</b>: ordinary table</li> <li>• <b>i</b>: index</li> <li>• <b>I</b>: global index of a partitioned table</li> <li>• <b>s</b>: sequence</li> <li>• <b>L</b>: long sequence</li> <li>• <b>v</b>: view</li> <li>• <b>c</b>: composite type</li> <li>• <b>t</b>: TOAST table</li> <li>• <b>f</b>: foreign table</li> <li>• <b>m</b>: materialized view</li> </ul> |
| relnatts         | smallint | Number of user columns in the relationship (excluding system columns). <a href="#">PG_ATTRIBUTE</a> has the same number of rows as the user columns.                                                                                                                                                                                                                                                       |
| relchecks        | smallint | Number of check constraints in the table. For details, see the system catalog <a href="#">PG_CONSTRAINT</a> .                                                                                                                                                                                                                                                                                              |
| relhasoids       | Boolean  | Its value is <b>true</b> if an OID is generated for each row of the relationship. Otherwise, the value is <b>false</b> .                                                                                                                                                                                                                                                                                   |
| relhaspkey       | Boolean  | Its value is <b>true</b> if the table has (or once had) a primary key. Otherwise, the value is <b>false</b> .                                                                                                                                                                                                                                                                                              |
| relhasrules      | Boolean  | Its value is <b>true</b> if the table has rules. For details, see the system catalog <a href="#">PG_REWRITE</a> .                                                                                                                                                                                                                                                                                          |
| relhastriggers   | Boolean  | The value is <b>true</b> if the table has (or once had) triggers. Triggers of the table and view are recorded in the system catalog <a href="#">PG_TRIGGER</a> .                                                                                                                                                                                                                                           |
| relhassubclass   | Boolean  | Its value is <b>true</b> if the table has (or once had) any inheritance child table. Otherwise, the value is <b>false</b> .                                                                                                                                                                                                                                                                                |
| relhasclusterkey | Boolean  | Specifies whether the local cluster storage is used. <ul style="list-style-type: none"> <li>• <b>true</b>: yes</li> <li>• <b>false</b>: no</li> </ul>                                                                                                                                                                                                                                                      |
| relrowmovement   | Boolean  | Specifies whether row migration is allowed when the partitioned table is updated. <ul style="list-style-type: none"> <li>• <b>true</b>: Row migration is allowed.</li> <li>• <b>false</b>: Row migration is not allowed.</li> </ul>                                                                                                                                                                        |

| Name           | Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| parttype       | "char"    | Specifies whether the table or index has the property of a partitioned table. <ul style="list-style-type: none"> <li>• <b>p</b>: The table or index has the property of a partitioned table.</li> <li>• <b>n</b>: The table or index does not have the property of a partitioned table.</li> <li>• <b>v</b>: The table is a value partitioned table in HDFS.</li> <li>• <b>s</b>: The table is a level-2 partitioned table.</li> </ul>                                                                                               |
| relfrozenxid   | xid32     | All transaction IDs before this one have been replaced by a permanent ("frozen") transaction ID in the table. This column is used to track whether the table needs to be vacuumed to prevent transaction ID wraparound (or to allow PG_CLOG to be shrunk). The value is <b>0 (InvalidTransactionId)</b> if the relationship is not a table.<br><br>To ensure forward compatibility, this column is reserved. The <b>relfrozenxid64</b> column is added to record the information.                                                    |
| relacl         | aclitem[] | Access permissions.<br><br>The command output of the query is as follows:<br>rolename=xxxx/yyyy -- Assigning permissions to a role<br>=xxxx/yyyy -- Assigning the permission to public<br><br>xxxx indicates assigned permissions, and yyyy indicates roles with the assigned permissions. For details on permission descriptions, see <a href="#">Table 12-34</a> .                                                                                                                                                                 |
| reloptions     | text[]    | Table or index access method, using character strings in the format of "keyword=value"                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| relreplident   | "char"    | Identifier of a decoding column in logical decoding. <ul style="list-style-type: none"> <li>• <b>d</b>: default (primary key, if any)</li> <li>• <b>n</b>: none</li> <li>• <b>f</b>: all columns</li> <li>• <b>i</b>: The indisreplident of the index is specified or the default index is used.</li> </ul>                                                                                                                                                                                                                          |
| relfrozenxid64 | xid       | All transaction IDs before this one have been replaced by a permanent ("frozen") transaction ID in the table. This column is used to track whether the table needs to be vacuumed to prevent transaction ID wraparound (or to allow PG_CLOG to be shrunk). The value is <b>0 (InvalidTransactionId)</b> if the relationship is not a table.<br><br>For a global temporary table, this field is meaningless. You can view <b>relfrozenxid64</b> of the global temporary table of each session in the pg_catalog.pg_gtt_relstats view. |

| Name         | Type           | Description                                                                                                                                                                                                                                                                                                                         |
|--------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| relbucket    | oid            | Bucket information in <a href="#">PG_HASHBUCKET</a> .                                                                                                                                                                                                                                                                               |
| relbucketkey | int2vect<br>or | Column number of a hash partition.                                                                                                                                                                                                                                                                                                  |
| relminmxid   | xid            | All multi-transaction IDs before this one have been replaced by a transaction ID in the table. This is used to track whether the table needs to be vacuumed in order to prevent multi-transaction IDs wraparound or to allow PG_CLOG to be shrunk. The value is <b>0 (InvalidTransactionId)</b> if the relationship is not a table. |

**Table 12-34** Description of permissions

| Parameter | Description                                     |
|-----------|-------------------------------------------------|
| r         | SELECT (read)                                   |
| w         | UPDATE (write)                                  |
| a         | INSERT (insert)                                 |
| d         | DELETE                                          |
| D         | TRUNCATE                                        |
| x         | REFERENCES                                      |
| t         | TRIGGER                                         |
| X         | EXECUTE                                         |
| U         | USAGE                                           |
| C         | CREATE                                          |
| c         | CONNECT                                         |
| T         | TEMPORARY                                       |
| A         | ALTER                                           |
| P         | DROP                                            |
| m         | COMMENT                                         |
| i         | INDEX                                           |
| v         | VACUUM                                          |
| *         | Authorization options for preceding permissions |



## 12.2.34 PG\_COLLATION

PG\_COLLATION describes available collations, which are essentially mappings from an SQL name to OS local categories.

**Table 12-35** PG\_COLLATION columns

| Name          | Type    | Reference                         | Description                                                                                                                 |
|---------------|---------|-----------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| oid           | oid     | -                                 | Row identifier (hidden attribute, which must be specified)                                                                  |
| collname      | name    | -                                 | Collation name (unique per namespace and encoding)                                                                          |
| collnamespace | oid     | <b>oid</b> in <b>PG_NAMESPACE</b> | OID of the namespace that contains this collation                                                                           |
| collowner     | oid     | <b>oid</b> in <b>PG_AUTHID</b>    | Owner of the collation                                                                                                      |
| collencoding  | integer | -                                 | Encoding in which the collation is applicable, or <b>-1</b> if it works for any encoding. It is compatible with PostgreSQL. |
| collcollate   | name    | -                                 | <b>LC_COLLATE</b> for this collation object                                                                                 |
| collctype     | name    | -                                 | <b>LC_CTYPE</b> for this collation object                                                                                   |

## 12.2.35 PG\_CONSTRAINT

PG\_CONSTRAINT records check, primary key, and unique constraints on tables.

**Table 12-36** PG\_CONSTRAINT columns

| Name         | Type | Description                                                 |
|--------------|------|-------------------------------------------------------------|
| oid          | oid  | Row identifier (hidden attribute, which must be specified). |
| conname      | name | Constraint name (not necessarily unique).                   |
| connamespace | oid  | OID of the namespace that contains the constraint           |

| Name          | Type    | Description                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| contype       | "char"  | <ul style="list-style-type: none"><li>• <b>c</b>: check constraint</li><li>• <b>p</b>: primary key constraint</li><li>• <b>u</b>: unique constraint</li><li>• <b>t</b>: trigger constraint</li><li>• <b>x</b>: mutual exclusion constraint</li><li>• <b>f</b>: foreign key constraint</li><li>• <b>s</b>: clustering constraint</li><li>• <b>i</b>: invalid constraint</li></ul> |
| condeferrable | Boolean | Specifies whether the constraint is deferrable. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no.</li></ul>                                                                                                                                                                                                                                  |
| condeferred   | Boolean | Specifies whether the constraint can be deferrable by default. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no.</li></ul>                                                                                                                                                                                                                   |
| convalidated  | Boolean | Specifies whether the constraint is valid. Currently, it can be set to <b>false</b> only for foreign key and check constraints. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no.</li></ul>                                                                                                                                                  |
| conrelid      | oid     | Table containing this constraint ( <b>0</b> if it is not a table constraint).                                                                                                                                                                                                                                                                                                    |
| contypid      | oid     | Domain containing this constraint ( <b>0</b> if it is not a domain constraint).                                                                                                                                                                                                                                                                                                  |
| conindid      | oid     | ID of the index associated with the constraint.                                                                                                                                                                                                                                                                                                                                  |
| confrelid     | oid     | Referenced table if this constraint is a foreign key. Otherwise, the value is <b>0</b> .                                                                                                                                                                                                                                                                                         |
| confupdtype   | "char"  | Foreign key update action code. <ul style="list-style-type: none"><li>• <b>a</b>: no action.</li><li>• <b>r</b>: restriction.</li><li>• <b>c</b>: cascading.</li><li>• <b>n</b>: The parameter is set to null.</li><li>• <b>d</b>: The default value is used.</li></ul>                                                                                                          |

| Name          | Type       | Description                                                                                                                                                                                                                                                                    |
|---------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| confdeltype   | "char"     | Foreign key deletion action code <ul style="list-style-type: none"> <li>• <b>a</b>: no action.</li> <li>• <b>r</b>: restriction.</li> <li>• <b>c</b>: cascading.</li> <li>• <b>n</b>: The parameter is set to null.</li> <li>• <b>d</b>: The default value is used.</li> </ul> |
| confmatchtype | "char"     | Foreign key match type. <ul style="list-style-type: none"> <li>• <b>f</b>: full match.</li> <li>• <b>p</b>: partial match.</li> <li>• <b>u</b>: unspecified (The NULL value can be matched if <b>f</b> is specified.)</li> </ul>                                               |
| conislocal    | Boolean    | Specifies whether the constraint is defined locally for the relation. <ul style="list-style-type: none"> <li>• <b>true</b>: yes.</li> <li>• <b>false</b>: no.</li> </ul>                                                                                                       |
| coninhcount   | integer    | Number of direct inheritance parent tables that this constraint has. When the value is not <b>0</b> , the constraint cannot be deleted or renamed.                                                                                                                             |
| connoinherit  | Boolean    | Specifies whether the constraint can be inherited. <ul style="list-style-type: none"> <li>• <b>true</b>: yes.</li> <li>• <b>false</b>: no.</li> </ul>                                                                                                                          |
| consoft       | Boolean    | Specifies whether the column indicates an informational constraint. <ul style="list-style-type: none"> <li>• <b>true</b>: yes.</li> <li>• <b>false</b>: no.</li> </ul>                                                                                                         |
| conopt        | Boolean    | Specifies whether you can use the informational constraint to optimize the execution plan. <ul style="list-style-type: none"> <li>• <b>true</b>: yes.</li> <li>• <b>false</b>: no.</li> </ul>                                                                                  |
| conkey        | smallint[] | Column list of the constrained control if this column is a table constraint.                                                                                                                                                                                                   |
| confkey       | smallint[] | List of referenced columns if this column is a foreign key.                                                                                                                                                                                                                    |
| confeqop      | oid[]      | ID list of the equality operators for PK = FK comparisons if this column is a foreign key.                                                                                                                                                                                     |

| Name         | Type         | Description                                                                                                                                                     |
|--------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| conppeqop    | oid[]        | ID list of the equality operators for PK = PK comparisons if this column is a foreign key.                                                                      |
| conffeqop    | oid[]        | ID list of the equality operators for FK = FK comparisons if this column is a foreign key. The value is empty because foreign keys are not supported currently. |
| conexclp     | oid[]        | ID list of the per-column exclusion operators if this column is an exclusion constraint.                                                                        |
| conbin       | pg_node_tree | Internal representation of the expression if this column is a check constraint.                                                                                 |
| consrc       | text         | Readable representation of the expression if this column is a check constraint.                                                                                 |
| conincluding | smallint[]   | Not for constraint, but will be included in the attribute column of <b>INDEX</b> .                                                                              |

**NOTICE**

- **consrc** is not updated when referenced objects change and does not track new column names. Instead of relying on this column to update, you are advised to use `pg_get_constraintdef()` to extract the definition of a check constraint.
- **relchecks** of **PG\_CLASS** must agree with the number of check-constraint entries found in the table for each relationship.

## 12.2.36 PG\_CONVERSION

PG\_CONVERSION describes encoding conversion information.

**Table 12-37** PG\_CONVERSION columns

| Name           | Type    | Reference                         | Description                                                |
|----------------|---------|-----------------------------------|------------------------------------------------------------|
| oid            | oid     | -                                 | Row identifier (hidden attribute, which must be specified) |
| conname        | name    | -                                 | Conversion name (unique within a namespace)                |
| connamespace   | oid     | <b>oid</b> in <b>PG_NAMESPACE</b> | OID of the namespace that contains this conversion         |
| conowner       | oid     | <b>oid</b> in <b>PG_AUTHID</b>    | Owner of the conversion                                    |
| conforencoding | integer | -                                 | Source encoding ID                                         |

| Name          | Type    | Reference                                | Description                                                                                          |
|---------------|---------|------------------------------------------|------------------------------------------------------------------------------------------------------|
| contoencoding | integer | -                                        | Destination encoding ID                                                                              |
| conproc       | regproc | <b>prname</b> in <a href="#">PG_PROC</a> | Conversion procedure                                                                                 |
| condefault    | Boolean | -                                        | If this is the default conversion, the value is <b>true</b> . Otherwise, the value is <b>false</b> . |

## 12.2.37 PG\_DATABASE

PG\_DATABASE records information about available databases.

**Table 12-38** PG\_DATABASE columns

| Name          | Type    | Description                                                                                                                                                                                                           |
|---------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid     | Row identifier (hidden attribute, which must be specified).                                                                                                                                                           |
| datname       | name    | Database name                                                                                                                                                                                                         |
| datdba        | oid     | Owner of the database, usually the user who created it.                                                                                                                                                               |
| encoding      | integer | Character encoding for the database.                                                                                                                                                                                  |
| datcollate    | name    | Sequence used by the database.                                                                                                                                                                                        |
| datctype      | name    | Character type used by the database.                                                                                                                                                                                  |
| datistemplate | Boolean | Specifies whether the database can be used as a template database.                                                                                                                                                    |
| datallowconn  | Boolean | If the value is <b>true</b> , users can connect to the database. If the value is <b>false</b> , no one can connect to this database. This column is used to protect the <b>template0</b> database from being altered. |
| datconnlimit  | integer | Maximum number of concurrent connections allowed on this database. The value <b>-1</b> indicates no limit.                                                                                                            |
| datlastsysoid | oid     | Last system OID in the database.                                                                                                                                                                                      |

| Name             | Type      | Description                                                                                                                                                                                                                                                                                                                                       |
|------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datfrozenxid     | xid32     | Tracks whether the database needs to be vacuumed to prevent transaction ID wraparound. This column is deprecated in the current version. To ensure forward compatibility, this column is reserved. The <b>datfrozenxid64</b> column is added to record the information.                                                                           |
| dattablespace    | oid       | Default tablespace of the database.                                                                                                                                                                                                                                                                                                               |
| datcompatibility | name      | Database compatibility mode. Currently, four compatibility modes are supported: A, B, C, and PG, indicating that the Oracle, MySQL, Teradata, and Postgres databases are compatible.                                                                                                                                                              |
| datacl           | aclitem[] | Access permission.                                                                                                                                                                                                                                                                                                                                |
| datfrozenxid64   | xid       | Tracks whether the database needs to be vacuumed to prevent transaction ID wraparound.                                                                                                                                                                                                                                                            |
| datminmxid       | xid       | All multi-transaction IDs before this one have been replaced by a transaction ID in the database. This is used to track whether the database needs to be vacuumed in order to prevent transaction IDs wraparound or to allow PG_CLOG to be shrunk. It is the minimum value of <b>relminmxid</b> in <b>PG_CLASS</b> of all tables in the database. |

## 12.2.38 PG\_DB\_ROLE\_SETTING

**PG\_DB\_ROLE\_SETTING** records the default values of configuration items bound to each role and database when the database is running.

**Table 12-39** PG\_DB\_ROLE\_SETTING columns

| Name        | Type   | Description                                                                               |
|-------------|--------|-------------------------------------------------------------------------------------------|
| setdatabase | oid    | Database corresponding to the configuration items ( <b>0</b> if no database is specified) |
| setrole     | oid    | Role corresponding to the configuration items ( <b>0</b> if no role is specified)         |
| setconfig   | text[] | Default value of GUC items.                                                               |

## 12.2.39 PG\_DEFAULT\_ACL

**PG\_DEFAULT\_ACL** records initial permissions assigned to newly created objects.

**Table 12-40** PG\_DEFAULT\_ACL columns

| Name            | Type      | Description                                                                                                                                                                                                                                                                                                                                            |
|-----------------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid             | oid       | Row identifier (hidden attribute, which must be specified)                                                                                                                                                                                                                                                                                             |
| defaclrole      | oid       | ID of the role associated with the permission                                                                                                                                                                                                                                                                                                          |
| defaclnamespace | oid       | Namespace associated with the permission ( <b>0</b> if no ID)                                                                                                                                                                                                                                                                                          |
| defaclobjtype   | "char"    | Object type of the permission <ul style="list-style-type: none"><li>• <b>r</b> indicates a table or view.</li><li>• <b>S</b> indicates a sequence.</li><li>• <b>f</b> indicates a function.</li><li>• <b>T</b> indicates a type.</li><li>• <b>K</b> indicates the client master key.</li><li>• <b>k</b> indicates the column encryption key.</li></ul> |
| defaclacl       | aclitem[] | Access permissions that this type of object should have on creation                                                                                                                                                                                                                                                                                    |

## 12.2.40 PG\_DEPEND

**PG\_DEPEND** records the dependency between database objects. This information allows **DROP** commands to find which other objects must be dropped by **DROP CASCADE** or prevent dropping in the **DROP RESTRICT** case.

See also [PG\\_SHDEPEND](#), which performs a similar function for dependencies involving objects that are shared across databases.

**Table 12-41** PG\_DEPEND columns

| Name     | Type    | Reference                              | Description                                                                                                                                                   |
|----------|---------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| classid  | oid     | <b>oid</b> in <a href="#">PG_CLASS</a> | OID of the system catalog where a dependent object resides                                                                                                    |
| objid    | oid     | Any OID column                         | OID of the dependent object                                                                                                                                   |
| objsubid | integer | -                                      | For a table column, this is the column number ( <b>objid</b> and <b>classid</b> refer to the table itself). The value is <b>0</b> for all other object types. |

| Name        | Type    | Reference              | Description                                                                                                                                                         |
|-------------|---------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| refclassid  | oid     | oid in <b>PG_CLASS</b> | OID of the system catalog where a referenced object resides                                                                                                         |
| refobjid    | oid     | Any OID column         | OID of the referenced object                                                                                                                                        |
| refobjsubid | integer | -                      | For a table column, this is the column number ( <b>refobjid</b> and <b>refclassid</b> refer to the table itself). The value is <b>0</b> for all other object types. |
| deptype     | "char"  | -                      | A code defining the specific semantics of this dependency                                                                                                           |

In all cases, a PG\_DEPEND entry indicates that the referenced object cannot be dropped without also dropping the dependent object. However, there are several subflavors identified by **deptype**:

- **DEPENDENCY\_NORMAL** (n): A normal relationship between separately created objects. The dependent object can be dropped without affecting the referenced object. The referenced object can only be dropped by specifying **CASCADE**, in which case the dependent object is dropped too. Example: a table column has a normal dependency on its data type.
- **DEPENDENCY\_AUTO** (a): The dependent object can be dropped separately from the referenced object, and should be automatically dropped (regardless of **RESTRICT** or **CASCADE** mode) if the referenced object is dropped. Example: a named constraint on a table is made autodependent on the table, so that it will go away if the table is dropped.
- **DEPENDENCY\_INTERNAL** (i): The dependent object was created as part of creation of the referenced object, and is only a part of its internal implementation. A **DROP** of the dependent object will be disallowed outright (We'll tell the user to issue a **DROP** against the referenced object, instead). A **DROP** of the referenced object will be propagated through to drop the dependent object whether **CASCADE** is specified or not.
- **DEPENDENCY\_EXTENSION** (e): The dependent object is a member of the extension of the referenced object. The dependent object can be dropped only via **DROP EXTENSION** on the referenced object. Functionally this dependency type acts the same as an internal dependency, but it is kept separate for clarity and to simplify **GS\_DUMP**.
- **DEPENDENCY\_PIN** (p): There is no dependent object; this type of entry is a signal that the system itself depends on the referenced object, and so that object must never be deleted. Entries of this type are created only by **initdb**. The columns for the dependent object contain zeroes.

## 12.2.41 PG\_DESCRIPTION

**PG\_DESCRIPTION** records optional descriptions (comments) for each database object. Descriptions of many built-in system objects are provided in the initial contents of **PG\_DESCRIPTION**.



See also [PG\\_SHDESCRIPTION](#), which provides a similar function for descriptions involving objects that are shared within the entire database.

**Table 12-42** PG\_DESCRIPTION columns

| Name        | Type    | Reference                              | Description                                                                                                                                       |
|-------------|---------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| objoid      | oid     | Any OID column                         | OID of the object that this description pertains to                                                                                               |
| classoid    | oid     | <b>oid</b> in <a href="#">PG_CLASS</a> | OID of the system catalog where the object appears                                                                                                |
| objsubid    | integer | -                                      | Column number for a comment on a table column ( <b>objoid</b> and <b>classoid</b> refer to the table itself); <b>0</b> for all other object types |
| description | text    | -                                      | Arbitrary text that serves as the description of the object                                                                                       |

## 12.2.42 PG\_DIRECTORY

**PG\_DIRECTORY** stores directory objects added by users. You can execute the CREATE DIRECTORY statement to add records to this system catalog. When **enable\_access\_server\_directory** is set to **off**, only the initial user can create directory objects. When **enable\_access\_server\_directory** is set to **on**, users with the SYSADMIN permission and users inheriting permissions of the built-in role gs\_role\_directory\_create can create directory objects.

**Table 12-43** PG\_DIRECTORY columns

| Name    | Type      | Description                                                |
|---------|-----------|------------------------------------------------------------|
| oid     | oid       | Row identifier (hidden attribute, which must be specified) |
| dirname | name      | Name of a directory object                                 |
| owner   | oid       | Owner of a directory object                                |
| dirpath | text      | Directory path.                                            |
| diracl  | aclitem[] | Access permissions.                                        |

## 12.2.43 PG\_ENUM

**PG\_ENUM** contains entries showing the values and labels for each enumerated type. The internal representation of a given enumerated value is actually the OID of its associated row in PG\_ENUM.

**Table 12-44** PG\_ENUM columns

| Name          | Type | Reference                      | Description                                                            |
|---------------|------|--------------------------------|------------------------------------------------------------------------|
| oid           | oid  | -                              | Row identifier (hidden attribute, which must be specified).            |
| enumtypid     | oid  | oid in <a href="#">PG_TYPE</a> | OID of the <a href="#">PG_TYPE</a> entry owning this enumerated value. |
| enumsortorder | real | -                              | Sort position of this enumerated value within its enumerated type.     |
| enumlabel     | name | -                              | Textual label for this enumerated value.                               |

The OIDs for **PG\_ENUM** rows follow a special rule: even-numbered OIDs are guaranteed to be ordered in the same way as the sort ordering of their enumerated type. If two even OIDs belong to the same enumerated type, the smaller OID must have the smaller **enumsortorder** value. Odd-numbered OID values need bear no relationship to the sort order. This rule allows the enumerated comparison routines to avoid catalog lookups in many common cases. The routines that create and alter enumerated types attempt to assign even OIDs to enumerated values whenever possible.

When an enumerated type is created, its members are assigned sort-order positions from 1 to  $n$ . However, members added later might be given negative or fractional values of **enumsortorder**. The only requirement on these values is that they be correctly ordered and unique within each enumerated type.

## 12.2.44 PG\_FOREIGN\_SERVER

**PG\_FOREIGN\_SERVER** records foreign server definitions. A foreign server describes a source of external data, such as a remote server. Foreign servers are accessed via foreign-data wrappers.

**Table 12-45** PG\_FOREIGN\_SERVER columns

| Name     | Type | Reference                                      | Description                                                |
|----------|------|------------------------------------------------|------------------------------------------------------------|
| oid      | oid  | -                                              | Row identifier (hidden attribute, which must be specified) |
| srvname  | name | -                                              | Name of a foreign server                                   |
| srvowner | oid  | OID in <a href="#">PG_AUTHID</a>               | Owner of the foreign server                                |
| srvfdw   | oid  | oid in <a href="#">PG_FOREIGN_DATA_WRAPPER</a> | OID of the foreign-data wrapper on a foreign server        |
| srvtype  | text | -                                              | Type of the server (optional)                              |

| Name       | Type      | Reference | Description                                                                           |
|------------|-----------|-----------|---------------------------------------------------------------------------------------|
| srvversion | text      | -         | Version of the server (optional)                                                      |
| srvacl     | aclitem[] | -         | Access permissions                                                                    |
| srvoptions | text[]    | -         | Option used for foreign servers, expressed in a string in the format of keyword=value |

## 12.2.45 PG\_HASHBUCKET

**PG\_HASHBUCKET** records hash bucket information.

**Table 12-46** PG\_HASHBUCKET columns

| Name           | Type             | Description                                                                                                                                                                                           |
|----------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid            | oid              | Row identifier (hidden attribute, which must be specified)                                                                                                                                            |
| bucketid       | oid              | Hash value calculated for a bucket vector. The hash value can be used to accelerate the search for a bucket vector.                                                                                   |
| bucketcnt      | integer          | Number of shards                                                                                                                                                                                      |
| bucketmap size | integer          | Total number of shards on all DN's                                                                                                                                                                    |
| bucketref      | integer          | Reserved column with 1 as its default value                                                                                                                                                           |
| bucketvector   | oidvector_extend | Records all bucket IDs contained in the bucket information in this row. A unique index is created in this column. Tables with the same bucket ID share the <b>PG_HASHBUCKET</b> data in the same row. |

## 12.2.46 PG\_INDEX

**PG\_INDEX** records part of index information. The rest is mostly recorded in **PG\_CLASS**.

**Table 12-47** PG\_INDEX columns

| Name       | Type | Description                                     |
|------------|------|-------------------------------------------------|
| indexrelid | oid  | OID of the <b>PG_CLASS</b> entry for the index. |

| Name           | Type     | Description                                                                                                                                                                                                                                                                                                                                    |
|----------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| indrelid       | oid      | OID of the <b>PG_CLASS</b> entry for the table that uses the index.                                                                                                                                                                                                                                                                            |
| indnatts       | smallint | Number of columns in the index.                                                                                                                                                                                                                                                                                                                |
| indisunique    | Boolean  | Specifies whether the index is unique. <ul style="list-style-type: none"> <li>• <b>true</b>: The index is unique.</li> <li>• <b>false</b>: The index is not unique.</li> </ul>                                                                                                                                                                 |
| indisprimary   | Boolean  | Specifies whether the index is the primary key of the table. <ul style="list-style-type: none"> <li>• <b>true</b>: The index is the primary key of the table. <b>indisunique</b> should always be <b>true</b> when the value of this column is <b>true</b>.</li> <li>• <b>false</b>: The index is not the primary key of the table.</li> </ul> |
| indisexclusion | Boolean  | Specifies whether the index supports exclusive constraints. <ul style="list-style-type: none"> <li>• <b>true</b>: The index supports exclusive constraints.</li> <li>• <b>false</b>: The index does not support exclusive constraints.</li> </ul>                                                                                              |
| indimmediate   | Boolean  | Specifies whether to check the uniqueness of data to be inserted. <ul style="list-style-type: none"> <li>• <b>true</b>: The uniqueness check is performed immediately when data is inserted.</li> <li>• <b>false</b>: The uniqueness check is not performed when data is inserted.</li> </ul>                                                  |
| indisclustered | Boolean  | Specifies whether the table is clustered on the index. <ul style="list-style-type: none"> <li>• <b>true</b>: The table is clustered on the index.</li> <li>• <b>false</b>: The table is not clustered on the index.</li> </ul>                                                                                                                 |
| indisusable    | Boolean  | Specifies whether the index is available for insert and select operations. <ul style="list-style-type: none"> <li>• <b>true</b>: The index is available for insert and select operations.</li> <li>• <b>false</b>: The index is unavailable for insert and select operations.</li> </ul>                                                       |
| indisvalid     | Boolean  | <ul style="list-style-type: none"> <li>• <b>true</b>: The index can be used for query.</li> <li>• <b>false</b>: The index is possibly incomplete and must still be modified by INSERT or UPDATE operations, but it cannot be securely used for queries. If it is a unique index, the uniqueness property is also not <b>true</b>.</li> </ul>   |

| Name           | Type         | Description                                                                                                                                                                                                                                                                                                                                                             |
|----------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| indcheckxmin   | Boolean      | <ul style="list-style-type: none"> <li><b>true:</b> Queries must not use the index until the xmin of this row in PG_INDEX is lower than their <b>TransactionXmin</b>, because the table may contain broken HOT chains with incompatible rows that they can see.</li> <li><b>false:</b> Indexes can be used for query.</li> </ul>                                        |
| indisready     | Boolean      | <ul style="list-style-type: none"> <li><b>true:</b> The index is available for inserting data.</li> <li><b>false:</b> The index is ignored when data is inserted or modified.</li> </ul>                                                                                                                                                                                |
| indkey         | int2vector   | This is an array of <b>indnatts</b> values indicating that this index creates table columns. For example, a value of <b>1 3</b> indicates that the first and the third columns make up the index key. The value <b>0</b> in this array indicates that the corresponding index attribute is an expression over the table columns, rather than a simple column reference. |
| indcollation   | oidvector    | ID of each column used by the index.                                                                                                                                                                                                                                                                                                                                    |
| indclass       | oidvector    | For each column in the index key, this contains the OID of the operator class to use. See <a href="#">PG_OPCLASS</a> for details.                                                                                                                                                                                                                                       |
| indoption      | int2vector   | Array of values that store per-column flag bits. The meaning of the bits is defined by the index's access method.                                                                                                                                                                                                                                                       |
| indexprs       | pg_node_tree | Expression trees (in <b>nodeToString()</b> representation) for index attributes that are not simple column references. It is a list with one element for each zero entry in indkey. The value is null if all index attributes are simple references.                                                                                                                    |
| indpred        | pg_node_tree | Expression tree (in <b>nodeToString()</b> representation) for partial index predicate. If the index is not a partial index, this column is an empty string.                                                                                                                                                                                                             |
| indisreplident | Boolean      | <p>Specifies whether the column of this index is a decoded column of logical decoding.</p> <ul style="list-style-type: none"> <li><b>true:</b> The column of this index is a decoded column of logical decoding.</li> <li><b>false:</b> The column of this index is not a decoded column of logical decoding.</li> </ul>                                                |
| indnkeyatts    | smallint     | Total number of columns in the index. The columns that exceed the value of <b>indnatts</b> are not involved in the index query.                                                                                                                                                                                                                                         |

## 12.2.47 PG\_INHERITS

**PG\_INHERITS** records information about table inheritance hierarchies. There is one entry for each direct child table in the database. Indirect inheritance can be determined by following chains of entries.

**Table 12-48** PG\_INHERITS columns

| Name      | Type    | Reference                    | Description                                                                                                                                                                              |
|-----------|---------|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| inhrelid  | oid     | <a href="#">PG_CLASS.oid</a> | OID of a child table                                                                                                                                                                     |
| inhparent | oid     | <a href="#">PG_CLASS.oid</a> | OID of a parent table                                                                                                                                                                    |
| inhseqno  | integer | -                            | If there is more than one direct parent for a child table (multiple inheritances), this number tells the order in which the inherited columns are to be arranged. The count starts at 1. |

## 12.2.48 PG\_JOB

**PG\_JOB** records detailed information about jobs created by users. Dedicated threads poll the system catalog **PG\_JOB** and trigger jobs based on scheduled job execution time, and update job status in **PG\_JOB**. This system catalog belongs to the Shared Relation category. All job records are visible to all databases.

**Table 12-49** PG\_JOB columns

| Name                 | Type   | Description                                                                                                                                                      |
|----------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid                  | oid    | Row identifier (hidden attribute, which must be specified)                                                                                                       |
| job_id               | bigint | Job ID, primary key, unique (with a unique index)                                                                                                                |
| current_postgres_pid | bigint | If the current job has been executed, the thread ID of this job is recorded. The default value is <b>-1</b> , indicating that the job has not yet been executed. |
| log_user             | name   | Username of the job creator                                                                                                                                      |
| priv_user            | name   | Username of the job executor                                                                                                                                     |
| dbname               | name   | Name of the database in which the job will be executed                                                                                                           |
| node_name            | name   | Primary database node on which the job will be executed                                                                                                          |

| Name            | Type                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| job_status      | "char"                      | <p>Execution status of the current task. The default value is 's'. The options are as follows:</p> <ul style="list-style-type: none"><li>• 'r': running</li><li>• 's': successfully finished</li><li>• 'f': job failed</li><li>• 'd': disable</li></ul> <p>If a job fails to be executed for 16 consecutive times, <b>job_status</b> is automatically set to 'd', and no more attempt will be made on this job.</p> <p>Note: When you disable a scheduled task (by setting the GUC parameter <b>job_queue_processes</b> to 0), the thread that monitors the job execution is not started, and the job status will not be updated. You can ignore this status. Only when the scheduled task function is enabled (the GUC parameter <b>job_queue_processes</b> is not set to 0), the system updates the value of this column based on the real-time job status.</p> |
| start_date      | timestamp without time zone | Start time of the first job execution, accurate to millisecond                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| next_run_date   | timestamp without time zone | Time when a scheduled job is executed next time. The time is accurate to milliseconds.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| failure_count   | smallint                    | Number of times the job has started and failed. If a job fails to be executed for 16 consecutive times, no more attempt will be made on it.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| interval        | text                        | Job execution interval                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| last_start_date | timestamp without time zone | Start time of the last job execution, accurate to millisecond                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| last_end_date   | timestamp without time zone | End time of the last job execution, accurate to millisecond                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| last_suc_date   | timestamp without time zone | Start time of the last successful job execution, accurate to millisecond                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| this_run_date   | timestamp without time zone | Start time of the ongoing job execution, accurate to millisecond                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| nspname         | name                        | Name of the schema used for job execution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| job_name        | text                        | Name of the DBE_SCHEDULER scheduled task.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

| Name            | Type                        | Description                                                                                                                                                                           |
|-----------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| end_date        | timestamp without time zone | Expiration time of the DBE_SCHEDULER scheduled task, accurate to millisecond.                                                                                                         |
| enable          | Boolean                     | The DBE_SCHEDULER scheduled task enabling status. The options are as follows: <ul style="list-style-type: none"><li>• <b>true</b>: enabled</li><li>• <b>false</b>: disabled</li></ul> |
| failure_message | text                        | Error information about the latest task execution.                                                                                                                                    |

## 12.2.49 PG\_JOB\_PROC

PG\_JOB\_PROC records the content of each job in the [PG\\_JOB](#) table, including the PL/SQL code blocks and anonymous blocks. Storing such information in the system catalog PG\_JOB and loading it to the shared memory will result in excessive memory usage. Therefore, such information is stored in a separate table and is retrieved when needed.

**Table 12-50** PG\_JOB\_PROC columns

| Name     | Type    | Description                                                                                    |
|----------|---------|------------------------------------------------------------------------------------------------|
| oid      | oid     | Row identifier (hidden attribute, which must be specified).                                    |
| job_id   | integer | Foreign key, which is joined with <b>job_id</b> in the system catalog <a href="#">PG_JOB</a> . |
| what     | text    | Job content, which is the program content in the DBE_SCHEDULER scheduled task.                 |
| job_name | text    | Name of the DBE_SCHEDULER scheduled task or program.                                           |

## 12.2.50 PG\_LANGUAGE

PG\_LANGUAGE registers programming languages. You can use them and APIs to write functions or stored procedures.

**Table 12-51** PG\_LANGUAGE columns

| Name | Type | Reference | Description                                                 |
|------|------|-----------|-------------------------------------------------------------|
| oid  | oid  | -         | Row identifier (hidden attribute, which must be specified). |



| Name          | Type      | Reference                           | Description                                                                                                                                                                                                                                                                                                                 |
|---------------|-----------|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| lanname       | name      | -                                   | Language name.                                                                                                                                                                                                                                                                                                              |
| lanowner      | oid       | oid in<br><a href="#">PG_AUTHID</a> | Owner of the language.                                                                                                                                                                                                                                                                                                      |
| lanispl       | Boolean   | -                                   | <ul style="list-style-type: none"><li>• <b>true</b>: user-defined language.</li><li>• <b>false</b>: internal language, for example, SQL.</li></ul> Currently, <b>gs_dump</b> still uses this column to determine which languages need to be dumped, but this might be replaced by a different mechanism in the future.      |
| lanpltrusted  | Boolean   | -                                   | <ul style="list-style-type: none"><li>• <b>true</b>: The language is trusted, which means that it is believed not to grant access to anything outside the normal SQL execution environment.</li><li>• <b>false</b>: The language is untrusted. Only the initial user can create functions in untrusted languages.</li></ul> |
| lanplcallfoid | oid       | oid in<br><a href="#">PG_PROC</a>   | For non-internal languages, this column references the language handler, which is a special function responsible for executing all functions that are written in the particular language.                                                                                                                                   |
| laninline     | oid       | oid in<br><a href="#">PG_PROC</a>   | This column references a function responsible for executing "inline" anonymous code blocks (DO blocks). The value is <b>0</b> if inline blocks are not supported.                                                                                                                                                           |
| lanvalidator  | oid       | oid in<br><a href="#">PG_PROC</a>   | This column references a language validator function responsible for checking the syntax and validity of new functions when they are created. The value is <b>0</b> if no validator is provided.                                                                                                                            |
| lanacl        | aclitem[] | -                                   | Access permissions                                                                                                                                                                                                                                                                                                          |

## 12.2.51 PG\_LARGEOBJECT

**PG\_LARGEOBJECT** records data making up large objects. A large object is identified by an OID assigned when it is created. Each large object is broken into segments or "pages" small enough to be conveniently stored as rows in **PG\_LARGEOBJECT**. The amount of data per page is defined as **LOBLKSIZE**.

This system catalog is accessible only to users with the SYSADMIN permission.

**Table 12-52** PG\_LARGEOBJECT columns

| Name   | Type    | Reference                                   | Description                                                                                             |
|--------|---------|---------------------------------------------|---------------------------------------------------------------------------------------------------------|
| loid   | oid     | <a href="#">PG_LARGEOBJECT_METADATA.oid</a> | Identifier of the large object that includes this page                                                  |
| pageno | integer | -                                           | Page number of this page within its large object (counting from zero)                                   |
| data   | bytea   | -                                           | Data stored in the large object. This will never be more than <b>LOBLKSIZE</b> bytes and might be less. |

Each row of **PG\_LARGEOBJECT** holds data for one page of a large object, beginning at byte offset (**pageno** \* **LOBLKSIZE**) within the object. The implementation allows sparse storage: pages might be missing, and might be shorter than **LOBLKSIZE** bytes even if they are not the last page of the object. Missing regions within a large object read as zeroes.

## 12.2.52 PG\_LARGEOBJECT\_METADATA

**PG\_LARGEOBJECT\_METADATA** records metadata associated with large objects. The actual large object data is stored in **PG\_LARGEOBJECT**.

**Table 12-53** PG\_LARGEOBJECT\_METADATA columns

| Name     | Type      | Reference                     | Description                                                |
|----------|-----------|-------------------------------|------------------------------------------------------------|
| oid      | oid       | -                             | Row identifier (hidden attribute, which must be specified) |
| lomowner | oid       | <a href="#">PG_AUTHID.oid</a> | Owner of the large object                                  |
| lomacl   | aclitem[] | -                             | Access permissions                                         |

## 12.2.53 PG\_NAMESPACE

**PG\_NAMESPACE** records namespaces, that is, schema-related information.

**Table 12-54** PG\_NAMESPACE columns

| Name | Type | Description                                                |
|------|------|------------------------------------------------------------|
| oid  | oid  | Row identifier (hidden attribute, which must be specified) |

| Name              | Type      | Description                                                                                                                                                                              |
|-------------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| nspname           | name      | Name of a namespace                                                                                                                                                                      |
| nspowner          | oid       | Owner of a namespace                                                                                                                                                                     |
| nsptimeline       | bigint    | Timeline when a namespace is created on the database node. This column is for internal use and valid only on the database node.                                                          |
| nspacl            | aclitem[] | Access permission                                                                                                                                                                        |
| in_redistribution | "char"    | Specifies whether the content is in the redistribution state.                                                                                                                            |
| nspblockchain     | boolean   | <ul style="list-style-type: none"><li>If the value is <b>true</b>, the tamper-proof schema is used.</li><li>If the value is <b>false</b>, the non-tamper-proof schema is used.</li></ul> |

## 12.2.54 PG\_OBJECT

**PG\_OBJECT** records the creator, creation time, and last modification time of objects of specified types (ordinary tables, indexes, sequences, views, stored procedures, and functions).

**Table 12-55** PG\_OBJECT columns

| Name        | Type                     | Description                                                                                                                                                                                                    |
|-------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| object_oid  | oid                      | Object identifier                                                                                                                                                                                              |
| object_type | "char"                   | Object type <ul style="list-style-type: none"><li><b>r</b>: ordinary table</li><li><b>i</b>: index</li><li><b>s</b>: sequence</li><li><b>v</b>: view</li><li><b>p</b>: stored procedure and function</li></ul> |
| creator     | oid                      | ID of a creator                                                                                                                                                                                                |
| ctime       | timestamp with time zone | Creation time of an object                                                                                                                                                                                     |
| mtime       | timestamp with time zone | Last modification time of an object. The modification operations include <b>ALTER</b> , <b>GRANT</b> , and <b>REVOKE</b> .                                                                                     |
| createcsn   | bigint                   | CSN when an object is created.                                                                                                                                                                                 |

| Name      | Type   | Description                                                   |
|-----------|--------|---------------------------------------------------------------|
| changeagn | bigint | CSN when DDL operations are performed on a table or an index. |

**NOTICE**

- Objects created or modified during database initialization (initdb) cannot be recorded. **PG\_OBJECT** does not contain these object records.
- When an object created before the upgrade is modified again, the modification time (specified by **mtime**) is recorded. When DDL operations are performed on a table or an index, the transaction commit sequence number (specified by **changeagn**) of the transaction to which the table or index belongs is recorded. Because the creation time of the object cannot be obtained, **ctime** and **createagn** are empty.
- The time recorded by **ctime** and **mtime** is the start time of the transaction to which the current operation belongs.
- The time of object modification due to capacity expansion is also recorded.
- **createagn** and **changeagn** record the transaction commit sequence number of the transaction to which the current operation belongs.
- When **enable\_gtt\_concurrent\_truncate** is set to **on**, the **mtime** field is not updated when the global temporary table is truncated.

## 12.2.55 PG\_OPCLASS

**PG\_OPCLASS** defines index access method operator classes.

Each operator class defines semantics for index columns of a particular data type and a particular index access method. An operator class essentially specifies that a particular operator family is applicable to a particular indexable column data type. The set of operators from the family that are actually usable with the indexed column are whichever ones accept the column's data type as their left-hand input.

**Table 12-56** PG\_OPCLASS columns

| Name         | Type | Reference                           | Description                                                    |
|--------------|------|-------------------------------------|----------------------------------------------------------------|
| oid          | oid  | -                                   | Row identifier (hidden attribute, which must be specified)     |
| opcmeth      | oid  | OID in <a href="#">PG_AM</a>        | Index access method operator class served by an operator class |
| opcname      | name | -                                   | Name of the operator class                                     |
| opcnamespace | oid  | OID in <a href="#">PG_NAMESPACE</a> | Namespace of the operator class                                |
| opcowner     | oid  | OID in <a href="#">PG_AUTHID</a>    | Owner of the operator class                                    |

| Name       | Type    | Reference                          | Description                                                                           |
|------------|---------|------------------------------------|---------------------------------------------------------------------------------------|
| opcfamily  | oid     | OID in <a href="#">PG_OPFAMILY</a> | Operator family containing the operator class                                         |
| opcintype  | oid     | OID in <a href="#">PG_TYPE</a>     | Data type that the operator class indexes                                             |
| opcdefault | boolean | -                                  | The value is <b>true</b> if this operator class is the default for <b>opcintype</b> . |
| opckeytype | oid     | OID in <a href="#">PG_TYPE</a>     | Type of data stored in an index, or zero if same as <b>opcintype</b>                  |

An operator class's **opcmethod** must match the **opfmetho**d of its containing operator family.

## 12.2.56 PG\_OPERATOR

PG\_OPERATOR records information about operators.

**Table 12-57** PG\_OPERATOR columns

| Name         | Type    | Reference                           | Description                                                                                                                                                |
|--------------|---------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid          | oid     | -                                   | Row identifier (hidden attribute, which must be specified)                                                                                                 |
| oprname      | name    | -                                   | Name of an operator                                                                                                                                        |
| oprnamespace | oid     | oid in <a href="#">PG_NAMESPACE</a> | OID of the namespace that contains the operator.                                                                                                           |
| oprowner     | oid     | oid in <a href="#">PG_AUTHID</a>    | Owner of the operator.                                                                                                                                     |
| oprkind      | "char"  | -                                   | <ul style="list-style-type: none"><li>• <b>b</b>: infix ("both").</li><li>• <b>l</b>: prefix ("left").</li><li>• <b>r</b>: postfix ("right").</li></ul>    |
| oprcanmerge  | Boolean | -                                   | Specifies whether the operator supports merge joins. <ul style="list-style-type: none"><li>• <b>t (true)</b>: yes</li><li>• <b>f (false)</b>: no</li></ul> |

| Name       | Type    | Reference                          | Description                                                                                                                                               |
|------------|---------|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| oprcanhash | Boolean | -                                  | Specifies whether the operator supports hash joins. <ul style="list-style-type: none"><li>• <b>t (true)</b>: yes</li><li>• <b>f (false)</b>: no</li></ul> |
| oprleft    | oid     | oid in <a href="#">PG_TYPE</a>     | Type of the left operand.                                                                                                                                 |
| oprright   | oid     | oid in <a href="#">PG_TYPE</a>     | Type of the right operand.                                                                                                                                |
| oprresult  | oid     | oid in <a href="#">PG_TYPE</a>     | Type of the result.                                                                                                                                       |
| oprcom     | oid     | oid in <a href="#">PG_OPERATOR</a> | If it exists, the value is the exchange character of this operator. If it does not exist, the value is <b>0</b> .                                         |
| oprnegate  | oid     | oid in <a href="#">PG_OPERATOR</a> | If it exists, the value is the invertor of this operator. If it does not exist, the value is <b>0</b> .                                                   |
| oprcode    | regproc | prname in <a href="#">PG_PROC</a>  | Function that implements the operator                                                                                                                     |
| oprrest    | regproc | prname in <a href="#">PG_PROC</a>  | Restriction selectivity estimation function for the operator.                                                                                             |
| oprjoin    | regproc | prname in <a href="#">PG_PROC</a>  | Join selectivity estimation function for the operator.                                                                                                    |

## 12.2.57 PG\_OPFAMILY

**PG\_OPFAMILY** defines operator families.

Each operator family is a collection of operators and associated support routines that implement semantics specified for a particular index access method. Furthermore, the operators in a family are all compatible, in a way that is specified by the access method. The operator family allows cross-data-type operators to be used with indexes and to be reasoned about using knowledge of access method semantics.

**Table 12-58** PG\_OPFAMILY columns

| Name | Type | Reference | Description                                                |
|------|------|-----------|------------------------------------------------------------|
| oid  | oid  | -         | Row identifier (hidden attribute, which must be specified) |

| Name         | Type | Reference                        | Description                                    |
|--------------|------|----------------------------------|------------------------------------------------|
| opfmethod    | oid  | <a href="#">PG_AM.oid</a>        | Index access method used by an operator family |
| opfname      | name | -                                | Name of the operator family                    |
| opfnamespace | oid  | <a href="#">PG_NAMESPACE.oid</a> | Namespace of the operator family               |
| opfowner     | oid  | <a href="#">PG_AUTHID.oid</a>    | Owner of the operator family                   |

The majority of the information defining an operator family is not in its **PG\_OPFAMILY** row, but in the associated rows in [PG\\_AMOP](#), [PG\\_AMPROC](#), and [PG\\_OPCLASS](#).

## 12.2.58 PG\_PARTITION

PG\_PARTITION records all partitioned tables, table partitions, TOAST tables on table partitions, and index partitions in the database. Partitioned index information is not stored in the system catalog PG\_PARTITION.

**Table 12-59** PG\_PARTITION columns

| Name        | Type    | Description                                                                                                                                                                                                                                |
|-------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid         | oid     | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                |
| relname     | name    | Names of the partitioned tables, table partitions, TOAST tables on table partitions, and index partitions.                                                                                                                                 |
| parttype    | "char"  | Object type. <ul style="list-style-type: none"><li>● <b>r</b>: partitioned table</li><li>● <b>p</b>: table partition</li><li>● <b>s</b>: table subpartition</li><li>● <b>x</b>: index partition</li><li>● <b>t</b>: TOAST table.</li></ul> |
| parentid    | oid     | OID of the partitioned table in PG_CLASS when the object is a partitioned table or a partition.<br>OID of the partitioned index when the object is an index partition.                                                                     |
| rangenum    | integer | Reserved column.                                                                                                                                                                                                                           |
| intervalnum | integer | Reserved column.                                                                                                                                                                                                                           |

| Name           | Type             | Description                                                                                                                                                                                                                                                                                                    |
|----------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| partstrategy   | "char"           | Partition policy of the partitioned table <ul style="list-style-type: none"><li>• <b>r</b>: range partition</li><li>• <b>v</b>: numeric partition</li><li>• <b>i</b>: interval partition</li><li>• <b>l</b>: list partition</li><li>• <b>h</b>: hash partition</li><li>• <b>n</b>: invalid partition</li></ul> |
| relfilenode    | oid              | Physical storage locations of the table partition, index partition, and TOAST table on the table partition                                                                                                                                                                                                     |
| reltablespace  | oid              | OID of the tablespace containing the table partition, index partition, and TOAST table on the table partition                                                                                                                                                                                                  |
| relpages       | double precision | Statistics: numbers of data pages of the table partition and index partition                                                                                                                                                                                                                                   |
| reltuples      | double precision | Statistics: numbers of tuples of the table partition and index partition                                                                                                                                                                                                                                       |
| relallvisible  | integer          | Statistics: number of visible data pages of the table partition and index partition                                                                                                                                                                                                                            |
| reltoastrelid  | oid              | OID of the TOAST table corresponding to the table partition                                                                                                                                                                                                                                                    |
| reltoastidxid  | oid              | OID of the TOAST table index corresponding to the table partition                                                                                                                                                                                                                                              |
| indextblid     | oid              | OID of the table partition corresponding to the index partition                                                                                                                                                                                                                                                |
| indisusable    | boolean          | Specifies whether the index partition is available.                                                                                                                                                                                                                                                            |
| reldeltarelid  | oid              | OID of a Delta table                                                                                                                                                                                                                                                                                           |
| reldeltaidx    | oid              | OID of the index for a Delta table                                                                                                                                                                                                                                                                             |
| relcudescrelid | oid              | OID of a CU description table                                                                                                                                                                                                                                                                                  |
| relcudescidx   | oid              | OID of the index for a CU description table                                                                                                                                                                                                                                                                    |
| relfrozenxid   | xid32            | Frozen transaction ID<br>To ensure forward compatibility, this column is reserved. The <b>relfrozenxid64</b> column is added to record the information.                                                                                                                                                        |
| intspnum       | integer          | Number of tablespaces that the interval partition belongs to                                                                                                                                                                                                                                                   |



| Name               | Type       | Description                                                                                                                                                                        |
|--------------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| partkey            | int2vector | Column number of the partition key                                                                                                                                                 |
| intervaltablespace | oidvector  | Tablespace that the interval partition belongs to. Interval partitions fall in the tablespaces in the round-robin manner.                                                          |
| interval           | text[]     | Interval value of the interval partition                                                                                                                                           |
| boundaries         | text[]     | Upper boundary of the range partition and interval partition                                                                                                                       |
| transit            | text[]     | Transit of the interval partition                                                                                                                                                  |
| reloptions         | text[]     | Storage property of a partition used for collecting online scale-out information. Same as <b>pg_class.reloptions</b> , it is expressed in a string in the format of keyword=value. |
| relfrozenxid64     | xid        | Frozen transaction ID                                                                                                                                                              |
| relminmxid         | xid        | Frozen multi-transaction ID                                                                                                                                                        |

## 12.2.59 PG\_PLTEMPLATE

PG\_PLTEMPLATE records template information for procedural languages.

**Table 12-60** PG\_PLTEMPLATE columns

| Name          | Type    | Description                                                                                                              |
|---------------|---------|--------------------------------------------------------------------------------------------------------------------------|
| tmplname      | name    | Name of the language for which this template is used.                                                                    |
| tmpltrusted   | Boolean | The value is <b>true</b> if the language is considered trusted. Otherwise, the value is <b>false</b> .                   |
| tmpldbcreate  | Boolean | The value is <b>true</b> if the language is created by the owner of the database. Otherwise, the value is <b>false</b> . |
| tmplhandler   | text    | Name of the call handler function                                                                                        |
| tmpline       | text    | Name of the anonymous block handler ( <b>NULL</b> if no name of the block handler exists).                               |
| tmplvalidator | text    | Name of the verification function ( <b>NULL</b> if no verification function is available).                               |
| tmplib        | text    | Path of the shared library that implements languages.                                                                    |

| Name    | Type      | Description                                     |
|---------|-----------|-------------------------------------------------|
| tmplacl | aclitem[] | Access permissions for template (not yet used). |

## 12.2.60 PG\_PROC

PG\_PROC records information about functions or procedures.

**Table 12-61** PG\_PROC columns

| Name         | Type    | Description                                                                                                                                                                        |
|--------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid          | oid     | Row identifier (hidden attribute, which must be specified).                                                                                                                        |
| proname      | name    | Function name.                                                                                                                                                                     |
| pronamespace | oid     | OID of the namespace that contains the function.                                                                                                                                   |
| proowner     | oid     | Owner of the function.                                                                                                                                                             |
| prolang      | oid     | Implementation language or call API of the function.                                                                                                                               |
| procost      | real    | Estimated execution cost.                                                                                                                                                          |
| prorows      | real    | Estimated number of rows that are influenced.                                                                                                                                      |
| provariadic  | oid     | Data type of parameter element.                                                                                                                                                    |
| protransform | regproc | Simplified call method for the function.                                                                                                                                           |
| proisagg     | Boolean | Specifies whether the function is an aggregate function. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                     |
| proiswindow  | Boolean | Specifies whether the function is a window function. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                         |
| prosecdef    | Boolean | Specifies whether the function is a security definer (or a setuid function). <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul> |

| Name            | Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| proleakproof    | Boolean   | Specifies whether the function has side effects. If no leakproof treatment is provided for parameters, the function throws errors. <ul style="list-style-type: none"> <li>• <b>t</b> (true): There is no side effect.</li> <li>• <b>f</b> (false): There are side effects.</li> </ul>                                                                                                                                                                                                                                                                                        |
| proisstrict     | Boolean   | <ul style="list-style-type: none"> <li>• <b>t</b> (true): When this function is used, the function does not call the function and returns null if the input parameter is null.</li> <li>• <b>f</b> (false): When this function is used, it is called even if the input parameter is null. Therefore, this function must process null input.</li> </ul>                                                                                                                                                                                                                       |
| proretset       | Boolean   | Specifies whether the return value of the function is a set (that is, multiple values of the specified data type). <ul style="list-style-type: none"> <li>• <b>true</b>: The return value of the function is a set.</li> <li>• <b>false</b>: The return value of the function is not a set.</li> </ul>                                                                                                                                                                                                                                                                       |
| provolatile     | "char"    | Specifies whether the function's result depends only on its input parameters, or is affected by outside factors. <ul style="list-style-type: none"> <li>• <b>i</b>: immutable functions, which always deliver the same result for the same inputs.</li> <li>• <b>s</b>: stable functions, whose results (for fixed inputs) do not change within a scan.</li> <li>• <b>v</b>: volatile functions, whose results may change at any time. It is also used for functions with side-effects, so that the engine cannot get optimized if volatile functions are called.</li> </ul> |
| pronargs        | smallint  | Number of parameters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| pronargdefaults | smallint  | Number of parameters that have default values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| prorettype      | oid       | Data type of return values.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| proargtypes     | oidvector | Array that stores the data types of function parameters. This array includes only input parameters (including <b>INOUT</b> parameters), and indicates the call signature (API) of the function.                                                                                                                                                                                                                                                                                                                                                                              |

| Name             | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| proallargtypes   | oid[]        | Array that contains the data types of function parameters. This array includes all parameter types (including <b>OUT</b> and <b>INOUT</b> parameters); however, if all the parameters are <b>IN</b> parameters, this column is null. Note that array subscripting is 1-based, whereas for historical reasons, <b>proargtypes</b> is subscripted from 0.                                                                                                                                       |
| proargmodes      | "char"[]     | Array with the modes of the function parameters, encoded as follows: <ul style="list-style-type: none"> <li>• <b>i</b> indicates the IN parameter.</li> <li>• <b>o</b> indicates the OUT parameter.</li> <li>• <b>b</b> indicates the INOUT parameter.</li> <li>• <b>v</b> indicates the VARIADIC parameter.</li> </ul> If all the parameters are <b>IN</b> parameters, this column is null. Note that subscripts correspond to positions of <b>proallargtypes</b> , not <b>proargtypes</b> . |
| proargnames      | text[]       | Array that stores the names of the function parameters. Parameters without a name are set to empty strings in the array. If none of the parameters have a name, this column is null. Note that subscripts correspond to positions of <b>proallargtypes</b> , not <b>proargtypes</b> .                                                                                                                                                                                                         |
| proargdefaults   | pg_node_tree | Expression tree of the default value. This is the list of <b>pronargdefaults</b> elements.                                                                                                                                                                                                                                                                                                                                                                                                    |
| prosrc           | text         | A definition that describes a function or stored procedure. In an interpreting language, it is the function source code, a connection symbol, a file name, or any body content specified when a function or stored procedure is created, depending on how a language or call is used.                                                                                                                                                                                                         |
| probin           | text         | Additional information about how to call the function. Again, the interpretation is language-specific.                                                                                                                                                                                                                                                                                                                                                                                        |
| proconfig        | text[]       | Function's local settings for run-time configuration variables.                                                                                                                                                                                                                                                                                                                                                                                                                               |
| proacl           | aclitem[]    | Access permissions. For details, see <a href="#">GRANT</a> and <a href="#">REVOKE</a> .                                                                                                                                                                                                                                                                                                                                                                                                       |
| prodefaultargpos | int2vector   | Position of the input parameter of a function with a default value.                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Name           | Type                 | Description                                                                                                                                                                                                                                                                                                                                                                            |
|----------------|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencedmode     | Boolean              | Execution mode of a function, indicating whether the function is executed in fence or not fence mode. <ul style="list-style-type: none"><li>• <b>true</b>: fence mode. In this mode, the function is executed in the fork process.</li><li>• <b>false</b>: not fence mode.</li></ul> For built-in functions, <b>fencedmode</b> is set to <b>false</b> , indicating the not fence mode. |
| proshippable   | Boolean              | Specifies whether the function can be pushed down to database nodes. The default value is <b>false</b> . <ul style="list-style-type: none"><li>• Functions of the IMMUTABLE type can always be pushed down to the database nodes.</li><li>• A STABLE or VOLATILE function can be pushed down to the database nodes only if SHIPPABLE is specified for it.</li></ul>                    |
| propackage     | Boolean              | Specifies whether the function supports overloading. The default value is <b>false</b> . <ul style="list-style-type: none"><li>• <b>t</b> (true): supported.</li><li>• <b>f</b> (false): not supported.</li></ul>                                                                                                                                                                      |
| prokind        | "char"               | Specifies whether the object is a function or a stored procedure. <ul style="list-style-type: none"><li>• <b>'f'</b>: The object is a function.</li><li>• <b>'p'</b>: The object is a stored procedure.</li></ul>                                                                                                                                                                      |
| proargsrc      | text                 | Describes the parameter input strings of functions or stored procedures that are A-compatible syntax, including parameter comments. The default value is <b>NULL</b> .                                                                                                                                                                                                                 |
| proisprivate   | Boolean              | Specifies whether a function is a private function in the package. The default value is <b>false</b> .                                                                                                                                                                                                                                                                                 |
| propackageid   | oid                  | OID of the package to which the function belongs. If the function is not in the package, the value is <b>0</b> .                                                                                                                                                                                                                                                                       |
| proargtypesext | oidvector<br>_extend | Data type array used to store function parameters when there are a large number of function parameters. This array includes only input parameters (including <b>INOUT</b> parameters), and indicates the call signature (API) of the function.                                                                                                                                         |

| Name                | Type              | Description                                                                                                                                                                                                                                 |
|---------------------|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| prodefaultargposext | int2vector_extend | Position of the input parameter with a default value when the function has a large number of parameters.                                                                                                                                    |
| allargtypes         | oidvector         | Array for storing the data types of stored procedure parameters, including all parameters (input parameters, output parameters, and <b>INOUT</b> parameters) of the stored procedure.                                                       |
| allargtypesext      | oidvector_extend  | Array for storing the data types of stored procedure parameters when the number of function parameters is greater than 666. The array contains all parameters (including input parameters, output parameters, and <b>INOUT</b> parameters). |

## 12.2.61 PG\_RANGE

PG\_RANGE records information about range types, except for records of types in [PG\\_TYPE](#).

**Table 12-62** PG\_RANGE columns

| Name         | Type    | Reference                                  | Description                                                                                                                                                                                                                           |
|--------------|---------|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rngtypid     | oid     | <b>oid</b> in <a href="#">PG_TYPE</a>      | OID of the range type.                                                                                                                                                                                                                |
| rngsubtype   | oid     | <b>oid</b> in <a href="#">PG_TYPE</a>      | OID of the element type (subtype) of this range type.                                                                                                                                                                                 |
| rngcollation | oid     | <b>oid</b> in <a href="#">PG_COLLATION</a> | OID of the collation used for range comparisons ( <b>0</b> if none).                                                                                                                                                                  |
| rngsubopc    | oid     | <b>oid</b> in <a href="#">PG_OPCLASS</a>   | OID of the subtype's operator class used for range comparisons.                                                                                                                                                                       |
| rngcanonical | regproc | <b>prname</b> in <a href="#">PG_PROC</a>   | Name of the function to convert a range value into canonical form ( <b>0</b> if none).                                                                                                                                                |
| rngsubdiff   | regproc | <b>prname</b> in <a href="#">PG_PROC</a>   | Name of the function used to calculate the difference between two elements in a range. The return value of the function is of the double-precision type. If the function does not exist, the value of <b>rngsubdiff</b> is <b>0</b> . |

If the element type is discrete, **rngcanonical** determines the collation used for the range type. If the element type is not collatable, **rngsubopc** determines the collation used for the range type. If the element type is collatable, **rngsubopc** and **rngcollation** determine the collation used for the range type.

## 12.2.62 PG\_REPLICATION\_ORIGIN

PG\_REPLICATION\_ORIGIN contains all created replication sources. This catalog is shared among all databases of a database instance. That is, there is only one copy for each instance, not for each database.

**Table 12-63** PG\_REPLICATION\_ORIGIN columns

| Name    | Type | Description                                                      |
|---------|------|------------------------------------------------------------------|
| roident | oid  | Unique replication source identifier within a database instance. |
| roname  | text | External user-defined replication source name.                   |

## 12.2.63 PG\_RESOURCE\_POOL

PG\_RESOURCE\_POOL provides information about database resource pools.

**Table 12-64** PG\_RESOURCE\_POOL columns

| Name              | Type    | Description                                                                               |
|-------------------|---------|-------------------------------------------------------------------------------------------|
| oid               | oid     | Row identifier (hidden attribute, which must be specified).                               |
| respool_name      | name    | Name of the resource pool.                                                                |
| mem_percent       | integer | Percentage of the memory configuration.                                                   |
| cpu_affinity      | bigint  | Value of cores bound to the CPU.                                                          |
| control_group     | name    | Name of the Cgroup where the resource pool is located.                                    |
| active_statements | integer | Maximum number of concurrent statements in the resource pool.                             |
| max_dop           | integer | Maximum scanning concurrency during data redistribution. This column is used for scaling. |
| memory_limit      | name    | Maximum memory of the resource pool.                                                      |
| parentid          | oid     | OID of the parent resource pool.                                                          |
| io_limits         | integer | Upper limit of IOPS. It is counted by 10,000 per second for row store.                    |

| Name        | Type    | Description                                                                                                                                                                                                                    |
|-------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| io_priority | name    | I/O priority set for jobs that consume many I/O resources. It takes effect when the I/O usage reaches 90%.                                                                                                                     |
| nodegroup   | name    | Name of the logical database to which the resource pool belongs.                                                                                                                                                               |
| is_foreign  | Boolean | Specifies whether the resource pool can be used for users outside the logical database. If it is set to <b>true</b> , the resource pool controls the resources of common users who do not belong to the current resource pool. |
| max_worker  | integer | Concurrency in a table during data redistribution. This column is used only for scaling.                                                                                                                                       |

Note: **max\_dop** and **max\_worker** are used for scaling and are not applicable to the centralized deployment.

## 12.2.64 PG\_REWRITE

PG\_REWRITE records rewrite rules defined for tables and views.

**Table 12-65** PG\_REWRITE columns

| Name     | Type     | Description                                                                                                                                                                  |
|----------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid      | oid      | Row identifier (hidden attribute, which must be specified)                                                                                                                   |
| rulename | name     | Rule name                                                                                                                                                                    |
| ev_class | oid      | Name of the table that uses the rule                                                                                                                                         |
| ev_attr  | smallint | Column to which this rule applies (always <b>0</b> to indicate the entire table)                                                                                             |
| ev_type  | "char"   | Event type for the rule. <ul style="list-style-type: none"><li>• <b>1</b>: SELECT</li><li>• <b>2</b>: UPDATE</li><li>• <b>3</b>: INSERT</li><li>• <b>4</b>: DELETE</li></ul> |



| Name       | Type         | Description                                                                                                                                                                                                                                                                                                    |
|------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ev_enabled | "char"       | Controls the mode in which the rule is triggered. <ul style="list-style-type: none"><li>• <b>O</b>: The rule is triggered in origin and local modes.</li><li>• <b>D</b>: The rule is disabled.</li><li>• <b>R</b>: The rule is triggered in replica mode.</li><li>• <b>A</b>: The rule always fires.</li></ul> |
| is_instead | Boolean      | The value is <b>true</b> if the rule is of the <b>INSTEAD</b> type. Otherwise, the value is <b>false</b> .                                                                                                                                                                                                     |
| ev_qual    | pg_node_tree | Expression tree (in the form of a nodeToString() representation) for the rule's qualifying condition.                                                                                                                                                                                                          |
| ev_action  | pg_node_tree | Query tree (in the form of a nodeToString() representation) for the rule's action.                                                                                                                                                                                                                             |

## 12.2.65 PG\_RLSPOLICY

PG\_RLSPOLICY records row-level security policies.

**Table 12-66** PG\_RLSPOLICY columns

| Name          | Type         | Description                                                                                                                                                                                                        |
|---------------|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid           | oid          | Row identifier (hidden attribute, which must be specified).                                                                                                                                                        |
| polname       | name         | Name of an access control policy.                                                                                                                                                                                  |
| polrelid      | oid          | OID of the table object on which the row-level security policy takes effect.                                                                                                                                       |
| polcmd        | "char"       | SQL operations affected by a row-level security policy.                                                                                                                                                            |
| polpermissive | Boolean      | Attribute of the row-level security policy. <ul style="list-style-type: none"><li>• <b>t</b>: OR condition concatenation of expressions.</li><li>• <b>f</b>: AND condition concatenation of expressions.</li></ul> |
| polroles      | oid[]        | OID list of users affected by the row-level security policy. If this parameter is not specified, all users are affected.                                                                                           |
| polqual       | pg_node_tree | Expression of a row-level security policy.                                                                                                                                                                         |

## 12.2.66 PG\_SECLABEL

PG\_SECLABEL records security labels on database objects.

See also [PG\\_SHSECLABEL](#), which provides a similar function for security labels of database objects that are shared within one database.

**Table 12-67** PG\_SECLABEL columns

| Name     | Type    | Reference                              | Description                                            |
|----------|---------|----------------------------------------|--------------------------------------------------------|
| objoid   | oid     | Any OID column                         | OID of the object that this security label pertains to |
| classoid | oid     | <b>oid</b> in <a href="#">PG_CLASS</a> | OID of the system catalog where the object appears     |
| objsubid | integer | -                                      | Column number for a security label on a table column   |
| provider | text    | -                                      | Label provider associated with the label               |
| label    | text    | -                                      | Security label applied to the object                   |

## 12.2.67 PG\_SHDEPEND

PG\_SHDEPEND records the dependency between database objects and shared objects, such as roles. Based on this information, GaussDB can ensure that those objects are unreferenced before attempting to delete them.

See also [PG\\_DEPEND](#), which provides a similar function for dependencies involving objects within a single database.

PG\_SHDEPEND is shared among all databases of a database instance. That is, there is only one PG\_SHDEPEND for each instance, not for each database.

**Table 12-68** PG\_SHDEPEND columns

| Name    | Type | Reference                                 | Description                                                                     |
|---------|------|-------------------------------------------|---------------------------------------------------------------------------------|
| dbid    | oid  | <b>oid</b> in <a href="#">PG_DATABASE</a> | OID of the database where a dependent object is ( <b>0</b> for a shared object) |
| classid | oid  | <b>oid</b> in <a href="#">PG_CLASS</a>    | OID of the system catalog where a dependent object resides                      |
| objid   | oid  | Any OID column                            | OID of the dependent object                                                     |

| Name       | Type    | Reference                     | Description                                                                                                                                       |
|------------|---------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| objsubid   | integer | -                             | Column number for a table column. ( <b>objid</b> and <b>classid</b> refer to the table itself.) The value is <b>0</b> for all other object types. |
| refclassid | oid     | <b>oid</b> in <b>PG_CLASS</b> | OID of the system catalog (must be a shared catalog) where a referenced object is located.                                                        |
| refobjid   | oid     | Any OID column                | OID of the referenced object                                                                                                                      |
| deptype    | "char"  | -                             | Code segment defining the specific semantics of this dependency relationship. See the following for details.                                      |

In all cases, a PG\_SHDEPEND entry indicates that the referenced object cannot be dropped without also dropping the dependent object. However, there are several subflavors identified by **deptype**:

- SHARED\_DEPENDENCY\_OWNER (o)  
The referenced object (which must be a role) is the owner of the dependent object.
- SHARED\_DEPENDENCY\_ACL (a)  
The referenced object (which must be a role) is mentioned in the access control list (ACL) of the dependent object. SHARED\_DEPENDENCY\_ACL does not add a record to the owner of the object because the owner will have a SHARED\_DEPENDENCY\_OWNER record.
- SHARED\_DEPENDENCY\_PIN (p)  
This type of record indicates that the system itself depends on the depended object. Therefore, such an object cannot be deleted. Entries of this type are created only by **initdb**. The columns for the dependent object contain zeroes.
- SHARED\_DEPENDENCY\_DBPRIV(d)  
The referenced object (must be a role) has the ANY permission on the dependent object (the specified OID of the dependent object corresponds to a row in the **GS\_DB\_PRIVILEGE** system catalog).

## 12.2.68 PG\_SHDESCRIPTION

PG\_SHDESCRIPTION records optional comments for shared database objects. Descriptions can be manipulated with the **COMMENT** command and viewed with the **\d** command.

See also PG\_DESCRIPTION, which provides a similar function for descriptions involving objects within a single database.

PG\_SHDESCRIPTION is shared among all databases of a database instance. That is, there is only one PG\_SHDESCRIPTION for each instance, not for each database.

**Table 12-69** PG\_SHDESCRIPTION columns

| Name        | Type | Reference                              | Description                                         |
|-------------|------|----------------------------------------|-----------------------------------------------------|
| objoid      | oid  | Any OID column                         | OID of the described object.                        |
| classoid    | oid  | <b>oid</b> in <a href="#">PG_CLASS</a> | OID of the system catalog where the object appears. |
| description | text | -                                      | Description of the object.                          |

## 12.2.69 PG\_SHSECLABEL

**PG\_SHSECLABEL** records security labels on shared database objects. Security labels can be manipulated with the **SECURITY LABEL** command.

For an easier way to view security labels, see [PG\\_SECLABELS](#).

See also [PG\\_SECLABEL](#), which provides a similar function for security labels involving objects within a single database.

Unlike most system catalogs, **PG\_SHSECLABEL** is shared across all databases in the system. There is only one copy of **PG\_SHSECLABEL** in the GaussDB system, not one per database.

**Table 12-70** PG\_SHSECLABEL columns

| Name     | Type | Reference                    | Description                                            |
|----------|------|------------------------------|--------------------------------------------------------|
| objoid   | oid  | Any OID column               | OID of the object that this security label pertains to |
| classoid | oid  | <a href="#">PG_CLASS.oid</a> | OID of the system catalog where the object appears     |
| provider | text | -                            | Label provider associated with the label               |
| label    | text | -                            | Security label applied to the object                   |

## 12.2.70 PG\_STATISTIC

**PG\_STATISTIC** stores statistics about tables and index columns in a database. By default, only users with the **SYSADMIN** permission can access this system catalog. Common users can access it only after being authorized.

**Table 12-71** PG\_STATISTIC columns

| Name     | Type | Description                                          |
|----------|------|------------------------------------------------------|
| starelid | oid  | Table or index that the described column belongs to. |

| Name            | Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| starelkind      | "char"   | Type of an object.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| staattnum       | smallint | Number of the described column in the table, starting from 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| stainherit      | Boolean  | Specifies whether to collect statistics for objects that have inheritance relationship.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| stanullfrac     | real     | Percentage of column entries that are null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| stawidth        | integer  | Average stored width, in bytes, of non-null entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| stadistinct     | real     | <p>Number of distinct, non-null data values in the column for database nodes.</p> <ul style="list-style-type: none"> <li>A value greater than 0 is the actual number of distinct values.</li> <li>A value less than 0 is the ratio of the <b>distinct</b> value to the total number of rows. For example, if <b>stadistinct</b> is <b>-0.5</b>, the actual <b>distinct</b> value is the total number of rows multiplied by 0.5.</li> <li>The value <b>0</b> indicates that the number of distinct values is unknown.</li> </ul> |
| stakindN        | smallint | <p>Code number stating that the type of statistics is stored in slot <i>N</i> of the <i>pg_statistic</i> row.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                                                                              |
| staopN          | oid      | <p>Operator used to generate the statistics stored in slot <i>N</i>. For example, a histogram slot shows the &lt; operator that defines the sort order of the data.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                        |
| stanumbers<br>N | real[]   | <p>Numerical statistics of the appropriate type for slot <i>N</i>. The value is <b>NULL</b> if the slot does not involve numerical values.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                                                 |
| stavaluesN      | anyarray | <p>Column data values of the appropriate type for slot <i>N</i>. The value is <b>NULL</b> if the slot type does not store any data values. Each array's element values are actually of the specific column's data type, so there is no way to define these columns' type more specifically than anyarray.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                  |

| Name          | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stadndistinct | real | Number of unique non-null data values in the <b>DN1</b> column. <ul style="list-style-type: none"><li>A value greater than 0 is the actual number of distinct values.</li><li>A value less than 0 is the ratio of the <b>distinct</b> value to the total number of rows. For example, if <b>stadistinct</b> is <b>-0.5</b>, the actual <b>distinct</b> value is the total number of rows multiplied by 0.5.</li><li>The value <b>0</b> indicates that the number of distinct values is unknown.</li></ul> |
| staextinfo    | text | Information about extension statistics. This is reserved.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**NOTICE**

PG\_STATISTIC stores sensitive information about statistical objects, such as MCVs. Users with the SYSADMIN permission and authorized users can access the PG\_STATISTIC system catalog to query the sensitive information about the statistical objects.

## 12.2.71 PG\_STATISTIC\_EXT

PG\_STATISTIC\_EXT records extended statistics about tables in a database. You can specify the extended statistics to be collected. This system catalog is accessible only to users with the SYSADMIN permission.

**Table 12-72** PG\_STATISTIC\_EXT columns

| Name        | Type    | Description                                                                                                                    |
|-------------|---------|--------------------------------------------------------------------------------------------------------------------------------|
| starelid    | oid     | Table or index that the described column belongs to.                                                                           |
| starelkind  | char    | Type of the object to which a table belongs. <b>c</b> indicates an ordinary table, and <b>p</b> indicates a partitioned table. |
| stainherit  | Boolean | Specifies whether to collect statistics for objects that have inheritance relationship.                                        |
| stanullfrac | real    | Percentage of column entries that are null.                                                                                    |
| stawidth    | integer | Average stored width, in bytes, of non-null entries.                                                                           |

| Name          | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| stadistinct   | real         | <p>Number of distinct, non-null data values in the column for database nodes.</p> <ul style="list-style-type: none"> <li>A value greater than 0 is the actual number of distinct values.</li> <li>A value less than 0 is the ratio of the <b>distinct</b> value to the total number of rows. For example, if <b>stadistinct</b> is <b>-0.5</b>, the actual <b>distinct</b> value is the total number of rows multiplied by 0.5.</li> <li>The value <b>0</b> indicates that the number of distinct values is unknown.</li> </ul> |
| stadndistinct | real         | <p>Number of unique non-null data values in the <b>DN1</b> column.</p> <ul style="list-style-type: none"> <li>A value greater than 0 is the actual number of distinct values.</li> <li>A value less than 0 is the ratio of the <b>distinct</b> value to the total number of rows. For example, if <b>stadistinct</b> is <b>-0.5</b>, the actual <b>distinct</b> value is the total number of rows multiplied by 0.5.</li> <li>The value <b>0</b> indicates that the number of distinct values is unknown.</li> </ul>            |
| stakindN      | smallint     | <p>Code number stating that the type of statistics is stored in slot <i>N</i> of the <i>pg_statistic</i> row.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                                                                              |
| staopN        | oid          | <p>Operator used to generate the statistics stored in slot <i>N</i>. For example, a histogram slot shows the &lt; operator that defines the sort order of the data.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                        |
| stakey        | int2vector   | Array of a column ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| stanumbersN   | real[]       | <p>Numerical statistics of the appropriate type for slot <i>N</i>. The value is <b>NULL</b> if the slot does not involve numerical values.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                                                                                                                                                                                 |
| stavaluesN    | anyarray     | <p>Column data values of the appropriate type for slot <i>N</i>. The value is <b>NULL</b> if the slot type does not store any data values. Each array's element values are actually of the specific column's data type, so there is no way to define these columns' type more specifically than anyarray.</p> <p>The value of <i>N</i> ranges from 1 to 5.</p>                                                                                                                                                                  |
| staexprs      | pg_node_tree | Expression corresponding to the extended statistics information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**NOTICE**

PG\_STATISTIC\_EXT stores sensitive information about statistical objects, such as MCVs. Users with the SYSADMIN permission and authorized users can access the PG\_STATISTIC\_EXT system catalog to query the sensitive information about the statistical objects.

## 12.2.72 PG\_SYNONYM

PG\_SYNONYM records the mapping between synonym object names and other database object names.

**Table 12-73** PG\_SYNONYM columns

| Name         | Type | Description                                                    |
|--------------|------|----------------------------------------------------------------|
| oid          | oid  | Database object ID                                             |
| synname      | name | Synonym name                                                   |
| synnamespace | oid  | OID of the namespace that contains a synonym                   |
| synowner     | oid  | Owner of a synonym, usually the OID of the user who created it |
| synobjschema | name | Schema name specified by an associated object                  |
| synobjname   | name | Name of an associated object                                   |

## 12.2.73 PG\_TABLESPACE

PG\_TABLESPACE records tablespace information.

**Table 12-74** PG\_TABLESPACE columns

| Name       | Type      | Description                                                                             |
|------------|-----------|-----------------------------------------------------------------------------------------|
| oid        | oid       | Row identifier (hidden attribute, which must be specified).                             |
| spcname    | name      | Tablespace name.                                                                        |
| spcowner   | oid       | Owner of the tablespace, usually the user who created it.                               |
| spcacl     | aclitem[] | Access permissions. For details, see <a href="#">GRANT</a> and <a href="#">REVOKE</a> . |
| spcoptions | text[]    | Options of the tablespace.                                                              |



| Name       | Type    | Description                                                                                                                                                                              |
|------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| spcmaxsize | text    | Maximum size of the available disk space, in bytes.                                                                                                                                      |
| relative   | Boolean | Specifies whether the storage path specified by the tablespace is a relative path. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul> |

## 12.2.74 PG\_TRIGGER

PG\_TRIGGER records trigger information.

**Table 12-75** PG\_TRIGGER columns

| Name           | Type     | Description                                                                                                                                                                                                                                                             |
|----------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid            | oid      | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                                             |
| tgrelid        | oid      | OID of the table where the trigger is located.                                                                                                                                                                                                                          |
| tgname         | name     | Trigger name.                                                                                                                                                                                                                                                           |
| tgfoid         | oid      | Function to be called by the trigger.                                                                                                                                                                                                                                   |
| tgtype         | smallint | Trigger type.                                                                                                                                                                                                                                                           |
| tgenabled      | "char"   | <ul style="list-style-type: none"><li>• <b>O</b>: The trigger is enabled in origin or local mode.</li><li>• <b>D</b>: The trigger is disabled.</li><li>• <b>R</b>: The trigger is enabled in replica mode.</li><li>• <b>A</b>: The trigger is always enabled.</li></ul> |
| tgisinternal   | Boolean  | Internal trigger ID. If the value is <b>true</b> , it indicates an internal trigger.                                                                                                                                                                                    |
| tgconstrrelid  | oid      | Table referenced by the integrity constraint.                                                                                                                                                                                                                           |
| tgconstrindid  | oid      | Index of the integrity constraint.                                                                                                                                                                                                                                      |
| tgconstraint   | oid      | OID of the constraint trigger in <a href="#">PG_CONSTRAINT</a> .                                                                                                                                                                                                        |
| tgdeferrable   | Boolean  | Specifies whether the constraint trigger is of the DEFERRABLE type.                                                                                                                                                                                                     |
| tginitdeferred | Boolean  | Specifies whether the trigger is of the INITIALLY DEFERRED type.                                                                                                                                                                                                        |
| tgargs         | smallint | Number of input parameters of the trigger function.                                                                                                                                                                                                                     |

| Name    | Type         | Description                                                                            |
|---------|--------------|----------------------------------------------------------------------------------------|
| tgattr  | int2vector   | Column ID specified by the trigger. If no column is specified, an empty array is used. |
| tgargs  | bytea        | Parameter specified in the trigger.                                                    |
| tgqual  | pg_node_tree | WHEN condition of the trigger ( <b>NULL</b> if the WHEN condition does not exist).     |
| tgowner | oid          | Trigger owner.                                                                         |

## 12.2.75 PG\_TS\_CONFIG

**PG\_TS\_CONFIG** contains entries representing text search configurations. A configuration specifies a particular text search parser and a list of dictionaries to use for each of the parser's output token types.

The parser is shown in the **PG\_TS\_CONFIG** entry, but the token-to-dictionary mapping is defined by subsidiary entries in [PG\\_TS\\_CONFIG\\_MAP](#).

**Table 12-76** PG\_TS\_CONFIG columns

| Name         | Type   | Reference                         | Description                                                |
|--------------|--------|-----------------------------------|------------------------------------------------------------|
| oid          | oid    | -                                 | Row identifier (hidden attribute, which must be specified) |
| cfgname      | name   | -                                 | Text search configuration name                             |
| cfgnamespace | oid    | <a href="#">PG_NAMESPACE</a> .oid | OID of the namespace that contains the configuration       |
| cfgowner     | oid    | <a href="#">PG_AUTHID</a> .oid    | Owner of the configuration                                 |
| cfgparser    | oid    | <a href="#">PG_TS_PARSER</a> .oid | OID of the text search parser for this configuration       |
| cfoptions    | text[] | -                                 | Configuration options                                      |

## 12.2.76 PG\_TS\_CONFIG\_MAP

**PG\_TS\_CONFIG\_MAP** contains entries showing which text search dictionaries should be consulted, and in what order, for each output token type of each text search configuration's parser.

**Table 12-77** PG\_TS\_CONFIG\_MAP columns

| Name         | Type    | Reference                           | Description                                                                                          |
|--------------|---------|-------------------------------------|------------------------------------------------------------------------------------------------------|
| mapcfg       | oid     | OID in <a href="#">PG_TS_CONFIG</a> | OID of the <a href="#">PG_TS_CONFIG</a> entry owning this map entry                                  |
| maptokentype | integer | -                                   | Token type generated by the configuration's parser                                                   |
| mapseqno     | integer | -                                   | Sequence number of a token type when the values of <b>mapcfg</b> or <b>maptokentype</b> are the same |
| mapdict      | oid     | OID in <a href="#">PG_TS_DICT</a>   | OID of the text search dictionary to consult                                                         |

## 12.2.77 PG\_TS\_DICT

**PG\_TS\_DICT** contains entries that define text search dictionaries. A dictionary depends on a text search template, which specifies all the implementation functions needed; the dictionary itself provides values for the user-settable parameters supported by the template.

This division of labor allows dictionaries to be created by unprivileged users. The parameters are specified by a text string **dictinitoption**, whose format and meaning vary depending on the template.

**Table 12-78** PG\_TS\_DICT columns

| Name           | Type | Reference                           | Description                                                |
|----------------|------|-------------------------------------|------------------------------------------------------------|
| oid            | oid  | -                                   | Row identifier (hidden attribute, which must be specified) |
| dictname       | name | -                                   | Text search dictionary name                                |
| dictnamespace  | oid  | <a href="#">PG_NAMESPACE</a> .oid   | OID of the namespace that contains the dictionary          |
| dictowner      | oid  | <a href="#">PG_AUTHID</a> .oid      | Owner of the dictionary                                    |
| dicttemplate   | oid  | <a href="#">PG_TS_TEMPLATE</a> .oid | OID of the text search template for the dictionary         |
| dictinitoption | text | -                                   | Initialization option string for the template              |

## 12.2.78 PG\_TS\_PARSER

**PG\_TS\_PARSER** contains entries defining text search parsers. A parser is responsible for splitting input text into lexemes and assigning a token type to each lexeme. The new parser must be created by the database system administrator.

**Table 12-79** PG\_TS\_PARSER columns

| Name         | Type    | Reference                         | Description                                                |
|--------------|---------|-----------------------------------|------------------------------------------------------------|
| oid          | oid     | -                                 | Row identifier (hidden attribute, which must be specified) |
| prsname      | name    | -                                 | Text search parser name                                    |
| prsnamespace | oid     | <a href="#">PG_NAMESPACE</a> .oid | OID of the namespace that contains the parser              |
| prsstart     | regproc | <a href="#">PG_PROC</a> .praname  | Name of the parser's startup function                      |
| prstoken     | regproc | <a href="#">PG_PROC</a> .praname  | Name of the parser's next-token function                   |
| prsend       | regproc | <a href="#">PG_PROC</a> .praname  | Name of the parser's shutdown function                     |
| prsheadline  | regproc | <a href="#">PG_PROC</a> .praname  | Name of the parser's headline function                     |
| prsllextype  | regproc | <a href="#">PG_PROC</a> .praname  | Name of the parser's lextype function                      |

## 12.2.79 PG\_TS\_TEMPLATE

**PG\_TS\_TEMPLATE** contains entries defining text search templates. A template provides a framework for text search dictionaries. Since a template must be implemented by C-language-level functions, templates can be created only by database administrators.

**Table 12-80** PG\_TS\_TEMPLATE columns

| Name          | Type | Reference                         | Description                                                |
|---------------|------|-----------------------------------|------------------------------------------------------------|
| oid           | oid  | -                                 | Row identifier (hidden attribute, which must be specified) |
| tmplname      | name | -                                 | Text search template name                                  |
| tmplnamespace | oid  | <a href="#">PG_NAMESPACE</a> .oid | OID of the namespace that contains the template            |

| Name       | Type    | Reference                        | Description                                    |
|------------|---------|----------------------------------|------------------------------------------------|
| tmplinit   | regproc | <a href="#">PG_PROC</a> .proname | Name of the template's initialization function |
| tmpllexize | regproc | <a href="#">PG_PROC</a> .proname | Name of the template's lexize function         |

## 12.2.80 PG\_TYPE

PG\_TYPE stores information about data types.

**Table 12-81** PG\_TYPE columns

| Name         | Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid          | oid      | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                                                                                                                                                                                                                                                        |
| typename     | name     | Data type name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| typnamespace | oid      | OID of the namespace that contains the type.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| typowner     | oid      | Owner of the type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| typlen       | smallint | <p>Number of bytes in the internal representation of the type for a fixed-size type. It is a negative number for a variable-length type.</p> <ul style="list-style-type: none"><li>• <b>-1</b>: a "varlena" type (one that has a length word).</li><li>• <b>-2</b>: a null-terminated C string.</li></ul>                                                                                                                                                                          |
| typbyval     | Boolean  | <ul style="list-style-type: none"><li>• <b>true</b>: A value of this type is passed by value internally.</li><li>• <b>false</b>: A value of this type is passed by reference internally.</li></ul> <p>If <b>typlen</b> of this type is not <b>1</b>, <b>2</b>, <b>4</b>, or <b>8</b>, you are advised to transfer a reference value for <b>typbyval</b>. You can also transfer a value for <b>typbyval</b>. Variable-length types are usually passed by reference or by value.</p> |
| typtype      | "char"   | <ul style="list-style-type: none"><li>• <b>b</b>: basic type.</li><li>• <b>c</b>: composite type (for example, the row type of a table).</li><li>• <b>d</b>: domain type.</li><li>• <b>p</b>: pseudo type.</li><li>• <b>o</b>: set type.</li></ul> <p>For details, see <b>typrelid</b> and <b>typbasetype</b>.</p>                                                                                                                                                                 |

| Name           | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| typcategory    | "char"  | Fuzzy classification of data types, which can be used by parsers as the basis for data conversion.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| typispreferred | Boolean | <ul style="list-style-type: none"><li>• <b>true</b>: Data is converted when it complies with the conversion rule specified by <b>typcategory</b>.</li><li>• <b>false</b>: Data is not converted.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                             |
| typisdefined   | Boolean | <ul style="list-style-type: none"><li>• <b>true</b>: The type has been defined.</li><li>• <b>false</b>: The placeholder of an undefined type is used. In this case, there is no reliable information except the name, namespace, and OID of the type.</li></ul>                                                                                                                                                                                                                                                                                                                                                        |
| typdelim       | "char"  | Character that separates two values of this type when parsing an array input. Note that the delimiter is associated with the array element data type, not the array data type.                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| typrelid       | oid     | If this is a composite type (see <b>typtype</b> ), then this column points to the <b>PG_CLASS</b> entry that defines the corresponding table. For a free-standing composite type, the PG_CLASS entry does not represent a table, but it is required for the type's <b>PG_ATTRIBUTE</b> entries to link to. It is <b>0</b> for non-composite type.                                                                                                                                                                                                                                                                      |
| typelem        | oid     | If <b>typelem</b> is not <b>0</b> , it identifies another row in PG_TYPE. The current type can be described as an array yielding values of type <b>typelem</b> . A "true" array type has a variable length ( <b>typlen</b> = <b>-1</b> ), but some fixed-length types ( <b>typlen</b> > <b>0</b> ) also have non-zero <b>typelem</b> , for example <b>name</b> and <b>point</b> . If a fixed-length type has a <b>typelem</b> , its internal representation must be a number of values of the <b>typelem</b> data type with no other data. Variable-length array types have a header defined by the array subroutines. |
| typarray       | oid     | If the value is not <b>0</b> , the corresponding type record is available in PG_TYPE.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| typinput       | regproc | Input conversion function (text format).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| typoutput      | regproc | Output conversion function (text format).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| typreceive     | regproc | Input conversion function (binary format); <b>0</b> for non-input conversion function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| typsend        | regproc | Output conversion function (binary format); <b>0</b> for non-output conversion function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| typmodin       | regproc | Input type modifier function; <b>0</b> if none.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| typmodout   | regproc | Output type modifier function; <b>0</b> if none.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| typanalyze  | regproc | Custom <b>ANALYZE</b> function; <b>0</b> if the standard function is used.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| typalign    | "char"  | <p>Alignment required when storing a value of this type. It applies to storage on disks as well as most representations of the value. When multiple values are stored consecutively, such as in the representation of a complete row on disk, padding is inserted before a data of this type so that it begins on the specified boundary. The alignment reference is the beginning of the first datum in the sequence. Possible values are:</p> <ul style="list-style-type: none"><li>• <b>c</b>: char alignment, that is, no alignment required</li><li>• <b>s</b>: short alignment (2 bytes on most machines)</li><li>• <b>i</b>: integer alignment (4 bytes on most machines)</li><li>• <b>d</b>: double alignment (8 bytes on many machines, but by no means all)</li></ul> <p><b>NOTICE</b><br/>For types used in system catalogs, the size and alignment defined in PG_TYPE must agree with the way that the compiler lays out the column in a structure representing a table row.</p> |
| typstorage  | "char"  | <p><b>typstorage</b> tells for varlena types (those with <b>typlen = -1</b>) if the type is prepared for toasting and what the default strategy for attributes of this type should be. Possible values are:</p> <ul style="list-style-type: none"><li>• <b>p</b>: Values are always stored plain.</li><li>• <b>e</b>: Values can be stored in a secondary relationship (if the relation has one, see <a href="#">PG_CLASS.reltoastrelid</a>).</li><li>• <b>m</b>: Values can be stored compressed inline.</li><li>• <b>x</b>: Values can be stored compressed inline or stored in secondary storage.</li></ul> <p><b>NOTICE</b><br/><b>m</b> domains can also be moved out to secondary storage, but only as a last resort (<b>e</b> and <b>x</b> domains are moved first).</p>                                                                                                                                                                                                              |
| typnotnull  | Boolean | Specifies whether the type has a NOTNULL constraint. Currently, it is used for domains only.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| typbasetype | oid     | If this is a domain (see <b>typtype</b> ), then <b>typbasetype</b> identifies the type that this one is based on. The value is <b>0</b> if this type is not a derived type.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

| Name          | Type         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| typtypmod     | integer      | Records the <b>typtypmod</b> to be applied to domains' base types by domains (the value is <b>-1</b> if the base type does not use <b>typmod</b> ). This is <b>-1</b> if this type is not a domain.                                                                                                                                                                                                                                                                                                |
| typndims      | integer      | If a field is an array, <b>typndims</b> is the number of array dimensions. This is <b>0</b> for types other than domains over array types.                                                                                                                                                                                                                                                                                                                                                         |
| typcollation  | oid          | Sequence rule for specified types. For details about the value, see the <a href="#">PG_COLLATION</a> system catalog. ( <b>0</b> if sequencing is not supported)                                                                                                                                                                                                                                                                                                                                    |
| typdefaultbin | pg_node_tree | nodeToString() representation of a default expression for the type if the value is non-null. Currently, this column is only used for domains.                                                                                                                                                                                                                                                                                                                                                      |
| typdefault    | text         | The value is <b>NULL</b> if a type has no associated default value. <ul style="list-style-type: none"><li>• If <b>typdefaultbin</b> is not <b>NULL</b>, <b>typdefault</b> must contain a default expression represented by <b>typdefaultbin</b>.</li><li>• If <b>typdefaultbin</b> is <b>NULL</b> and <b>typdefault</b> is not, then <b>typdefault</b> is the external representation of the type's default value, which can be fed to the type's input converter to produce a constant.</li></ul> |
| typacl        | aclitem[]    | Access permission.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## 12.2.81 PG\_USER\_MAPPING

**PG\_USER\_MAPPING** records mappings from local users to remote.

This system catalog is accessible only to users with the SYSADMIN permission. Common users can query the [PG\\_USER\\_MAPPINGS](#) view.

**Table 12-82** PG\_USER\_MAPPING columns

| Name     | Type | Reference                              | Description                                                                  |
|----------|------|----------------------------------------|------------------------------------------------------------------------------|
| oid      | oid  | -                                      | Row identifier (hidden attribute, which must be specified)                   |
| umuser   | oid  | <a href="#">PG_AUTHID</a> .oid         | OID of the local role being mapped ( <b>0</b> if the user mapping is public) |
| umserver | oid  | <a href="#">PG_FOREIGN_SERVER</a> .oid | OID of the foreign server that contains the mapping                          |



| Name      | Type   | Reference | Description                                                                         |
|-----------|--------|-----------|-------------------------------------------------------------------------------------|
| umoptions | text[] | -         | User mapping specific options, expressed in a string in the format of keyword=value |

## 12.2.82 PG\_USER\_STATUS

PG\_USER\_STATUS records the states of users who access the database. Only users with the SYSADMIN permission can access this system catalog.

**Table 12-83** PG\_USER\_STATUS columns

| Name             | Type                     | Description                                                                                                                                                                                                                                                              |
|------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid              | oid                      | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                                              |
| roloid           | oid                      | ID of the role.                                                                                                                                                                                                                                                          |
| failcount        | integer                  | Number of failed attempts.                                                                                                                                                                                                                                               |
| locktime         | timestamp with time zone | Time at which the role is locked.                                                                                                                                                                                                                                        |
| rolstatus        | smallint                 | Role state. <ul style="list-style-type: none"><li>• <b>0</b>: normal</li><li>• <b>1</b>: The role is locked for a specific period of time because the failed login attempts exceed the threshold.</li><li>• <b>2</b>: The role is locked by the administrator.</li></ul> |
| permpspace       | bigint                   | Size of the permanent table storage space used by the role.                                                                                                                                                                                                              |
| tempSPACE        | bigint                   | Size of the temporary table storage space used by the role.                                                                                                                                                                                                              |
| password expired | smallint                 | Specifies whether a password is valid. <ul style="list-style-type: none"><li>• <b>0</b>: The password is valid.</li><li>• <b>1</b>: The password is invalid.</li></ul>                                                                                                   |

## 12.2.83 PG\_WORKLOAD\_GROUP

PG\_WORKLOAD\_GROUP provides workload group information in the database.

**Table 12-84** PG\_WORKLOAD\_GROUP columns

| Name            | Type    | Description                                                |
|-----------------|---------|------------------------------------------------------------|
| oid             | oid     | Row identifier (hidden attribute, which must be specified) |
| workload_gpname | name    | Workload group name                                        |
| respool_oid     | oid     | ID bound to the resource pool                              |
| act_statements  | integer | Maximum number of active statements in the workload group  |

## 12.2.84 PGXC\_CLASS

PGXC\_CLASS records the replicated or distributed information for each table. **PGXC\_CLASS** can only be used to query table definitions in centralized mode.

**Table 12-85** PGXC\_CLASS columns

| Name            | Type             | Description                                                                                                                                                                                                                  |
|-----------------|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pcrelid         | oid              | Table OID.                                                                                                                                                                                                                   |
| pclocatortype   | "char"           | Locator type. <ul style="list-style-type: none"><li>• <b>H</b>: Hash</li><li>• <b>G</b>: Range</li><li>• <b>L</b>: List</li><li>• <b>M</b>: Modulo</li><li>• <b>N</b>: Round Robin</li><li>• <b>R</b>: Replication</li></ul> |
| pchashalgorithm | smallint         | Distributed tuple using the hash algorithm.                                                                                                                                                                                  |
| pchashbuckets   | smallint         | Value of a hash container.                                                                                                                                                                                                   |
| pgroup          | name             | Name of a node group.                                                                                                                                                                                                        |
| redistributed   | "char"           | Indicates that a table has been redistributed.                                                                                                                                                                               |
| redis_order     | integer          | Redistribution sequence. Tables whose values are <b>0</b> will not be redistributed in this round of redistribution.                                                                                                         |
| pcattnum        | int2vector       | Column number used as a distributed key.                                                                                                                                                                                     |
| nodeoids        | oidvector_extend | List of distributed table node OIDs.                                                                                                                                                                                         |
| options         | text             | Extension status information. This is a reserved column in the system.                                                                                                                                                       |

## 12.2.85 PGXC\_GROUP

PGXC\_GROUP records information about node groups. PGXC\_GROUP can only be used to query table definitions in the centralized system.

**Table 12-86** PGXC\_GROUP columns

| Name              | Type              | Description                                                                                                                                                                                                                                                                                                           |
|-------------------|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid               | oid               | Row identifier (hidden attribute, which must be specified).                                                                                                                                                                                                                                                           |
| group_name        | name              | Node group name.                                                                                                                                                                                                                                                                                                      |
| in_redistribution | "char"            | Specifies whether redistribution is required. Valid value: <ul style="list-style-type: none"><li>• <b>n</b>: The node group is not redistributed.</li><li>• <b>y</b>: The source node group is in redistribution.</li><li>• <b>t</b>: The destination node group is in redistribution.</li></ul>                      |
| group_members     | oidvector_ext end | Node OID list of the node group.                                                                                                                                                                                                                                                                                      |
| group_buckets     | text              | Distributed data bucket group.                                                                                                                                                                                                                                                                                        |
| is_installation   | Boolean           | Specifies whether to install a sub-database instance. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                                                                                                                                                           |
| group_acl         | aclitem[]         | Access permission.                                                                                                                                                                                                                                                                                                    |
| group_kind        | "char"            | Node group type. The options are as follows: <ul style="list-style-type: none"><li>• <b>i</b>: installation node group</li><li>• <b>n</b>: node group in a common non-logical database instance</li><li>• <b>v</b>: node group in a logical database instance</li><li>• <b>e</b>: elastic database instance</li></ul> |
| group_parent      | oid               | For a child node group, this field indicates the OID of the parent node group. For a parent node group, this field is left blank.                                                                                                                                                                                     |

## 12.2.86 PGXC\_NODE

PGXC\_NODE stores information about database instance nodes. **PGXC\_NODE** can only be used to query table definitions in centralized mode.

**Table 12-87** PGXC\_NODE columns

| Name             | Type    | Description                                                                                                                                                                                                            |
|------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| oid              | oid     | Row identifier (hidden attribute, which must be specified)                                                                                                                                                             |
| node_name        | name    | Node name                                                                                                                                                                                                              |
| node_type        | "char"  | Node type <ul style="list-style-type: none"><li>● <b>C</b>: CN</li><li>● <b>D</b>: DN</li><li>● <b>S</b>: standby node</li></ul>                                                                                       |
| node_port        | integer | Port number of the node                                                                                                                                                                                                |
| node_host        | name    | Host name or IP address of a node. (If a virtual IP address is configured, its value is a virtual IP address.)                                                                                                         |
| node_port1       | integer | Port number of a replication node                                                                                                                                                                                      |
| node_host1       | name    | Host name or IP address of a replication node. (If a virtual IP address is configured, its value is a virtual IP address.)                                                                                             |
| hostis_primary   | boolean | Specifies whether a primary/standby switchover occurs on the current node. <ul style="list-style-type: none"><li>● <b>t</b> (true): yes</li><li>● <b>f</b> (false): no</li></ul>                                       |
| nodeis_primary   | boolean | Specifies whether the current node is preferred to execute non-query operations in the <b>replication</b> table. <ul style="list-style-type: none"><li>● <b>t</b> (true): yes</li><li>● <b>f</b> (false): no</li></ul> |
| nodeis_preferred | boolean | Specifies whether the current node is preferred to execute queries in the <b>replication</b> table. <ul style="list-style-type: none"><li>● <b>t</b> (true): yes</li><li>● <b>f</b> (false): no</li></ul>              |
| node_id          | integer | Node ID. The value is obtain by calculating the value of <b>node_name</b> using the hash function.                                                                                                                     |
| sctp_port        | integer | Port used by the TCP proxy communications library of the primary node to listen to the data channel. (The SCTP communications library is no longer supported in the current version.)                                  |
| control_port     | integer | Port used by the TCP proxy communications library of the primary node to listen to the control channel                                                                                                                 |

| Name           | Type    | Description                                                                                                                                                                           |
|----------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sctp_port1     | integer | Port used by the TCP proxy communications library of the standby node to listen to the data channel. (The SCTP communications library is no longer supported in the current version.) |
| control_port1  | integer | Port used by the TCP proxy communications library of the standby node to listen to the control channel                                                                                |
| nodeis_central | boolean | Specifies whether the current node is a CN. It is invalid for DNs. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>              |
| nodeis_active  | boolean | Specifies whether the current node is normal. It is invalid for DNs. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>            |

## 12.2.87 PGXC\_SLICE

PGXC\_SLICE is a system catalog created for recording range distribution and list distribution details. Currently, range interval cannot be used to automatically scale out shards. It is reserved in the system catalog.

PGXC\_SLICE can only be used to query table definitions in centralized mode.

**Table 12-88** PGXC\_SLICE columns

| Name             | Type   | Description                                                                                                                                                                                      |
|------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| relname          | name   | Table name or shard name, which is distinguished by <b>type</b> .                                                                                                                                |
| type             | "char" | <ul style="list-style-type: none"><li>• <b>'t'</b>: <b>relname</b> is the table name.</li><li>• <b>'s'</b>: <b>relname</b> is the shard name.</li></ul>                                          |
| strategy         | "char" | <ul style="list-style-type: none"><li>• <b>'r'</b>: range distributed table.</li><li>• <b>'l'</b>: list distributed table.</li></ul> This value will be extended for subsequent interval shards. |
| relid            | oid    | OID of the distributed table to which the tuple belongs.                                                                                                                                         |
| referenc<br>eoid | oid    | OID of the referenced distributed table, which is used for slice reference table creation syntax.                                                                                                |

| Name            | Type    | Description                                                                                                            |
|-----------------|---------|------------------------------------------------------------------------------------------------------------------------|
| sindex          | integer | Position of the current boundary in a shard when the table is a list distributed table.                                |
| interval        | text[]  | Reserved column.                                                                                                       |
| transitboundary | text[]  | Reserved column.                                                                                                       |
| transitno       | integer | Reserved column.                                                                                                       |
| nodeoid         | oid     | When <b>relname</b> is set to a shard name, <b>nodeoid</b> indicates the OID of the DN where the shard data is stored. |
| boundaries      | text[]  | When <b>relname</b> is set to a shard name, this parameter indicates the boundary value of the shard.                  |
| specified       | Boolean | Specifies whether the DN corresponding to the current shard is explicitly specified in the DDL.                        |
| sliceorder      | integer | User-defined shard sequence.                                                                                           |

## 12.2.88 PLAN\_TABLE\_DATA

PLAN\_TABLE\_DATA stores plan information collected by EXPLAIN PLAN. Different from the PLAN\_TABLE view, the system catalog PLAN\_TABLE\_DATA stores EXPLAIN PLAN information collected by all sessions and users.

**Table 12-89** PLAN\_TABLE\_DATA columns

| Name         | Type         | Description                                                                                                                                              |
|--------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| session_id   | text         | Session that inserts the data. Its value consists of a worker thread start timestamp and a worker thread ID. Values are constrained by <b>NOT NULL</b> . |
| user_id      | oid          | ID of the user who inserts the data. Values are constrained by <b>NOT NULL</b> .                                                                         |
| statement_id | varchar2(30) | Query tag specified by a user.                                                                                                                           |
| plan_id      | bigint       | Query ID The ID is automatically generated in the plan generation phase and is used by kernel engineers for debugging.                                   |
| id           | integer      | Node ID in a plan                                                                                                                                        |
| operation    | varchar2(30) | Operation description.                                                                                                                                   |

| Name         | Type             | Description                                                 |
|--------------|------------------|-------------------------------------------------------------|
| options      | varchar2(255)    | Operation action.                                           |
| object_name  | name             | Name of an operated object. It is defined by users.         |
| object_type  | varchar2(30)     | Object type.                                                |
| object_owner | name             | Schema to which the object belongs. It is defined by users. |
| projection   | varchar2(4000)   | Returned column information.                                |
| cost         | double precision | Execution cost estimated by the optimizer for an operator   |
| cardinality  | double precision | Number of rows estimated by the optimizer for an operator   |

 NOTE

- PLAN\_TABLE\_DATA records data of all users and sessions on the current node. Only administrators can access all the data. Common users can view their own data in the [PLAN\\_TABLE](#) view.
- Data is automatically inserted into PLAN\_TABLE\_DATA after **EXPLAIN PLAN** is executed. Therefore, do not manually insert data into or update data in PLAN\_TABLE\_DATA. Otherwise, data in PLAN\_TABLE\_DATA may be disordered. To delete data, you are advised to use the [PLAN\\_TABLE](#) view.
- Information in the **statement\_id**, **object\_name**, **object\_owner**, and **projection** columns is stored in letter cases specified by users and information in other columns is stored in uppercase.

## 12.2.89 STATEMENT\_HISTORY

Displays information about statements executed on the current node. To query this system catalog, you must have the SYSADMIN permission. The result can be queried only in the system database but cannot be queried in the user database.

The constraints on the query of this system catalog are as follows:

- Data must be queried in the Postgres database. No data exists in other databases.
- This system catalog is controlled by **track\_stmt\_stat\_level**. The default value is **OFF,L0**, where the first part controls full SQL statements, and the second part controls slow SQL statements. For details about the record level of each field, see the following table.
- For slow SQL statements, if the value of **track\_stmt\_stat\_level** is not **OFF** and the SQL execution time exceeds the value of **log\_min\_duration\_statement**, the SQL statement is recorded as a slow SQL statement.

**Table 12-90** STATEMENT\_HISTORY columns

| Name               | Type                     | Description                                                                                                                                                                                                                           | Record Level |
|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| db_name            | name                     | Database name.                                                                                                                                                                                                                        | L0           |
| schema_name        | name                     | Schema name.                                                                                                                                                                                                                          | L0           |
| origin_node        | integer                  | Node name.                                                                                                                                                                                                                            | L0           |
| user_name          | name                     | Username.                                                                                                                                                                                                                             | L0           |
| application_name   | text                     | Name of the application that sends a request.                                                                                                                                                                                         | L0           |
| client_addr        | text                     | IP address of the client that sends a request.                                                                                                                                                                                        | L0           |
| client_port        | integer                  | Port number of the client that sends a request.                                                                                                                                                                                       | L0           |
| unique_query_id    | bigint                   | ID of the normalized SQL statement.                                                                                                                                                                                                   | L0           |
| debug_query_id     | bigint                   | ID of the unique SQL statement.                                                                                                                                                                                                       | L0           |
| query              | text                     | Normalized SQL statement.                                                                                                                                                                                                             | L0           |
| start_time         | timestamp with time zone | Time when a statement starts.                                                                                                                                                                                                         | L0           |
| finish_time        | timestamp with time zone | Time when a statement ends.                                                                                                                                                                                                           | L0           |
| slow_sql_threshold | bigint                   | Standard for slow SQL statement execution.                                                                                                                                                                                            | L0           |
| transaction_id     | bigint                   | Transaction ID.                                                                                                                                                                                                                       | L0           |
| thread_id          | bigint                   | ID of an execution thread.                                                                                                                                                                                                            | L0           |
| session_id         | bigint                   | Session ID of a user.                                                                                                                                                                                                                 | L0           |
| n_soft_parse       | bigint                   | Number of soft parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . | L0           |



| Name                | Type   | Description                                                                                                                                                                                                                           | Record Level |
|---------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| n_hard_parse        | bigint | Number of hard parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . | L0           |
| query_plan          | text   | Statement execution plan.                                                                                                                                                                                                             | L1           |
| n_returned_rows     | bigint | Number of rows in the result set returned by the <b>SELECT</b> statement.                                                                                                                                                             | L0           |
| n_tuples_fetched    | bigint | Number of rows randomly scanned.                                                                                                                                                                                                      | L0           |
| n_tuples_returned   | bigint | Number of rows sequentially scanned.                                                                                                                                                                                                  | L0           |
| n_tuples_inserted   | bigint | Number of rows inserted.                                                                                                                                                                                                              | L0           |
| n_tuples_updated    | bigint | Number of rows updated.                                                                                                                                                                                                               | L0           |
| n_tuples_deleted    | bigint | Number of rows deleted.                                                                                                                                                                                                               | L0           |
| n_blocks_fetched    | bigint | Number of buffer block access times.                                                                                                                                                                                                  | L0           |
| n_blocks_hit        | bigint | Number of buffer block hits.                                                                                                                                                                                                          | L0           |
| db_time             | bigint | Valid DB time, which is accumulated if multiple threads are involved (unit: $\mu$ s).                                                                                                                                                 | L0           |
| cpu_time            | bigint | CPU time (unit: $\mu$ s).                                                                                                                                                                                                             | L0           |
| execution_time      | bigint | Execution time in the executor (unit: $\mu$ s).                                                                                                                                                                                       | L0           |
| parse_time          | bigint | SQL parsing time (unit: $\mu$ s).                                                                                                                                                                                                     | L0           |
| plan_time           | bigint | SQL plan generation time (unit: $\mu$ s).                                                                                                                                                                                             | L0           |
| rewrite_time        | bigint | SQL rewriting time (unit: $\mu$ s).                                                                                                                                                                                                   | L0           |
| pl_execution_time   | bigint | Execution time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                           | L0           |
| pl_compilation_time | bigint | Compilation time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                         | L0           |
| data_io_time        | bigint | I/O time (unit: $\mu$ s).                                                                                                                                                                                                             | L0           |

| Name                 | Type   | Description                                                                                                                                                                                           | Record Level |
|----------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| net_send_info        | text   | Network status of messages sent through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.     | L0           |
| net_rcv_info         | text   | Network status of messages received through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode. | L0           |
| net_stream_send_info | text   | Network status of messages sent through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.      | L0           |
| net_stream_rcv_info  | text   | Network status of messages received through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.  | L0           |
| lock_count           | bigint | Number of locks.                                                                                                                                                                                      | L0           |
| lock_time            | bigint | Time required for locking.                                                                                                                                                                            | L1           |
| lock_wait_count      | bigint | Number of lock waits.                                                                                                                                                                                 | L0           |
| lock_wait_time       | bigint | Time required for lock waiting.                                                                                                                                                                       | L1           |
| lock_max_count       | bigint | Maximum number of locks.                                                                                                                                                                              | L0           |
| lwlock_count         | bigint | Number of lightweight locks (reserved).                                                                                                                                                               | L0           |
| lwlock_wait_count    | bigint | Number of lightweight lock waits.                                                                                                                                                                     | L0           |
| lwlock_time          | bigint | Time required for lightweight locking (reserved).                                                                                                                                                     | L1           |
| lwlock_wait_time     | bigint | Time required for lightweight lock waiting.                                                                                                                                                           | L1           |

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Record Level |
|-------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| details     | bytea   | List of statement lock events, which are recorded in time sequence. The number of records is affected by the <b>track_stmt_details_size</b> parameter. This field is in binary format and needs to be read by using the parsing function <code>pg_catalog.statement_detail_decode</code> . For details, see <a href="#">Table 7-51</a> .<br>Events include: <ul style="list-style-type: none"><li>• Start locking.</li><li>• Complete locking.</li><li>• Start lock waiting.</li><li>• Complete lock waiting.</li><li>• Start unlocking.</li><li>• Complete unlocking.</li><li>• Start lightweight lock waiting.</li><li>• Complete lightweight lock waiting.</li></ul> | L2           |
| is_slow_sql | Boolean | Specifies whether the SQL statement is a slow SQL statement. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | L0           |
| trace_id    | text    | Driver-specific trace ID, which is associated with an application request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | L0           |

| Name   | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | Record Level |
|--------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| advise | text | <p>Risks which may cause slow SQL statements. (Multiple risks may exist at the same time.)</p> <ul style="list-style-type: none"><li>• Cast Function Cause Index Miss. : Index matching may fail due to implicit conversion.</li><li>• Limit too much rows. : The SQL statement execution may slow down due to a large <b>limit</b> value.</li><li>• Proleakproof of function is false. : The <b>proleakproof</b> of the function is set to <b>false</b>. In this case, the function does not use statistics when generating a plan due to data leakage risks. As a result, the accuracy of the generated plan is affected and the SQL statement execution may slow down.</li></ul> | L0           |

## 12.2.90 STREAMING\_STREAM

**STREAMING\_STREAM** records the metadata of all STREAM objects.

**Table 12-91** STREAMING\_STREAM columns

| Name    | Type  | Description                                                 |
|---------|-------|-------------------------------------------------------------|
| relid   | oid   | STREAM object ID.                                           |
| queries | bytea | Bitmap mapping of the CONTVIEW corresponding to the STREAM. |

## 12.2.91 STREAMING\_CONT\_QUERY

**STREAMING\_CONT\_QUERY** records the metadata of all CONTVIEW objects.

**Table 12-92** STREAMING\_CONT\_QUERY columns

| Name | Type    | Description                               |
|------|---------|-------------------------------------------|
| id   | integer | Unique identifier of the CONTVIEW object. |

| Name        | Type       | Description                                                                                                                                                                 |
|-------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| type        | "char"     | CONTVIEW type. <ul style="list-style-type: none"><li>• <b>'r'</b>: CONTVIEW is based on the row-store model.</li></ul>                                                      |
| relid       | oid        | CONTVIEW object ID.                                                                                                                                                         |
| defrelid    | oid        | ID of the continuous computing rule view corresponding to CONTVIEW.                                                                                                         |
| active      | Boolean    | Determines whether the CONTVIEW is in the continuous computing state. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul> |
| streamrelid | oid        | ID of STREAM corresponding to CONTVIEW.                                                                                                                                     |
| matrelid    | oid        | ID of the materialized table corresponding to CONTVIEW.                                                                                                                     |
| lookupidxid | oid        | ID of GROUP LOOK UP INDEX corresponding to CONTVIEW. This column is for internal use and is available only in row-store tables.                                             |
| step_factor | smallint   | CONTVIEW step mode. The main values are <b>0</b> (no overlapping window) and <b>1</b> (sliding window, with one step).                                                      |
| ttl         | integer    | Value of <b>ttl_interval</b> set by CONTVIEW.                                                                                                                               |
| ttl_attno   | smallint   | Number of a time column corresponding to the TTL function set by CONTVIEW.                                                                                                  |
| dictrelid   | oid        | ID of the dictionary table corresponding to CONTVIEW.                                                                                                                       |
| grpnum      | smallint   | Number of dimension columns in the CONTVIEW continuous computing rule. This column is for internal use.                                                                     |
| grpidx      | int2vector | Index of the dimension column in TARGET LIST in the CONTVIEW continuous computing rule. This column is for internal use.                                                    |

## 12.2.92 STREAMING\_REAPER\_STATUS

**STREAMING\_REAPER\_STATUS** records the status information about the reaper thread of the streaming engine. (Due to specification changes, the current version no longer supports this feature. Do not use it.)

**Table 12-93** STREAMING\_REAPER\_STATUS columns

| Name                   | Type    | Description                                                                                                                                        |
|------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| id                     | integer | Unique identifier of the CONTVIEW object.                                                                                                          |
| contquery_name         | name    | Name of the CONTVIEW object.                                                                                                                       |
| gather_interval        | text    | Value of <b>gather_interval</b> (time parameter for automatically aggregating historical data before a specific time) set for the CONTVIEW object. |
| gather_completion_time | text    | Time when the latest GATHER (historical data aggregation) operation on the CONTVIEW object is completed.                                           |

## 12.3 System Views

### 12.3.1 ADM\_COL\_COMMENTS

ADM\_COL\_COMMENTS displays information about table column comments in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-94** ADM\_COL\_COMMENTS columns

| Name        | Type                  | Description  |
|-------------|-----------------------|--------------|
| owner       | character varying(64) | Table owner. |
| table_name  | character varying(64) | Table name.  |
| column_name | character varying(64) | Column name. |
| comments    | text                  | Comments.    |

### 12.3.2 ADM\_CONS\_COLUMNS

ADM\_CONS\_COLUMNS displays information about constraint columns in database tables. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-95** ADM\_CONS\_COLUMNS columns

| Name            | Type                  | Description                          |
|-----------------|-----------------------|--------------------------------------|
| owner           | character varying(64) | Constraint creator.                  |
| column_name     | character varying(64) | Name of a constraint-related column. |
| constraint_name | character varying(64) | Constraint name.                     |
| position        | smallint              | Position of a column in a table.     |
| table_name      | character varying(64) | Name of a constraint-related table.  |

### 12.3.3 ADM\_CONSTRAINTS

ADM\_CONSTRAINTS displays information about table constraints in database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-96** ADM\_CONSTRAINTS columns

| Name            | Type                  | Description                                                                                                                                                                                                                   |
|-----------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | character varying(64) | Constraint creator.                                                                                                                                                                                                           |
| constraint_name | character varying(64) | Constraint name.                                                                                                                                                                                                              |
| constraint_type | text                  | Constraint type. <ul style="list-style-type: none"><li>• <b>c</b>: check constraint.</li><li>• <b>f</b>: foreign key constraint.</li><li>• <b>p</b>: primary key constraint.</li><li>• <b>u</b>: unique constraint.</li></ul> |
| table_name      | character varying(64) | Name of a constraint-related table.                                                                                                                                                                                           |
| index_owner     | character varying(64) | Owner of a constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                              |
| index_name      | character varying(64) | Name of the constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                             |

## 12.3.4 ADM\_DATA\_FILES

ADM\_DATA\_FILES displays the description of database files. By default, only the SYSADMIN can access this view. Common users can access the view only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-97** ADM\_DATA\_FILES columns

| Name            | Type             | Description                                     |
|-----------------|------------------|-------------------------------------------------|
| tablespace_name | name             | Name of the tablespace to which a file belongs. |
| bytes           | double precision | Length of a file, in bytes.                     |

## 12.3.5 ADM\_HIST\_SNAPSHOT

ADM\_HIST\_SNAPSHOT records the index information and start time of WDR snapshots stored in the current system. By default, only the SYSADMIN can access this view. Common users can access the view only after being authorized.

**Table 12-98** ADM\_HIST\_SNAPSHOT columns

| Name                | Type      | Description                   |
|---------------------|-----------|-------------------------------|
| snap_id             | bigint    | WDR snapshot ID.              |
| begin_interval_time | timestamp | Start time of a WDR snapshot. |

## 12.3.6 ADM\_HIST\_SQL\_PLAN

ADM\_HIST\_SQL\_PLAN displays plan information collected by the current user by running the **EXPLAIN PLAN** statement. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized.

**Table 12-99** ADM\_HIST\_SQL\_PLAN columns

| Name            | Type                  | Description                                                                                                                                              |
|-----------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| sql_id          | character varying(30) | Session that inserts the data. Its value consists of a worker thread start timestamp and a worker thread ID. Values are constrained by <b>NOT NULL</b> . |
| plan_hash_value | bigint                | Query ID.                                                                                                                                                |



| Name        | Type                   | Description                                         |
|-------------|------------------------|-----------------------------------------------------|
| operation   | character varying(30)  | Operation description.                              |
| options     | character varying(255) | Operation option.                                   |
| object_name | name                   | Name of an operated object. It is defined by users. |

### 12.3.7 ADM\_HIST\_SQLSTAT

ADM\_HIST\_SQLSTAT displays information about statements executed on the current node. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized.

**Table 12-100** ADM\_HIST\_SQLSTAT columns

| Name                 | Type    | Description                                                                                                  |
|----------------------|---------|--------------------------------------------------------------------------------------------------------------|
| INSTANCE_NUMBER      | integer | Instance ID of a snapshot.                                                                                   |
| SQL_ID               | bigint  | Query ID.                                                                                                    |
| PLAN_HASH_VALUE      | bigint  | ID of the normalized SQL statement.                                                                          |
| MODULE               | text    | Name of the module that is executing when the SQL statement is first parsed.                                 |
| ELAPSED_TIME_DELTA   | bigint  | Valid DB time, which is accumulated if multiple threads are involved (unit: $\mu$ s).                        |
| CPU_TIME_DELTA       | bigint  | CPU time (unit: $\mu$ s)                                                                                     |
| EXECUTIONS_DELTA     | integer | Increment in the number of executions that have occurred on this object since it was brought into the cache. |
| IOWAIT_DELTA         | bigint  | I/O time (unit: $\mu$ s)                                                                                     |
| APWAIT_DELTA         | integer | Delta value of the application wait time.                                                                    |
| ROWS_PROCESSED_DELTA | bigint  | Number of rows in the result set returned by the SELECT statement.                                           |
| SNAP_ID              | integer | Unique snapshot ID.                                                                                          |

### 12.3.8 ADM\_IND\_COLUMNS

ADM\_IND\_COLUMNS displays column information about all indexes in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-101** ADM\_IND\_COLUMNS columns

| Name            | Type                  | Description                          |
|-----------------|-----------------------|--------------------------------------|
| index_owner     | character varying(64) | Index owner.                         |
| index_name      | character varying(64) | Index name.                          |
| table_owner     | character varying(64) | Table owner.                         |
| table_name      | character varying(64) | Table name.                          |
| column_name     | name                  | Column name.                         |
| column_position | smallint              | Position of the column in the index. |

### 12.3.9 ADM\_IND\_EXPRESSIONS

ADM\_IND\_EXPRESSIONS displays information about expression indexes in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-102** ADM\_IND\_EXPRESSIONS columns

| Name              | Type                  | Description                                            |
|-------------------|-----------------------|--------------------------------------------------------|
| table_owner       | character varying(64) | Table owner.                                           |
| table_name        | character varying(64) | Table name.                                            |
| index_owner       | character varying(64) | Index owner.                                           |
| index_name        | character varying(64) | Index name.                                            |
| column_expression | text                  | Function-based index expression of a specified column. |
| column_position   | smallint              | Position of the column in the index.                   |

### 12.3.10 ADM\_IND\_PARTITIONS

ADM\_IND\_PARTITIONS displays information about all index partitions in the database (excluding global indexes on partitioned tables). Each index partition of

a partitioned table in the database, if present, has a row of records in ADM\_IND\_PARTITIONS. By default, only the SYSADMIN can access this view. Common users can access the view only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-103** ADM\_IND\_PARTITIONS columns

| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index_owner            | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs.                                                                                                                                                                                                                                                                                                       |
| index_name             | character varying(64) | Name of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                                   |
| partition_name         | character varying(64) | Name of the index partition.                                                                                                                                                                                                                                                                                                                                                                |
| def_tablespace_name    | name                  | Tablespace name of the index partition.                                                                                                                                                                                                                                                                                                                                                     |
| high_value             | text                  | Limit of the partition corresponding to the index partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>• For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>• For list partitioning, the value list of each partition is displayed.</li><li>• For hash partitioning, the number of each partition is displayed.</li></ul> |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                                                                                                 |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                                 |
| high_value_length      | integer               | Character length of the limit of the partition corresponding to an index partition.                                                                                                                                                                                                                                                                                                         |

## 12.3.11 ADM\_IND\_SUBPARTITIONS

ADM\_IND\_SUBPARTITIONS displays information about level-2 partitions of all indexes (excluding global indexes of partitioned tables) in the database. Each level-2 index partition of a level-2 partitioned table in the database, if present, has a row of records in ADM\_IND\_SUBPARTITIONS. By default, only the SYSADMIN can access this view. Common users can access the view only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-104** ADM\_IND\_SUBPARTITIONS columns

| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                   |
|------------------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index_owner            | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs.                                                                                                                                                                                                                                         |
| index_name             | character varying(64) | Name of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                     |
| partition_name         | character varying(64) | Name of the partition where an index is located.                                                                                                                                                                                                                                                                              |
| subpartition_name      | character varying(64) | Name of the level-2 partition where an index is located.                                                                                                                                                                                                                                                                      |
| def_tablespace_name    | name                  | Tablespace name of an index partition.                                                                                                                                                                                                                                                                                        |
| high_value             | text                  | Limit of the partition corresponding to an index partition.<br><b>NOTE</b><br>For range partitioning and interval partitioning, the upper limit of each partition is displayed.<br>For list partitioning, the value list of each partition is displayed.<br>For hash partitioning, the number of each partition is displayed. |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                                   |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                   |

| Name              | Type    | Description                                                                         |
|-------------------|---------|-------------------------------------------------------------------------------------|
| high_value_length | integer | Character length of the limit of the partition corresponding to an index partition. |

## 12.3.12 ADM\_INDEXES

ADM\_INDEXES displays all indexes in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-105** ADM\_INDEXES columns

| Name        | Type                  | Description                                                         |
|-------------|-----------------------|---------------------------------------------------------------------|
| owner       | character varying(64) | Index owner.                                                        |
| index_name  | character varying(64) | Index name.                                                         |
| table_name  | character varying(64) | Name of the table corresponding to the index.                       |
| uniqueness  | text                  | Specifies whether an index is unique.                               |
| partitioned | character(3)          | Specifies whether the index has the property of partitioned tables. |
| generated   | character varying(1)  | Specifies whether the index name is generated by the system.        |

## 12.3.13 ADM\_OBJECTS

ADM\_OBJECTS displays all database objects in the database. Only the SYSADMIN can access this view. Common users can access the view only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-106** ADM\_OBJECTS columns

| Name        | Type | Description   |
|-------------|------|---------------|
| owner       | name | Object owner. |
| object_name | name | Object name.  |

| Name          | Type                     | Description                                         |
|---------------|--------------------------|-----------------------------------------------------|
| object_id     | oid                      | Object OID.                                         |
| object_type   | name                     | Object type. For example, table, schema, and index. |
| namespace     | oid                      | Namespace containing an object.                     |
| created       | timestamp with time zone | Creation time of the object.                        |
| last_ddl_time | timestamp with time zone | Last modification time of the object.               |

**NOTICE**

For details about the value ranges of **created** and **last\_ddl\_time**, see [PG\\_OBJECT](#).

## 12.3.14 ADM\_PART\_INDEXES

ADM\_PART\_INDEXES displays information about all partitioned table indexes (excluding global indexes of partitioned tables) in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-107** ADM\_PART\_INDEXES columns

| Name                   | Type                  | Description                                              |
|------------------------|-----------------------|----------------------------------------------------------|
| def_tablespace_name    | name                  | Tablespace name of a partitioned table index.            |
| index_owner            | character varying(64) | Owner name of a partitioned table index.                 |
| index_name             | character varying(64) | Name of a partitioned table index.                       |
| partition_count        | bigint                | Number of index partitions of a partitioned table index. |
| partitioning_key_count | integer               | Number of partition keys of the partitioned table.       |

| Name                      | Type                  | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| partitioning_type         | text                  | Partitioning policy of the partitioned table.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> .                                                                                                     |
| schema                    | character varying(64) | Name of the schema to which the partitioned table index belongs.                                                                                                                                                                                                                                         |
| table_name                | character varying(64) | Name of the partitioned table to which a partitioned table index belongs.                                                                                                                                                                                                                                |
| subpartitioning_type      | text                  | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer               | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer               | Number of level-2 partition keys of the partitioned table.                                                                                                                                                                                                                                               |

### 12.3.15 ADM\_PART\_TABLES

ADM\_PART\_TABLES displays information about all partitioned tables in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-108** ADM\_PART\_TABLES columns

| Name        | Type                  | Description                        |
|-------------|-----------------------|------------------------------------|
| table_owner | character varying(64) | Owner name of a partitioned table. |

| Name                      | Type                     | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_name                | character<br>varying(64) | Name of the partitioned table.                                                                                                                                                                                                                                                                           |
| partitioning_type         | text                     | Partitioning policy of the partitioned table.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> .                                                                                                     |
| partition_count           | bigint                   | Number of partitions of the partitioned table.                                                                                                                                                                                                                                                           |
| partitioning_key_count    | integer                  | Number of partition keys of the partitioned table.                                                                                                                                                                                                                                                       |
| def_tablespace_name       | name                     | Tablespace name of the partitioned table.                                                                                                                                                                                                                                                                |
| schema                    | character<br>varying(64) | Schema of the partitioned table.                                                                                                                                                                                                                                                                         |
| subpartitioning_type      | text                     | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer                  | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer                  | Number of level-2 partition keys of the partitioned table.                                                                                                                                                                                                                                               |

## 12.3.16 ADM\_PROCEEDURES

ADM\_PROCEEDURES displays information about all stored procedures or functions in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.



**Table 12-109** ADM\_PROCEDURES columns

| Name            | Type                  | Description                                       |
|-----------------|-----------------------|---------------------------------------------------|
| owner           | character varying(64) | Owner of a stored procedure or function.          |
| object_name     | character varying(64) | Name of a stored procedure or function.           |
| argument_number | smallint              | Number of input parameters in a stored procedure. |

### 12.3.17 ADM\_SCHEDULER\_JOBS

ADM\_SCHEDULER\_JOBS displays information about all DBE\_SCHEDULER scheduled jobs in the database.

**Table 12-110** ADM\_SCHEDULER\_JOBS columns

| Name                | Type                        | Description                                         |
|---------------------|-----------------------------|-----------------------------------------------------|
| owner               | name                        | Owner of a scheduled job.                           |
| job_name            | text                        | Name of a scheduled job.                            |
| job_style           | text                        | Action mode of a scheduled job.                     |
| job_creator         | name                        | Creator of a scheduled job.                         |
| program_name        | text                        | Name of the program referenced by a scheduled job.  |
| job_action          | text                        | Program content of a scheduled job.                 |
| number_of_arguments | text                        | Number of parameters of a scheduled job.            |
| schedule_name       | text                        | Name of the schedule referenced by a scheduled job. |
| start_date          | timestamp without time zone | Start time of a scheduled job.                      |
| repeat_interval     | text                        | Period of a scheduled job.                          |
| end_date            | timestamp without time zone | End time of a scheduled job.                        |

| Name            | Type                        | Description                                                       |
|-----------------|-----------------------------|-------------------------------------------------------------------|
| job_class       | text                        | Name of the scheduled job class to which a scheduled job belongs. |
| enabled         | Boolean                     | Status of a scheduled job.                                        |
| auto_drop       | text                        | Status of the automatic deletion function of a scheduled job.     |
| state           | "char"                      | Status of a scheduled job.                                        |
| failure_count   | smallint                    | Number of scheduled job failures.                                 |
| last_start_date | timestamp without time zone | Last time when a scheduled job was started.                       |
| next_run_date   | timestamp without time zone | Next execution time of a scheduled job.                           |
| destination     | text                        | Target name of a scheduled job.                                   |
| credential_name | text                        | Credential name of a scheduled job.                               |
| comments        | text                        | Comments of a scheduled job.                                      |

## 12.3.18 ADM\_SEQUENCES

ADM\_SEQUENCES displays information about all sequences in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-111** ADM\_SEQUENCES columns

| Name           | Type                  | Description                    |
|----------------|-----------------------|--------------------------------|
| sequence_owner | character varying(64) | Owner of a sequence.           |
| sequence_name  | character varying(64) | Name of the sequence.          |
| min_value      | int16                 | Minimum value of the sequence. |
| max_value      | int16                 | Maximum value of the sequence. |

| Name         | Type         | Description                                                                                                                                                                                                                     |
|--------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| increment_by | int16        | Increment of the sequence.                                                                                                                                                                                                      |
| last_number  | int16        | Value of the previous sequence.                                                                                                                                                                                                 |
| cache_size   | int16        | Size of the sequence disk cache.                                                                                                                                                                                                |
| cycle_flag   | character(1) | Specifies whether the sequence is a cycle sequence. The value can be <b>Y</b> or <b>N</b> . <ul style="list-style-type: none"><li>• <b>Y</b>: It is a cycle sequence.</li><li>• <b>N</b>: It is not a cycle sequence.</li></ul> |

### 12.3.19 ADM\_SOURCE

ADM\_SOURCE displays all stored procedures or functions in the database and provides columns defined by the stored procedures or functions. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-112** ADM\_SOURCE columns

| Name  | Type                  | Description                                   |
|-------|-----------------------|-----------------------------------------------|
| owner | character varying(64) | Owner of a stored procedure or function.      |
| name  | character varying(64) | Name of a stored procedure or function.       |
| text  | text                  | Definition of a stored procedure or function. |

### 12.3.20 ADM\_SYNONYMS

ADM\_SYNONYMS displays all synonyms in the database. This view is accessible only to users with SYSADMIN permissions. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-113** ADM\_SYNONYMS columns

| Name              | Type | Description                                                                                                                                                                                                                          |
|-------------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner             | text | Owner of a synonym.                                                                                                                                                                                                                  |
| schema_name       | text | Name of the schema to which the synonym belongs.                                                                                                                                                                                     |
| synonym_name      | text | Synonym name.                                                                                                                                                                                                                        |
| table_owner       | text | Owner of the associated object. Although the column is called <b>table_owner</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym.             |
| table_name        | text | Name of the associated object. Although the column is called <b>table_name</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym.               |
| table_schema_name | text | Schema name of the associated object. Although the column is called <b>table_schema_name</b> , the associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym. |

### 12.3.21 ADM\_TAB\_COLUMNS

ADM\_TAB\_COLUMNS displays information about the columns of tables. Each column of each table in the database has a row of data in ADM\_TAB\_COLUMNS. By default, only the SYSADMIN can access this view. Common users can access the view only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-114** ADM\_TAB\_COLUMNS columns

| Name        | Type                   | Description                   |
|-------------|------------------------|-------------------------------|
| owner       | character varying(64)  | Table owner.                  |
| table_name  | character varying(64)  | Table name.                   |
| column_name | character varying(64)  | Column name.                  |
| data_type   | character varying(128) | Data type of the column.      |
| data_length | integer                | Length of a column, in bytes. |

| Name           | Type    | Description                                                                                                    |
|----------------|---------|----------------------------------------------------------------------------------------------------------------|
| data_precision | integer | Precision of a data type. This column is valid for the numeric data type and <b>NULL</b> for other data types. |
| data_scale     | integer | Number of decimal places. This column is valid for the numeric data type and <b>0</b> for other data types.    |
| nullable       | bpchar  | Specifies whether a column can be empty ( <b>n</b> for the primary key constraint and NOT NULL constraint).    |
| column_id      | integer | Sequence number of a column when a table is created.                                                           |
| avg_col_len    | numeric | Average length of a column, in bytes.                                                                          |
| char_length    | numeric | Column length (in characters) which is valid only for the varchar, nvarchar2, bpchar, and char types.          |
| comments       | text    | Comments.                                                                                                      |

### 12.3.22 ADM\_TAB\_COMMENTS

ADM\_TAB\_COMMENTS displays comments about all tables and views in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-115** ADM\_TAB\_COMMENTS columns

| Name       | Type                  | Description                |
|------------|-----------------------|----------------------------|
| owner      | character varying(64) | Owner of a table or view.  |
| table_name | character varying(64) | Name of the table or view. |
| comments   | text                  | Comments.                  |

### 12.3.23 ADM\_TAB\_PARTITIONS

ADM\_TAB\_PARTITIONS displays information about all partitioned tables in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-116** ADM\_TAB\_PARTITIONS columns

| Name               | Type                  | Description                                                                                                                                                                                                                                                                                                                                          |
|--------------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_owner        | character varying(64) | Role name.                                                                                                                                                                                                                                                                                                                                           |
| table_name         | character varying(64) | Relational table name.                                                                                                                                                                                                                                                                                                                               |
| partition_name     | character varying(64) | Partition name.                                                                                                                                                                                                                                                                                                                                      |
| high_value         | text                  | Limit of a partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>• For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>• For list partitioning, the value list of each partition is displayed.</li><li>• For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name    | name                  | Tablespace name of the partitioned table.                                                                                                                                                                                                                                                                                                            |
| schema             | character varying(64) | Name of a namespace.                                                                                                                                                                                                                                                                                                                                 |
| subpartition_count | bigint                | Number of level-2 partitions.                                                                                                                                                                                                                                                                                                                        |
| high_value_length  | integer               | Length of the partition boundary value expression.                                                                                                                                                                                                                                                                                                   |

### 12.3.24 ADM\_TAB\_SUBPARTITIONS

ADM\_TAB\_SUBPARTITIONS displays information about all level-2 partitions in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-117** ADM\_TAB\_SUBPARTITIONS columns

| Name              | Type                  | Description                  |
|-------------------|-----------------------|------------------------------|
| table_owner       | character varying(64) | Role name.                   |
| table_name        | character varying(64) | Relational table name.       |
| partition_name    | character varying(64) | Partition name.              |
| subpartition_name | character varying(64) | Name of a level-2 partition. |

| Name              | Type                  | Description                                                                                                                                                                                                                                                                                                                                            |
|-------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| high_value        | text                  | Limit of a level-2 partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name   | name                  | Tablespace name of a level-2 partitioned table.                                                                                                                                                                                                                                                                                                        |
| schema            | character varying(64) | Name of a namespace.                                                                                                                                                                                                                                                                                                                                   |
| high_value_length | integer               | Character length of the limit of a level-2 partition.                                                                                                                                                                                                                                                                                                  |

### 12.3.25 ADM\_TABLES

ADM\_TABLES displays information about all partitioned tables in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-118** ADM\_TABLES columns

| Name            | Type                  | Description                                                                                                                                                      |
|-----------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | character varying(64) | Table owner.                                                                                                                                                     |
| table_name      | character varying(64) | Table name.                                                                                                                                                      |
| tablespace_name | character varying(64) | Tablespace name of the table.                                                                                                                                    |
| dropped         | character varying     | Specifies whether the current table is deleted. <ul style="list-style-type: none"><li><b>YES:</b> It is deleted.</li><li><b>NO:</b> It is not deleted.</li></ul> |
| num_rows        | numeric               | Estimated number of rows in the table.                                                                                                                           |

| Name      | Type                 | Description                                                                                                                                                                                            |
|-----------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status    | character varying(8) | Specifies whether the current table is valid. <ul style="list-style-type: none"><li>• <b>VALID</b>: The current table is valid.</li><li>• <b>UNUSABLE</b>: The current table is unavailable.</li></ul> |
| temporary | character(1)         | Specifies whether the table is a temporary table. <ul style="list-style-type: none"><li>• <b>Y</b>: The table is a temporary table.</li><li>• <b>N</b>: The table is not a temporary table.</li></ul>  |

### 12.3.26 ADM\_TABLESPACES

ADM\_TABLESPACES displays information about available tablespaces. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-119** ADM\_TABLESPACES columns

| Name            | Type                  | Description     |
|-----------------|-----------------------|-----------------|
| tablespace_name | character varying(64) | Tablespace name |

### 12.3.27 ADM\_TRIGGERS

ADM\_TRIGGERS displays information about triggers in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-120** ADM\_TRIGGERS columns

| Name         | Type                  | Description   |
|--------------|-----------------------|---------------|
| trigger_name | character varying(64) | Trigger name. |
| table_owner  | character varying(64) | Role name.    |



| Name       | Type                  | Description                                                                                                                                                                                                                                                             |
|------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_name | character varying(64) | Relational table name.                                                                                                                                                                                                                                                  |
| status     | character varying(64) | <ul style="list-style-type: none"><li>• <b>O</b>: The trigger is enabled in origin or local mode.</li><li>• <b>D</b>: The trigger is disabled.</li><li>• <b>R</b>: The trigger is enabled in replica mode.</li><li>• <b>A</b>: The trigger is always enabled.</li></ul> |

### 12.3.28 ADM\_TYPE\_ATTRS

ADM\_TYPE\_ATTRS displays the attributes of the current database object type. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized.

**Table 12-121** ADM\_TYPE\_ATTRS columns

| Name            | Type     | Description                                                                                                                                                                                                                                                                                                                              |
|-----------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | oid      | Owner of the type.                                                                                                                                                                                                                                                                                                                       |
| type_name       | name     | Data type name.                                                                                                                                                                                                                                                                                                                          |
| attr_name       | name     | Column name.                                                                                                                                                                                                                                                                                                                             |
| attr_type_mod   | integer  | Type-specific data supplied at the table creation time (for example, the maximum length of a varchar column). This column is used as the third parameter when passing to type-specific input functions and length coercion functions. The value will generally be <b>-1</b> for types that do not need ATTTYPMOD.                        |
| attr_type_owner | oid      | Owner of an attribute of this type.                                                                                                                                                                                                                                                                                                      |
| attr_type_name  | name     | Data type attribute name.                                                                                                                                                                                                                                                                                                                |
| length          | smallint | Number of bytes in the internal representation of the type for a fixed-size type. It is a negative number for a variable-length type. <ul style="list-style-type: none"><li>• The value <b>-1</b> indicates a "varlena" type (one that has a length word).</li><li>• The value <b>-2</b> indicates a NULL-terminated C string.</li></ul> |
| precision       | integer  | Precision of the numeric type.                                                                                                                                                                                                                                                                                                           |
| scale           | integer  | Range of the numeric type                                                                                                                                                                                                                                                                                                                |

| Name               | Type         | Description                                                                                |
|--------------------|--------------|--------------------------------------------------------------------------------------------|
| character_set_name | character(1) | Character set name of an attribute ( <b>c</b> or <b>n</b> ).                               |
| attr_no            | smallint     | Attribute number                                                                           |
| inherited          | character(1) | Specifies whether the attribute is inherited from the super type ( <b>Y</b> or <b>N</b> ). |

## 12.3.29 ADM\_USERS

ADM\_USERS displays information about all database users. This view is accessible only to users with SYSADMIN permissions. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-122** ADM\_USERS columns

| Name     | Type                  | Description |
|----------|-----------------------|-------------|
| username | character varying(64) | Username.   |

## 12.3.30 ADM\_VIEWS

ADM\_VIEWS displays views in the database. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-123** ADM\_VIEWS columns

| Name      | Type                  | Description      |
|-----------|-----------------------|------------------|
| owner     | character varying(64) | Owner of a view. |
| view_name | character varying(64) | View name.       |

## 12.3.31 DB\_ALL\_TABLES

DB\_ALL\_TABLES displays tables or views accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-124** DB\_ALL\_TABLES columns

| Name            | Type | Description                                    |
|-----------------|------|------------------------------------------------|
| owner           | name | Owner of a table or view.                      |
| table_name      | name | Name of the table or view.                     |
| tablespace_name | name | Tablespace where the table or view is located. |

### 12.3.32 DB\_COL\_COMMENTS

DB\_COL\_COMMENTS displays comment information about table columns accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-125** DB\_COL\_COMMENTS columns

| Name        | Type                  | Description  |
|-------------|-----------------------|--------------|
| owner       | character varying(64) | Table owner. |
| table_name  | character varying(64) | Table name.  |
| column_name | character varying(64) | Column name. |
| comments    | text                  | Comments.    |

### 12.3.33 DB\_CONS\_COLUMNS

DB\_CONS\_COLUMNS displays information about constraint columns accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-126** DB\_CONS\_COLUMNS columns

| Name            | Type                  | Description                          |
|-----------------|-----------------------|--------------------------------------|
| owner           | character varying(64) | Constraint creator.                  |
| constraint_name | character varying(64) | Constraint name.                     |
| table_name      | character varying(64) | Name of a constraint-related table.  |
| column_name     | character varying(64) | Name of a constraint-related column. |
| position        | smallint              | Position of a column in a table.     |

## 12.3.34 DB\_CONSTRAINTS

DB\_CONSTRAINTS displays information about constraints accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-127** DB\_CONSTRAINTS columns

| Name            | Type                  | Description                                                                                                                                                                                                                   |
|-----------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | character varying(64) | Constraint creator.                                                                                                                                                                                                           |
| constraint_name | character varying(64) | Constraint name.                                                                                                                                                                                                              |
| constraint_type | text                  | Constraint type. <ul style="list-style-type: none"><li>• <b>c</b>: check constraint.</li><li>• <b>f</b>: foreign key constraint.</li><li>• <b>p</b>: primary key constraint.</li><li>• <b>u</b>: unique constraint.</li></ul> |
| table_name      | character varying(64) | Name of a constraint-related table.                                                                                                                                                                                           |
| index_owner     | character varying(64) | Owner of a constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                              |
| index_name      | character varying(64) | Name of the constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                             |

## 12.3.35 DB\_DEPENDENCIES

DB\_DEPENDENCIES displays dependencies between functions and advanced packages accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

---

**NOTICE**

In GaussDB, this table is empty without any record due to information constraints.

---

**Table 12-128** DB\_DEPENDENCIES columns

| Name  | Type                  | Description   |
|-------|-----------------------|---------------|
| owner | character varying(30) | Object owner. |
| name  | character varying(30) | Object name.  |

| Name                 | Type                   | Description                                                                                          |
|----------------------|------------------------|------------------------------------------------------------------------------------------------------|
| type                 | character varying(17)  | Object type.                                                                                         |
| referenced_owner     | character varying(30)  | Owner of a referenced object.                                                                        |
| referenced_name      | character varying(64)  | Name of a referenced object.                                                                         |
| referenced_type      | character varying(17)  | Type of a referenced object.                                                                         |
| referenced_link_name | character varying(128) | Name of the link to a referenced object.                                                             |
| schemaid             | numeric                | ID of the current schema.                                                                            |
| dependency_type      | character varying(4)   | Dependency type ( <b>REF</b> indicates soft reference and <b>HARD</b> indicates direct description). |

### 12.3.36 DB\_IND\_COLUMNS

DB\_IND\_COLUMNS displays all index columns accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-129** DB\_IND\_COLUMNS columns

| Name            | Type                  | Description                          |
|-----------------|-----------------------|--------------------------------------|
| index_owner     | character varying(64) | Index owner.                         |
| index_name      | character varying(64) | Index name.                          |
| table_owner     | character varying(64) | Table owner.                         |
| table_name      | character varying(64) | Table name.                          |
| column_name     | name                  | Column name.                         |
| column_position | smallint              | Position of the column in the index. |

### 12.3.37 DB\_IND\_EXPRESSIONS

DB\_IND\_EXPRESSIONS displays information about expression indexes accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-130** DB\_IND\_EXPRESSIONS columns

| Name              | Type                  | Description                                            |
|-------------------|-----------------------|--------------------------------------------------------|
| index_owner       | character varying(64) | Index owner.                                           |
| index_name        | character varying(64) | Index name.                                            |
| table_owner       | character varying(64) | Table owner.                                           |
| table_name        | character varying(64) | Table name.                                            |
| column_expression | text                  | Function-based index expression of a specified column. |
| column_position   | smallint              | Position of the column in the index.                   |

### 12.3.38 DB\_IND\_PARTITIONS

DB\_IND\_PARTITIONS displays information about partitioned table indexes (excluding global indexes of partitioned tables) accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-131** DB\_IND\_PARTITIONS columns

| Name                | Type                  | Description                                                                           |
|---------------------|-----------------------|---------------------------------------------------------------------------------------|
| index_owner         | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs. |
| index_name          | character varying(64) | Name of the partitioned table index to which the index partition belongs.             |
| partition_name      | character varying(64) | Name of the partition where an index is located.                                      |
| def_tablespace_name | name                  | Tablespace name of the index partition.                                               |

| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| high_value             | text                  | Limit of the partition corresponding to the index partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li><b>t</b> (true): yes.</li><li><b>f</b> (false): no.</li></ul>                                                                                                                                                                                                                               |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                           |
| high_value_length      | integer               | Character length of the limit of the partition corresponding to an index partition.                                                                                                                                                                                                                                                                                                   |

### 12.3.39 DB\_IND\_SUBPARTITIONS

DB\_IND\_SUBPARTITIONS displays information about level-2 partitions of indexes (excluding global indexes of partitioned tables) accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-132** DB\_IND\_SUBPARTITIONS columns

| Name           | Type                  | Description                                                                           |
|----------------|-----------------------|---------------------------------------------------------------------------------------|
| index_owner    | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs. |
| index_name     | character varying(64) | Name of the partitioned table index to which the index partition belongs.             |
| partition_name | character varying(64) | Name of the index partition.                                                          |

| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subpartition_name      | character varying(64) | Name of a level-2 index partition.                                                                                                                                                                                                                                                                                                                                                         |
| def_tablespace_name    | name                  | Tablespace name of an index partition.                                                                                                                                                                                                                                                                                                                                                     |
| high_value             | text                  | Limit of the partition corresponding to an index partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>• For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>• For list partitioning, the value list of each partition is displayed.</li><li>• For hash partitioning, the number of each partition is displayed.</li></ul> |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                                                                                                 |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                                |
| high_value_length      | integer               | Character length of the limit of the partition corresponding to an index partition.                                                                                                                                                                                                                                                                                                        |

## 12.3.40 DB\_INDEXES

DB\_INDEXES displays information about indexes accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-133** DB\_INDEXES columns

| Name       | Type                  | Description                                   |
|------------|-----------------------|-----------------------------------------------|
| owner      | character varying(64) | Index owner.                                  |
| index_name | character varying(64) | Index name.                                   |
| table_name | character varying(64) | Name of the table corresponding to the index. |



| Name        | Type                 | Description                                                         |
|-------------|----------------------|---------------------------------------------------------------------|
| uniqueness  | text                 | Specifies whether the index is unique.                              |
| partitioned | character(3)         | Specifies whether the index has the property of partitioned tables. |
| generated   | character varying(1) | Specifies whether the name of an index is generated by the system.  |

### 12.3.41 DB\_OBJECTS

DB\_OBJECTS displays all database objects accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-134** DB\_OBJECTS columns

| Name          | Type                     | Description                                   |
|---------------|--------------------------|-----------------------------------------------|
| owner         | name                     | Object owner.                                 |
| object_name   | name                     | Object name.                                  |
| object_id     | oid                      | Object OID.                                   |
| object_type   | name                     | Object type.                                  |
| namespace     | oid                      | ID of the namespace where the object resides. |
| created       | timestamp with time zone | Time when the object was created.             |
| last_ddl_time | timestamp with time zone | Last time when the object was modified.       |

#### NOTICE

For details about the value ranges of **created** and **last\_ddl\_time**, see [PG\\_OBJECT](#).

### 12.3.42 DB\_PART\_INDEXES

DB\_PART\_INDEXES displays information about partitioned table indexes (excluding global indexes of partitioned tables) accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-135** DB\_PART\_INDEXES columns

| Name                      | Type                  | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| def_tablespace_name       | name                  | Tablespace name of a partitioned table index                                                                                                                                                                                                                                                             |
| index_owner               | character varying(64) | Owner name of a partitioned table index                                                                                                                                                                                                                                                                  |
| index_name                | character varying(64) | Name of a partitioned table index                                                                                                                                                                                                                                                                        |
| partition_count           | bigint                | Number of index partitions of a partitioned table index                                                                                                                                                                                                                                                  |
| partitioning_key_count    | integer               | Number of partition keys of a partitioned table                                                                                                                                                                                                                                                          |
| partitioning_type         | text                  | Partitioning policy of a partitioned table<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> .                                                                                                        |
| schema                    | character varying(64) | Name of the schema to which a partitioned table index belongs                                                                                                                                                                                                                                            |
| table_name                | character varying(64) | Name of the partitioned table to which a partitioned table index belongs                                                                                                                                                                                                                                 |
| subpartitioning_type      | text                  | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer               | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer               | Number of level-2 partition keys of a partitioned table                                                                                                                                                                                                                                                  |

## 12.3.43 DB\_PART\_TABLES

DB\_PART\_TABLES displays information about partitioned tables accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-136** DB\_PART\_TABLES columns

| Name                      | Type                  | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_owner               | character varying(64) | Owner name of a partitioned table.                                                                                                                                                                                                                                                                       |
| table_name                | character varying(64) | Name of the partitioned table.                                                                                                                                                                                                                                                                           |
| partitioning_type         | text                  | Partitioning policy of the partitioned table.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> .                                                                                                     |
| partition_count           | bigint                | Number of partitions of the partitioned table.                                                                                                                                                                                                                                                           |
| partitioning_key_count    | integer               | Number of partition keys of the partitioned table.                                                                                                                                                                                                                                                       |
| def_tablespace_name       | name                  | Tablespace name of the partitioned table.                                                                                                                                                                                                                                                                |
| schema                    | character varying(64) | Schema of the partitioned table.                                                                                                                                                                                                                                                                         |
| subpartitioning_type      | text                  | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer               | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer               | Number of level-2 partition keys of the partitioned table.                                                                                                                                                                                                                                               |

## 12.3.44 DB\_PROCEEDURES

DB\_PROCEEDURES displays information about all stored procedures or functions accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-137** DB\_PROCEEDURES columns

| Name        | Type | Description   |
|-------------|------|---------------|
| owner       | name | Object owner. |
| object_name | name | Object name.  |

## 12.3.45 DB\_SEQUENCES

DB\_SEQUENCES displays all sequences accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-138** DB\_SEQUENCES columns

| Name           | Type         | Description                                                                                                                                                                                                                     |
|----------------|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sequence_owner | name         | Owner of a sequence.                                                                                                                                                                                                            |
| sequence_name  | name         | Name of the sequence.                                                                                                                                                                                                           |
| min_value      | int16        | Minimum value of the sequence.                                                                                                                                                                                                  |
| max_value      | int16        | Maximum value of the sequence.                                                                                                                                                                                                  |
| increment_by   | int16        | Value by which the sequence is incremented.                                                                                                                                                                                     |
| cycle_flag     | character(1) | Specifies whether the sequence is a cycle sequence. The value can be <b>Y</b> or <b>N</b> . <ul style="list-style-type: none"><li>• <b>Y</b>: It is a cycle sequence.</li><li>• <b>N</b>: It is not a cycle sequence.</li></ul> |
| last_number    | int16        | Value of the previous sequence.                                                                                                                                                                                                 |
| cache_size     | int16        | Size of the sequence disk cache.                                                                                                                                                                                                |

## 12.3.46 DB\_SOURCE

DB\_SOURCE displays information about stored procedures or functions accessible to the current user, and provides columns defined by the stored procedures and functions. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-139** DB\_SOURCE columns

| Name  | Type | Description        |
|-------|------|--------------------|
| owner | name | Object owner.      |
| name  | name | Object name.       |
| type  | name | Object class.      |
| text  | text | Object definition. |

## 12.3.47 DB\_SYNONYMS

DB\_SYNONYMS displays all synonyms accessible to the current user.

**Table 12-140** DB\_SYNONYMS columns

| Name              | Type | Description                                                                                                                                                                                                                          |
|-------------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner             | text | Owner of a synonym.                                                                                                                                                                                                                  |
| schema_name       | text | Name of the schema to which the synonym belongs.                                                                                                                                                                                     |
| synonym_name      | text | Synonym name.                                                                                                                                                                                                                        |
| table_owner       | text | Owner of the associated object. Although the column is called <b>table_owner</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym.             |
| table_name        | text | Name of the associated object. Although the column is called <b>table_name</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym.               |
| table_schema_name | text | Schema name of the associated object. Although the column is called <b>table_schema_name</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym. |

## 12.3.48 DB\_TAB\_COLUMNS

DB\_TAB\_COLUMNS displays description information about columns of tables and views accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-141** DB\_TAB\_COLUMNS columns

| Name           | Type                   | Description                                                                                                    |
|----------------|------------------------|----------------------------------------------------------------------------------------------------------------|
| owner          | character varying(64)  | Table owner.                                                                                                   |
| table_name     | character varying(64)  | Table name.                                                                                                    |
| column_name    | character varying(64)  | Column name.                                                                                                   |
| data_type      | character varying(128) | Data type of the column.                                                                                       |
| data_length    | integer                | Length of a column, in bytes.                                                                                  |
| data_precision | integer                | Precision of a data type. This column is valid for the numeric data type and <b>NULL</b> for other data types. |
| data_scale     | integer                | Number of decimal places. This column is valid for the numeric data type and <b>0</b> for other data types.    |
| nullable       | bpchar                 | Specifies whether a column can be empty ( <b>n</b> for a primary key constraint or NOT NULL constraint).       |
| column_id      | integer                | Column ID generated when an object is created or a column is added.                                            |
| char_length    | numeric                | Column length (in characters) which is valid only for varchar, nvarchar2, bpchar, and char types.              |
| avg_col_len    | numeric                | Average length of a column, in bytes.                                                                          |
| comments       | text                   | Comments.                                                                                                      |

## 12.3.49 DB\_TAB\_COMMENTS

DB\_TAB\_COMMENTS displays comments about all tables and views accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-142** DB\_TAB\_COMMENTS columns

| Name       | Type                  | Description                |
|------------|-----------------------|----------------------------|
| owner      | character varying(64) | Owner of a table or view.  |
| table_name | character varying(64) | Name of the table or view. |
| comments   | text                  | Comments.                  |

## 12.3.50 DB\_TAB\_PARTITIONS

DB\_TAB\_PARTITIONS displays information about partitioned tables accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-143** DB\_TAB\_PARTITIONS columns

| Name               | Type                  | Description                                                                                                                                                                                                                                                                                                                                    |
|--------------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_owner        | character varying(64) | Role name.                                                                                                                                                                                                                                                                                                                                     |
| table_name         | character varying(64) | Relational table name.                                                                                                                                                                                                                                                                                                                         |
| partition_name     | character varying(64) | Partition name.                                                                                                                                                                                                                                                                                                                                |
| high_value         | text                  | Limit of a partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name    | name                  | Tablespace name of the partitioned table.                                                                                                                                                                                                                                                                                                      |
| schema             | character varying(64) | Name of a namespace.                                                                                                                                                                                                                                                                                                                           |
| subpartition_count | bigint                | Number of level-2 partitions.                                                                                                                                                                                                                                                                                                                  |
| high_value_length  | integer               | Length of the partition boundary value expression.                                                                                                                                                                                                                                                                                             |

## 12.3.51 DB\_TAB\_SUBPARTITIONS

DB\_TAB\_SUBPARTITIONS displays information about level-2 partitioned tables accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-144** DB\_TAB\_SUBPARTITIONS columns

| Name              | Type                  | Description                                                                                                                                                                                                                                                                                                                                            |
|-------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_owner       | character varying(64) | Role name.                                                                                                                                                                                                                                                                                                                                             |
| table_name        | character varying(64) | Relational table name.                                                                                                                                                                                                                                                                                                                                 |
| partition_name    | character varying(64) | Partition name.                                                                                                                                                                                                                                                                                                                                        |
| subpartition_name | character varying(64) | Name of a level-2 partition.                                                                                                                                                                                                                                                                                                                           |
| high_value        | text                  | Limit of a level-2 partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name   | name                  | Tablespace name of a level-2 partitioned table.                                                                                                                                                                                                                                                                                                        |
| schema            | character varying(64) | Name of a namespace.                                                                                                                                                                                                                                                                                                                                   |
| high_value_length | integer               | Length of the boundary value expression in level-2 partition.                                                                                                                                                                                                                                                                                          |

## 12.3.52 DB\_TABLES

DB\_TABLES displays all tables accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-145** DB\_TABLES columns

| Name            | Type                  | Description                            |
|-----------------|-----------------------|----------------------------------------|
| owner           | character varying(64) | Table owner.                           |
| table_name      | character varying(64) | Table name.                            |
| tablespace_name | character varying(64) | Tablespace name of the table.          |
| num_rows        | numeric               | Estimated number of rows in the table. |



| Name      | Type                 | Description                                                                                                                                                                                            |
|-----------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| status    | character varying(8) | Specifies whether the current table is valid. <ul style="list-style-type: none"><li>● <b>VALID</b>: The current table is valid.</li><li>● <b>UNUSABLE</b>: The current table is unavailable.</li></ul> |
| temporary | character(1)         | Specifies whether the table is a temporary table. <ul style="list-style-type: none"><li>● <b>Y</b>: The table is a temporary table.</li><li>● <b>N</b>: The table is not a temporary table.</li></ul>  |
| dropped   | character varying    | Specifies whether the current table is deleted. <ul style="list-style-type: none"><li>● <b>YES</b>: It is deleted.</li><li>● <b>NO</b>: It is not deleted.</li></ul>                                   |

### 12.3.53 DB\_TRIGGERS

DB\_TRIGGERS displays information about triggers accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-146** DB\_TRIGGERS columns

| Name         | Type                  | Description                                                                                                                                                                                                                                                             |
|--------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| trigger_name | character varying(64) | Trigger name.                                                                                                                                                                                                                                                           |
| table_owner  | character varying(64) | Role name.                                                                                                                                                                                                                                                              |
| table_name   | character varying(64) | Relational table name.                                                                                                                                                                                                                                                  |
| status       | character varying(64) | <ul style="list-style-type: none"><li>● <b>O</b>: The trigger is enabled in origin or local mode.</li><li>● <b>D</b>: The trigger is disabled.</li><li>● <b>R</b>: The trigger is enabled in replica mode.</li><li>● <b>A</b>: The trigger is always enabled.</li></ul> |

### 12.3.54 DB\_USERS

DB\_USERS displays all users of the database visible to the current user. However, it does not describe the users. By default, only users with the SYSADMIN permission can access this view. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-147** DB\_USERS columns

| Name     | Type | Description    |
|----------|------|----------------|
| user_id  | oid  | OID of a user. |
| username | name | Username.      |

## 12.3.55 DB\_VIEWS

DB\_VIEWS displays the description about all views accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-148** DB\_VIEWS columns

| Name        | Type    | Description              |
|-------------|---------|--------------------------|
| owner       | name    | Owner of a view.         |
| view_name   | name    | View name.               |
| text        | text    | Text in the view.        |
| text_length | integer | Text length of the view. |

## 12.3.56 DV\_SESSIONS

DV\_SESSIONS displays all session information about the current session. By default, only users with the SYSADMIN permission can access this system view. Common users can access it only after being authorized.

**Table 12-149** DV\_SESSIONS columns

| Name    | Type    | Description                                                                                                              |
|---------|---------|--------------------------------------------------------------------------------------------------------------------------|
| sid     | bigint  | OID of the active backend thread of the current session.                                                                 |
| serial# | integer | Sequence number of the active backend thread, which is <b>0</b> in GaussDB                                               |
| user#   | oid     | OID of the user logged in to the backend thread. The OID is <b>0</b> if the backend thread is a global auxiliary thread. |

| Name     | Type | Description                                                                                                                                                                                                                                                                                                                                                          |
|----------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| username | name | Name of the user logged in to the backend thread.<br><b>username</b> is null if the backend thread is a global auxiliary thread.<br><b>application_name</b> can be identified by associating with <b>pg_stat_get_activity()</b> .<br>Example:<br>SELECT s.*,a.application_name FROM DV_SESSIONS AS s LEFT JOIN pg_stat_get_activity(NULL) AS a ON s.sid=a.sessionid; |

### 12.3.57 DV\_SESSION\_LONGOPS

DV\_SESSION\_LONGOPS displays the progress of ongoing operations. The view can be accessed only after being authorized.

**Table 12-150** DV\_SESSION\_LONGOPS columns

| Name      | Type    | Description                                                                   |
|-----------|---------|-------------------------------------------------------------------------------|
| sid       | bigint  | OID of the running backend process.                                           |
| serial#   | integer | Sequence number of the running backend process, which is <b>0</b> in GaussDB. |
| sofar     | integer | Completed workload, which is null in GaussDB.                                 |
| totalwork | integer | Total workload, which is null in GaussDB.                                     |

### 12.3.58 GET\_GLOBAL\_PREPARED\_XACTS (Deprecated)

This view is not supported in the centralized system.

### 12.3.59 GS\_ALL\_CONTROL\_GROUP\_INFO

This view is not supported in the centralized system.

### 12.3.60 GS\_AUDITING

GS\_AUDITING displays all audit information about database-related operations. Only users with the SYSADMIN or POLADMIN permission can access this view.

**Table 12-151** GS\_AUDITING columns

| Name    | Type | Description                        |
|---------|------|------------------------------------|
| polname | name | Policy name, which must be unique. |

| Name        | Type    | Description                                                                                                                                                                                                                  |
|-------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pol_type    | text    | Audit policy type. The value can be <b>access</b> or <b>privilege</b> . <ul style="list-style-type: none"><li>• <b>access</b>: DML operations are audited.</li><li>• <b>privilege</b>: DDL operations are audited.</li></ul> |
| polenabled  | Boolean | Specifies whether to enable a policy. <ul style="list-style-type: none"><li>• <b>t</b> (true): enabled</li><li>• <b>f</b> (false): disabled</li></ul>                                                                        |
| access_type | name    | DML database operation type. For example, SELECT, INSERT, and DELETE.                                                                                                                                                        |
| label_name  | name    | Resource label name. This column corresponds to the <b>polname</b> column in the GS_AUDITING_POLICY system catalog.                                                                                                          |
| priv_object | text    | Describes the path of the database asset.                                                                                                                                                                                    |
| filter_name | text    | Logical character string of a filter criterion.                                                                                                                                                                              |

### 12.3.61 GS\_AUDITING\_ACCESS

GS\_AUDITING\_ACCESS displays all audit information about database DML-related operations. Only users with the SYSADMIN or POLADMIN permission can access this view.

**Table 12-152** GS\_AUDITING\_ACCESS columns

| Name          | Type    | Description                                                                                                                                             |
|---------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| polname       | name    | Policy name, which must be unique.                                                                                                                      |
| pol_type      | text    | Audit policy type. The value ' <b>access</b> ' indicates that DML operations are audited.                                                               |
| polenabled    | Boolean | Specifies whether to enable a policy. <ul style="list-style-type: none"><li>• <b>t</b> (true): enabled.</li><li>• <b>f</b> (false): disabled.</li></ul> |
| access_type   | name    | DML database operation type. For example, SELECT, INSERT, and DELETE.                                                                                   |
| label_name    | name    | Resource label name. This column corresponds to the <b>polname</b> column in the GS_AUDITING_POLICY system catalog.                                     |
| access_object | text    | Describes the path of the database asset.                                                                                                               |
| filter_name   | text    | Logical character string of a filter criterion.                                                                                                         |

## 12.3.62 GS\_AUDITING\_PRIVILEGE

GS\_AUDITING\_PRIVILEGE displays all audit information about database DDL-related operations. Only users with the SYSADMIN or POLADMIN permission can access this view.

**Table 12-153** GS\_AUDITING\_PRIVILEGE columns

| Name        | Type    | Description                                                                                                                                             |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| polname     | name    | Policy name, which must be unique.                                                                                                                      |
| pol_type    | text    | Audit policy type. The value <b>privilege</b> indicates that DDL operations are audited.                                                                |
| polenabled  | Boolean | Specifies whether to enable a policy. <ul style="list-style-type: none"><li>• <b>t</b> (true): enabled.</li><li>• <b>f</b> (false): disabled.</li></ul> |
| access_type | name    | DDL database operation type. For example, CREATE, ALTER, and DROP.                                                                                      |
| label_name  | name    | Resource label name. This column corresponds to the <b>polname</b> column in the GS_AUDITING_POLICY system catalog.                                     |
| priv_object | text    | Full domain name of a database object.                                                                                                                  |
| filter_name | text    | Logical character string of a filter criterion.                                                                                                         |

## 12.3.63 GS\_CLUSTER\_RESOURCE\_INFO

This view is not supported in the centralized system.

## 12.3.64 GS\_COMM\_PROXY\_THREAD\_STATUS

GS\_COMM\_PROXY\_THREAD\_STATUS displays statistics on packets sent and received by the proxy communications library **comm\_proxy**. This view displays statistics about data sent and received by **comm\_proxy** only when the user-mode network deployment mode is enabled during the installation of the centralized database and **enable\_dfx** of **comm\_proxy\_attr** is set to **true**. In other scenarios, this view cannot be queried.

**Table 12-154** GS\_COMM\_PROXY\_THREAD\_STATUS columns

| Name             | Type   | Description                                                                                               |
|------------------|--------|-----------------------------------------------------------------------------------------------------------|
| ProxyThreadId    | bigint | ID of the current network proxy thread <b>comm_proxy</b> .                                                |
| ProxyCpuAffinity | text   | NUMA-CPU affinity of the current network proxy thread <b>comm_proxy</b> , indicating the NUMA and CPU ID. |

| Name            | Type   | Description                                                                          |
|-----------------|--------|--------------------------------------------------------------------------------------|
| ThreadStartTime | text   | Start time of the current network proxy thread <b>comm_proxy</b> .                   |
| RxPckNums       | bigint | Number of packets received by the current network proxy thread <b>comm_proxy</b> .   |
| TxPckNums       | bigint | Number of packets sent by the current network proxy thread <b>comm_proxy</b> .       |
| RxPcks          | float  | Number of packets received by the network proxy thread <b>comm_proxy</b> per second. |
| TxPcks          | float  | Number of packets sent by the network proxy thread <b>comm_proxy</b> per second.     |

### 12.3.65 GS\_DB\_PRIVILEGES

**GS\_DB\_PRIVILEGES** displays the granting of ANY permissions. Each record corresponds to a piece of authorization information.

**Table 12-155** GS\_DB\_PRIVILEGES columns

| Name           | Type | Description                                                                                                                                                                           |
|----------------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rolename       | name | Username.                                                                                                                                                                             |
| privilege_type | text | ANY permission of a user. For details about the value, see <a href="#">Table 7-108</a> .                                                                                              |
| admin_option   | text | Specifies whether the ANY permission recorded in the <b>privilege_type</b> column can be re-granted. <ul style="list-style-type: none"><li>• <b>yes</b></li><li>• <b>no</b></li></ul> |

### 12.3.66 GS\_FILE\_STAT

**GS\_FILE\_STAT** records statistics about data file I/Os to indicate I/O performance and detect performance problems such as abnormal I/O operations.

**Table 12-156** GS\_FILE\_STAT columns

| Name    | Type | Description     |
|---------|------|-----------------|
| filenum | oid  | File identifier |
| dbid    | oid  | Database ID     |

| Name      | Type   | Description                                              |
|-----------|--------|----------------------------------------------------------|
| spcid     | oid    | Tablespace ID                                            |
| phyrds    | bigint | Number of times of reading physical files                |
| phywrts   | bigint | Number of times of writing into physical files           |
| phyblkrd  | bigint | Number of times of reading physical file blocks          |
| phyblkwrt | bigint | Number of times of writing into physical file blocks     |
| readtim   | bigint | Total duration of reading, in microseconds               |
| writetim  | bigint | Total duration of writing, in microseconds               |
| avgiotim  | bigint | Average duration of reading and writing, in microseconds |
| lstiotim  | bigint | Duration of the last file reading, in microseconds       |
| miniotim  | bigint | Minimum duration of reading and writing, in microseconds |
| maxiowtm  | bigint | Maximum duration of reading and writing, in microseconds |

### 12.3.67 GS\_GET\_CONTROL\_GROUP\_INFO

This view is not supported in the centralized system.

### 12.3.68 GS\_GSC\_MEMORY\_DETAIL

GS\_GSC\_MEMORY\_DETAIL displays the memory usage of the global system cache of the current process on the current node. The data is displayed only when GSC is enabled.

Note that the query is separated by the database memory context. Therefore, some memory statistics are missing. The memory context corresponding to the missing memory statistics is **GlobalSysDBCache**.

**Table 12-157** GS\_GSC\_MEMORY\_DETAIL columns

| Name      | Type    | Description                                |
|-----------|---------|--------------------------------------------|
| db_id     | text    | Database ID.                               |
| totalsize | numeric | Total size of the shared memory, in bytes. |

| Name     | Type    | Description                                    |
|----------|---------|------------------------------------------------|
| freesize | numeric | Remaining size of the shared memory, in bytes. |
| usedsize | numeric | Used size of the shared memory, in bytes.      |

### 12.3.69 GS\_INSTANCE\_TIME

**GS\_INSTANCE\_TIME** records time consumption information of the current node. The time consumption information is classified into the following types:

- DB\_TIME: effective time spent by jobs in multi-core scenarios.
- CPU\_TIME: CPU time cost.
- EXECUTION\_TIME: time spent in the executor.
- PARSE\_TIME: time spent on parsing SQL statements.
- PLAN\_TIME: time spent on generating plans.
- REWRITE\_TIME: time spent on rewriting SQL statements.
- PL\_EXECUTION\_TIME: execution time of the PL/SQL stored procedure.
- PL\_COMPILATION\_TIME: compilation time of the PL/SQL stored procedure.
- NET\_SEND\_TIME: time spent on the network.
- DATA\_IO\_TIME: time spent on I/Os.

**Table 12-158** GS\_INSTANCE\_TIME columns

| Name      | Type    | Description            |
|-----------|---------|------------------------|
| stat_id   | integer | Statistics ID          |
| stat_name | text    | Type name              |
| value     | bigint  | Time value, in $\mu$ s |

### 12.3.70 GS\_LABELS

**GS\_LABELS** displays all configured resource labels. Only users with the SYSADMIN or POLADMIN permission can access this view.

**Table 12-159** GS\_LABELS columns

| Name      | Type | Description                                                                                                                          |
|-----------|------|--------------------------------------------------------------------------------------------------------------------------------------|
| labelname | name | Resource label name.                                                                                                                 |
| labeltype | name | Resource label type This parameter corresponds to the <b>labeltype</b> column in the <a href="#">GS_POLICY_LABEL</a> system catalog. |



| Name       | Type | Description                                                                                                          |
|------------|------|----------------------------------------------------------------------------------------------------------------------|
| fqdtype    | name | Database resource type. For example, table, schema, or index.                                                        |
| schemaname | name | Name of the schema to which the database resource belongs.                                                           |
| fqdnname   | name | Database resource name.                                                                                              |
| columnname | name | Name of the database resource column. If the marked database resource is not a column, this parameter is left blank. |

### 12.3.71 GS\_LSC\_MEMORY\_DETAIL

GS\_LSC\_MEMORY\_DETAIL displays statistics about the local SysCache memory usage of all threads based on the MemoryContext node. The statistics are available only when GSC is enabled.

**Table 12-160** GS\_LSC\_MEMORY\_DETAIL columns

| Name        | Type     | Description                                                                              |
|-------------|----------|------------------------------------------------------------------------------------------|
| threadid    | text     | Thread start time and ID (string: <i>timestamp.sessionid</i> ).                          |
| tid         | bigint   | Thread ID.                                                                               |
| thrdtype    | text     | Thread type. It can be any thread type in the system, such as postgresql and wlmmonitor. |
| contextname | text     | Name of the memory context.                                                              |
| level       | smallint | Hierarchy of the memory context.                                                         |
| parent      | text     | Name of the parent memory context.                                                       |
| totalsize   | bigint   | Total size of the memory context, in bytes.                                              |
| freesize    | bigint   | Total size of released memory in the memory context, in bytes.                           |
| usedsize    | bigint   | Total size of used memory in the memory context, in bytes.                               |

### 12.3.72 GS\_MASKING

GS\_MASKING displays all configured dynamic masking policies. Only users with the SYSADMIN or POLADMIN permission can access this view.

**Table 12-161** GS\_MATVIEWS columns

| Name           | Type    | Description                                              |
|----------------|---------|----------------------------------------------------------|
| polname        | name    | Name of the masking policy                               |
| polenabed      | Boolean | Specifies whether to enable the masking policy.          |
| maskaction     | name    | Masking function                                         |
| labelname      | name    | Name of the label to which the masking function applies. |
| masking_object | text    | Masking database resource object                         |
| filter_name    | text    | Logical expression of a filter criterion                 |

### 12.3.73 GS\_MATVIEWS

GS\_MATVIEWS displays information about each materialized view in the database.

**Table 12-162** GS\_MATVIEWS columns

| Name         | Type    | Reference                              | Description                                                                                                   |
|--------------|---------|----------------------------------------|---------------------------------------------------------------------------------------------------------------|
| schemaname   | name    | <a href="#">PG_NAMESPACE</a> .nspname  | Name of the schema of a materialized view                                                                     |
| matviewname  | name    | <a href="#">PG_CLASS</a> .relname      | Name of a materialized view                                                                                   |
| matviewowner | name    | <a href="#">PG_AUTHID</a> .rolname     | Owner of a materialized view                                                                                  |
| tablespace   | name    | <a href="#">PG_TABLESPACE</a> .spcname | Tablespace name of a materialized view. If the default tablespace of the database is used, the value is null. |
| hasindexes   | boolean | -                                      | This column is true if a materialized view has (or has recently had) any indexes.                             |
| definition   | text    | -                                      | Definition of a materialized view (a reconstructed SELECT query)                                              |

### 12.3.74 GS\_OS\_RUN\_INFO

GS\_OS\_RUN\_INFO displays the running status of the OS.

**Table 12-163** GS\_OS\_RUN\_INFO columns

| Name       | Type    | Description                                                         |
|------------|---------|---------------------------------------------------------------------|
| id         | integer | ID                                                                  |
| name       | text    | Name of the OS running status.                                      |
| value      | numeric | Value of the OS running status.                                     |
| comments   | text    | Remarks of the OS running status.                                   |
| cumulative | boolean | Specifies whether the value of the OS running status is cumulative. |

## 12.3.75 GS\_REDO\_STAT

GS\_REDO\_STAT displays the log replay of session threads.

**Table 12-164** GS\_REDO\_STAT columns

| Name      | Type   | Description                                                               |
|-----------|--------|---------------------------------------------------------------------------|
| phywrts   | bigint | Number of times that data is written during log replay.                   |
| phyblkwrt | bigint | Number of data blocks written during log replay.                          |
| writetim  | bigint | Total time required for writing data during log replay.                   |
| avgiotim  | bigint | Average time required for writing data during log replay.                 |
| lstiotim  | bigint | Time consumed by the last data write operation during log replay.         |
| miniotim  | bigint | Minimum time consumed by a single data write operation during log replay. |
| maxiowtm  | bigint | Maximum time consumed by a single data write operation during log replay. |

## 12.3.76 GS\_SESSION\_MEMORY

GS\_SESSION\_MEMORY displays the memory usage at the session level, including all the memory allocated to GaussDB and Stream threads on DN for jobs currently executed by users. If the GUC parameter **enable\_memory\_limit** is set to **off**, this view is unavailable.

**Table 12-165** GS\_SESSION\_MEMORY columns

| Name     | Type    | Description                                                                           |
|----------|---------|---------------------------------------------------------------------------------------|
| sessid   | text    | Thread start time and ID                                                              |
| init_mem | integer | Memory allocated to the currently executed jobs before they enter the executor, in MB |
| used_mem | integer | Memory allocated to the currently executed jobs, in MB                                |
| peak_mem | integer | Peak memory allocated to the currently executed jobs, in MB                           |

### 12.3.77 GS\_SESSION\_MEMORY\_CONTEXT

GS\_SESSION\_MEMORY\_CONTEXT displays the memory usage of all sessions based on the MemoryContext node. If the GUC parameter **enable\_memory\_limit** or **enable\_thread\_pool** is set to **off**, this view is unavailable.

The memory context **TempSmallContextGroup** collects information about all memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the number of the collected memory contexts is recorded in the **usedsize** column. Therefore, the **totalsize** and **freesize** columns for **TempSmallContextGroup** in the view display the corresponding information about all the memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the **usedsize** column displays the number of these memory contexts.

**Table 12-166** GS\_SESSION\_MEMORY\_CONTEXT columns

| Name        | Type     | Description                                                      |
|-------------|----------|------------------------------------------------------------------|
| sessid      | text     | Session start time and ID (string: <i>timestamp.sessionid</i> ). |
| threadid    | bigint   | ID of the thread bound to a session (-1 if no thread is bound).  |
| contextname | text     | Name of the memory context.                                      |
| level       | smallint | Hierarchy of the memory context.                                 |
| parent      | text     | Name of the parent memory context.                               |
| totalsize   | bigint   | Total size of the memory context, in bytes.                      |
| freesize    | bigint   | Total size of released memory in the memory context, in bytes.   |

| Name     | Type   | Description                                                                                                                                               |
|----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| usedsize | bigint | Size of used memory in the memory context, in bytes. For <b>TempSmallContextGroup</b> , this parameter specifies the number of collected memory contexts. |

### 12.3.78 GS\_SESSION\_MEMORY\_DETAIL

GS\_SESSION\_MEMORY\_DETAIL displays the session memory usage based on the MemoryContext node. When **enable\_thread\_pool** is set to **on**, this view contains memory usage of all threads and sessions. If the GUC parameter **enable\_memory\_limit** is set to **off**, this view is unavailable.

The memory context **TempSmallContextGroup** collects information about all memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the number of the collected memory contexts is recorded in the **usedsize** column. Therefore, the **totalsize** and **freesize** columns for **TempSmallContextGroup** in the view display the corresponding information about all the memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the **usedsize** column displays the number of these memory contexts.

You can run the `SELECT * FROM gs_session_memctx_detail(threadid, '');` statement to record information about all memory contexts of a thread into the *threadid\_timestamp.log* file in the `$GAUSSLOG/pg_log/${node_name}/dumppmem` directory. *threadid* can be obtained from **sessid** in the following table.

**Table 12-167** GS\_SESSION\_MEMORY\_DETAIL columns

| Name        | Type     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sessid      | text     | <ol style="list-style-type: none"><li>1. When the thread pool is disabled (<b>enable_thread_pool</b> is set to <b>off</b>), this column indicates the thread start time + session ID (string: <i>timestamp.sessionid</i>).</li><li>2. When the thread pool is enabled (<b>enable_thread_pool</b> is set to <b>on</b>): If the memory context is at the thread level, this column indicates the thread start time + thread ID (string: <i>timestamp.threadid</i>). If the memory context is at the session level, the column indicates the thread start time + session ID (string: <i>timestamp.sessionid</i>).</li></ol> |
| sesstype    | text     | Thread name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| contextname | text     | Name of the memory context.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| level       | smallint | Hierarchy of the memory context.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| parent      | text     | Name of the parent memory context.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

| Name      | Type   | Description                                                                                                                                               |
|-----------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| totalsize | bigint | Total size of the memory context, in bytes.                                                                                                               |
| freesize  | bigint | Total size of released memory in the memory context, in bytes.                                                                                            |
| usedsize  | bigint | Size of used memory in the memory context, in bytes. For <b>TempSmallContextGroup</b> , this parameter specifies the number of collected memory contexts. |

### 12.3.79 GS\_SESSION\_STAT

GS\_SESSION\_STAT displays the session states based on session threads or the AutoVacuum thread.

**Table 12-168** GS\_SESSION\_STAT columns

| Name     | Type    | Description                      |
|----------|---------|----------------------------------|
| sessid   | text    | Thread ID and start time.        |
| statid   | integer | Statistics ID.                   |
| statname | text    | Name of the statistics session.  |
| statunit | text    | Unit of the statistics session.  |
| value    | bigint  | Value of the statistics session. |

### 12.3.80 GS\_SESSION\_TIME

GS\_SESSION\_TIME displays the running time of session threads and time consumed in each execution phase.

**Table 12-169** GS\_SESSION\_TIME columns

| Name      | Type    | Description                      |
|-----------|---------|----------------------------------|
| sessid    | text    | Thread ID and start time.        |
| stat_id   | integer | ID of the statistics session.    |
| stat_name | text    | Session type.                    |
| value     | bigint  | Value of the statistics session. |

## 12.3.81 GS\_SQL\_COUNT

GS\_SQL\_COUNT displays statistics about five types of running statements (SELECT, INSERT, UPDATE, DELETE, and MERGE INTO) on the current node of the database.

- When a common user queries the GS\_SQL\_COUNT view, statistics about the current node of the user are displayed. When an administrator queries the GS\_SQL\_COUNT view, statistics about the current node of all users are displayed.
- When the database or the node is restarted, the statistics are cleared and will be measured again.
- The system counts when a node receives a query, including a query inside the database.

**Table 12-170** GS\_SQL\_COUNT columns

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| node_name           | name   | Node name.                                                  |
| user_name           | name   | Username.                                                   |
| select_count        | bigint | Statistical result of the SELECT statement.                 |
| update_count        | bigint | Statistical result of the UPDATE statement.                 |
| insert_count        | bigint | Statistical result of the INSERT statement.                 |
| delete_count        | bigint | Statistical result of the DELETE statement.                 |
| mergeinto_count     | bigint | Statistical result of the <b>MERGE INTO</b> statement.      |
| ddl_count           | bigint | Number of DDL statements.                                   |
| dml_count           | bigint | Number of DML statements.                                   |
| dcl_count           | bigint | Number of DML statements.                                   |
| total_select_elapse | bigint | Total response time of SELECT statements (unit: $\mu$ s).   |
| avg_select_elapse   | bigint | Average response time of SELECT statements (unit: $\mu$ s). |
| max_select_elapse   | bigint | Maximum response time of SELECT statements (unit: $\mu$ s). |
| min_select_elapse   | bigint | Minimum response time of SELECT statements (unit: $\mu$ s). |
| total_update_elapse | bigint | Total response time of UPDATE statements (unit: $\mu$ s).   |

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| avg_update_elapse   | bigint | Average response time of UPDATE statements (unit: $\mu$ s). |
| max_update_elapse   | bigint | Maximum response time of UPDATE statements (unit: $\mu$ s). |
| min_update_elapse   | bigint | Minimum response time of UPDATE statements (unit: $\mu$ s). |
| total_insert_elapse | bigint | Total response time of INSERT statements (unit: $\mu$ s).   |
| avg_insert_elapse   | bigint | Average response time of INSERT statements (unit: $\mu$ s). |
| max_insert_elapse   | bigint | Maximum response time of INSERT statements (unit: $\mu$ s). |
| min_insert_elapse   | bigint | Minimum response time of INSERT statements (unit: $\mu$ s). |
| total_delete_elapse | bigint | Total response time of DELETE statements (unit: $\mu$ s).   |
| avg_delete_elapse   | bigint | Average response time of DELETE statements (unit: $\mu$ s). |
| max_delete_elapse   | bigint | Maximum response time of DELETE statements (unit: $\mu$ s). |
| min_delete_elapse   | bigint | Minimum response time of DELETE statements (unit: $\mu$ s). |

### 12.3.82 GS\_STAT\_SESSION\_CU

GS\_STAT\_SESSION\_CU queries the CU hit rate of running sessions on each node of the entire database. The data about a session is cleared after you exit this session or restart the database.

**Table 12-171** GS\_STAT\_SESSION\_CU columns

| Name          | Type    | Description                             |
|---------------|---------|-----------------------------------------|
| mem_hit       | integer | Number of memory hits.                  |
| hdd_sync_read | integer | Number of synchronous hard disk reads.  |
| hdd_asyn_read | integer | Number of asynchronous hard disk reads. |



### 12.3.83 GS\_THREAD\_MEMORY\_CONTEXT

GS\_THREAD\_MEMORY\_CONTEXT displays the memory usage of all threads based on the MemoryContext node. This view is equivalent to the GS\_SESSION\_MEMORY\_DETAIL view when **enable\_thread\_pool** is set to **off**. If the GUC parameter **enable\_memory\_limit** is set to **off**, this view is unavailable.

The memory context **TempSmallContextGroup** collects information about all memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the number of the collected memory contexts is recorded in the **usedsize** column. Therefore, the **totalsize** and **freesize** columns for **TempSmallContextGroup** in the view display the corresponding information about all the memory contexts whose value in the **totalsize** column is less than 8192 bytes in the current thread, and the **usedsize** column displays the number of these memory contexts.

**Table 12-172** GS\_THREAD\_MEMORY\_CONTEXT columns

| Name        | Type     | Description                                                                                                                                               |
|-------------|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| threadid    | text     | Thread start time and ID (string: <i>timestamp.sessionid</i> ).                                                                                           |
| tid         | bigint   | Thread ID.                                                                                                                                                |
| thrdtype    | text     | Thread type.                                                                                                                                              |
| contextname | text     | Name of the memory context.                                                                                                                               |
| level       | smallint | Hierarchy of the memory context.                                                                                                                          |
| parent      | text     | Name of the parent memory context.                                                                                                                        |
| totalsize   | bigint   | Total size of the memory context, in bytes.                                                                                                               |
| freesize    | bigint   | Total size of released memory in the memory context, in bytes.                                                                                            |
| usedsize    | bigint   | Size of used memory in the memory context, in bytes. For <b>TempSmallContextGroup</b> , this parameter specifies the number of collected memory contexts. |

### 12.3.84 GS\_TOTAL\_MEMORY\_DETAIL

GS\_TOTAL\_MEMORY\_DETAIL displays the memory usage of the current database node, in MB. If the GUC parameter **enable\_memory\_limit** is set to **off**, this view is unavailable.

**Table 12-173** GS\_TOTAL\_MEMORY\_DETAIL columns

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| nodename    | text    | Node name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| memorytype  | text    | Memory type. The value must be one of the following: <ul style="list-style-type: none"><li>• <b>max_process_memory</b>: memory occupied by the GaussDB instance.</li><li>• <b>process_used_memory</b>: memory occupied by the GaussDB process.</li><li>• <b>max_dynamic_memory</b>: maximum dynamic memory.</li><li>• <b>dynamic_used_memory</b>: used dynamic memory.</li><li>• <b>dynamic_peak_memory</b>: dynamic peak memory.</li><li>• <b>dynamic_used_shrctx</b>: maximum dynamic shared memory context.</li><li>• <b>dynamic_peak_shrctx</b>: dynamic peak value of the shared memory context.</li><li>• <b>max_backend_memory</b>: maximum memory that can be used when the HA port is used to execute services.</li><li>• <b>backend_used_memory</b>: memory that has been used when the HA port is used to execute services.</li><li>• <b>max_shared_memory</b>: maximum shared memory.</li><li>• <b>shared_used_memory</b>: used shared memory.</li><li>• <b>max_sctpcomm_memory</b>: maximum memory allowed for the communications library.</li><li>• <b>sctpcomm_used_memory</b>: memory used by the communications library.</li><li>• <b>sctpcomm_peak_memory</b>: memory peak of the communications library.</li><li>• <b>other_used_memory</b>: other used memory.</li></ul> |
| memorybytes | integer | Size of allocated memory-typed memory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

### 12.3.85 GS\_TOTAL\_NODEGROUP\_MEMORY\_DETAIL

GS\_TOTAL\_NODEGROUP\_MEMORY\_DETAIL returns the memory usage (in MB) of the current logical instance of the database. If the GUC parameter **enable\_memory\_limit** is set to **off**, this view is unavailable.

**Table 12-174** GS\_TOTAL\_NODEGROUP\_MEMORY\_DETAIL columns

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ngname      | text    | Logical instance name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| memorytype  | text    | Memory type. The value must be one of the following: <ul style="list-style-type: none"><li>• <b>ng_total_memory</b>: total memory of the logical instance.</li><li>• <b>ng_used_memory</b>: memory usage of the logical instance.</li><li>• <b>ng_estimate_memory</b>: estimated memory usage of the logical instance.</li><li>• <b>ng_foreignrp_memsize</b>: total memory of the external resource pool of the logical instance.</li><li>• <b>ng_foreignrp_usesize</b>: memory usage of the external resource pool of the logical instance.</li><li>• <b>ng_foreignrp_peaksize</b>: peak memory usage of the external resource pool of the logical instance.</li><li>• <b>ng_foreignrp_mempct</b>: percentage of the external resource pool of the logical instance to the total memory of the logical instance.</li><li>• <b>ng_foreignrp_estmsize</b>: estimated memory usage of the external resource pool of the logical instance.</li></ul> |
| memorybytes | integer | Size of allocated memory-typed memory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## 12.3.86 GV\_SESSION

GV\_SESSION displays all session information about the current session.

**Table 12-175** GV\_SESSION columns

| Name       | Type    | Description                                            |
|------------|---------|--------------------------------------------------------|
| sid        | bigint  | Name of the database where the hotspot key is located. |
| serial#    | integer | Name of the schema where the hotspot key is located.   |
| schemaname | name    | Name of the table where the hotspot key is located.    |
| user#      | oid     | Value of the hotspot key.                              |

| Name           | Type      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| username       | name      | Hash value of the hotspot key in the database. If the table is a list or range distributed table, the value of this column is <b>0</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| machine        | text      | Frequency of accessing the hotspot key.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| sql_id         | bigint    | SQL OID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| client_info    | text      | Client information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| event          | text      | Queuing status of a statement. The value can be: <ul style="list-style-type: none"><li>• <b>waiting in queue</b>: The statement is in the queue.</li><li>• <b>Empty</b>: The statement is running.</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| sql_exec_start | timestamp | Start time of SQL statement execution.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| program        | text      | Name of the application connected to the backend.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| status         | text      | Overall status of this backend. The value can be: <ul style="list-style-type: none"><li>• <b>active</b>: The backend is executing a query.</li><li>• <b>idle</b>: The backend is waiting for a new client command.</li><li>• <b>idle in transaction</b>: The backend is in a transaction, but there is no statement being executed in the transaction.</li><li>• <b>idle in transaction (aborted)</b>: The backend is in a transaction, but there are statements failed in the transaction.</li><li>• <b>fastpath function call</b>: The backend is executing a fast-path function.</li><li>• <b>disabled</b>: This state is reported if <b>track_activities</b> is disabled in this backend.</li></ul> |

## 12.3.87 MPP\_TABLES

The following information is displayed in the MPP\_TABLES view.

**Table 12-176** MPP\_TABLES columns

| Name       | Type | Description                               |
|------------|------|-------------------------------------------|
| schemaname | name | Name of the schema that contains a table. |
| tablename  | name | Table name.                               |

| Name       | Type             | Description                          |
|------------|------------------|--------------------------------------|
| tableowner | name             | Table owner.                         |
| tablespace | name             | Tablespace containing the table.     |
| pgroup     | name             | Name of a node group.                |
| nodeoids   | oidvector_extend | List of distributed table node OIDs. |

## 12.3.88 MY\_COL\_COMMENTS

MY\_COL\_COMMENTS displays column comments of the table accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-177** MY\_COL\_COMMENTS columns

| Name        | Type                  | Description         |
|-------------|-----------------------|---------------------|
| owner       | character varying(64) | Constraint creator. |
| owner       | character varying(64) | Table owner.        |
| table_name  | character varying(64) | Table name.         |
| column_name | character varying(64) | Column name.        |
| comments    | text                  | Comments.           |

## 12.3.89 MY\_CONS\_COLUMNS

MY\_CONS\_COLUMNS displays information about primary key constraint columns in tables accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-178** MY\_CONS\_COLUMNS columns

| Name        | Type                  | Description                          |
|-------------|-----------------------|--------------------------------------|
| owner       | character varying(64) | Constraint creator.                  |
| table_name  | character varying(64) | Name of a constraint-related table.  |
| column_name | character varying(64) | Name of a constraint-related column. |

| Name            | Type                  | Description                      |
|-----------------|-----------------------|----------------------------------|
| constraint_name | character varying(64) | Constraint name.                 |
| position        | smallint              | Position of a column in a table. |

## 12.3.90 MY\_CONSTRAINTS

MY\_CONSTRAINTS displays table constraint information accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-179** MY\_CONSTRAINTS columns

| Name            | Type                   | Description                                                                                                                                                                                                                   |
|-----------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | character varying(64)  | Constraint creator.                                                                                                                                                                                                           |
| constraint_name | vcharacter varying(64) | Constraint name.                                                                                                                                                                                                              |
| constraint_type | text                   | Constraint type. <ul style="list-style-type: none"><li>• <b>c</b>: check constraint.</li><li>• <b>f</b>: foreign key constraint.</li><li>• <b>p</b>: primary key constraint.</li><li>• <b>u</b>: unique constraint.</li></ul> |
| table_name      | character varying(64)  | Name of a constraint-related table.                                                                                                                                                                                           |
| index_owner     | character varying(64)  | Owner of a constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                              |
| index_name      | character varying(64)  | Name of the constraint-related index (only for the unique constraint and primary key constraint).                                                                                                                             |

## 12.3.91 MY\_IND\_COLUMNS

MY\_IND\_COLUMNS displays column information about all indexes accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-180** MY\_IND\_COLUMNS columns

| Name        | Type                  | Description  |
|-------------|-----------------------|--------------|
| index_owner | character varying(64) | Index owner. |

| Name            | Type                  | Description                          |
|-----------------|-----------------------|--------------------------------------|
| index_name      | character varying(64) | Index name.                          |
| table_owner     | character varying(64) | Table owner.                         |
| table_name      | character varying(64) | Table name.                          |
| column_name     | name                  | Column name.                         |
| column_position | smallint              | Position of the column in the index. |

## 12.3.92 MY\_IND\_EXPRESSIONS

MY\_IND\_EXPRESSIONS displays information about function-based expression indexes accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-181** MY\_IND\_EXPRESSIONS columns

| Name              | Type                  | Description                                            |
|-------------------|-----------------------|--------------------------------------------------------|
| table_owner       | character varying(64) | Table owner.                                           |
| table_name        | character varying(64) | Table name.                                            |
| index_owner       | character varying(64) | Index owner.                                           |
| index_name        | character varying(64) | Index name.                                            |
| column_expression | text                  | Function-based index expression of a specified column. |
| column_position   | smallint              | Position of the column in the index.                   |

## 12.3.93 MY\_IND\_PARTITIONS

MY\_IND\_PARTITIONS displays the partition information of indexes (excluding global indexes of partitioned tables) accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-182** MY\_IND\_PARTITIONS columns

| Name        | Type                  | Description                                                                           |
|-------------|-----------------------|---------------------------------------------------------------------------------------|
| index_owner | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs. |

| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index_name             | character varying(64) | Name of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                             |
| partition_name         | character varying(64) | Name of the index partition.                                                                                                                                                                                                                                                                                                                                                          |
| def_tablespace_name    | name                  | Tablespace name of the index partition.                                                                                                                                                                                                                                                                                                                                               |
| high_value             | text                  | Limit of the partition corresponding to the index partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li><b>t</b> (true): yes.</li><li><b>f</b> (false): no.</li></ul>                                                                                                                                                                                                                               |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                           |
| high_value_length      | integer               | Character length of the limit of the partition corresponding to an index partition.                                                                                                                                                                                                                                                                                                   |

### 12.3.94 MY\_IND\_SUBPARTITIONS

**MY\_IND\_SUBPARTITIONS** displays information about level-2 partitions of indexes (excluding global indexes of partitioned tables) owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-183** MY\_IND\_SUBPARTITIONS columns

| Name        | Type                  | Description                                                                           |
|-------------|-----------------------|---------------------------------------------------------------------------------------|
| index_owner | character varying(64) | Name of the owner of the partitioned table index to which an index partition belongs. |



| Name                   | Type                  | Description                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| index_name             | character varying(64) | Name of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                                  |
| partition_name         | character varying(64) | Name of the partition where an index is located                                                                                                                                                                                                                                                                                                                                            |
| subpartition_name      | character varying(64) | Name of the level-2 partition where an index is located.                                                                                                                                                                                                                                                                                                                                   |
| def_tablespace_name    | name                  | Tablespace name of an index partition.                                                                                                                                                                                                                                                                                                                                                     |
| high_value             | text                  | Limit of the partition corresponding to an index partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>• For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>• For list partitioning, the value list of each partition is displayed.</li><li>• For hash partitioning, the number of each partition is displayed.</li></ul> |
| index_partition_usable | Boolean               | Specifies whether the index partition is available. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                                                                                                                                                                                                                                  |
| schema                 | character varying(64) | Schema of the partitioned table index to which the index partition belongs.                                                                                                                                                                                                                                                                                                                |
| high_value_length      | integer               | Character length of the limit of the partition corresponding to an index partition.                                                                                                                                                                                                                                                                                                        |

### 12.3.95 MY\_INDEXES

MY\_INDEXES displays index information about the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-184** MY\_INDEXES columns

| Name        | Type                  | Description                                                         |
|-------------|-----------------------|---------------------------------------------------------------------|
| owner       | character varying(64) | Index owner.                                                        |
| index_name  | character varying(64) | Index name.                                                         |
| table_name  | character varying(64) | Name of the table corresponding to the index.                       |
| uniqueness  | text                  | Specifies whether the index is unique.                              |
| partitioned | character(3)          | Specifies whether the index has the property of partitioned tables. |
| generated   | character varying(1)  | Specifies whether the index name is generated by the system.        |

## 12.3.96 MY\_JOBS

MY\_JOBS displays details about the scheduled jobs owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-185** MY\_JOBS columns

| Name       | Type                        | Description                                         |
|------------|-----------------------------|-----------------------------------------------------|
| job        | bigint                      | Job ID.                                             |
| log_user   | name                        | Username of a job creator.                          |
| priv_user  | name                        | Username of a job executor.                         |
| dbname     | name                        | Name of the database where the job is created.      |
| start_date | timestamp without time zone | Job start time.                                     |
| start_suc  | text                        | Start time of the successful job execution.         |
| last_date  | timestamp without time zone | Start time of the last job execution.               |
| last_suc   | text                        | Start time of the last successful job execution.    |
| this_date  | timestamp without time zone | Start time of the ongoing job execution.            |
| this_suc   | text                        | Start time of the ongoing successful job execution. |

| Name      | Type                        | Description                                                                                                                                                                                                                                                                               |
|-----------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| next_date | timestamp without time zone | Schedule time of the next job execution.                                                                                                                                                                                                                                                  |
| next_suc  | text                        | Schedule time of the next successful job execution.                                                                                                                                                                                                                                       |
| broken    | text                        | If the value of the <b>status</b> column is 'd', the value is 'y'. Otherwise, the value is 'n'.                                                                                                                                                                                           |
| status    | "char"                      | Status of the current job. The value can be 'r', 's', 'f', or 'd'. The default value is 'r'. <ul style="list-style-type: none"><li>• <b>r</b>: running</li><li>• <b>s</b>: successful execution</li><li>• <b>f</b>: failed execution</li><li>• <b>d</b>: canceling an execution</li></ul> |
| interval  | text                        | Time expression used to calculate the next time the job will be executed. If this parameter is set to <b>null</b> , the job will be executed once only.                                                                                                                                   |
| failures  | smallint                    | Number of times the job has started and failed. If a job fails to be executed for 16 consecutive times, no more attempt will be made on it.                                                                                                                                               |
| what      | text                        | Executable job.                                                                                                                                                                                                                                                                           |

### 12.3.97 MY\_OBJECTS

MY\_OBJECTS displays database objects accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-186** MY\_OBJECTS columns

| Name        | Type | Description                                                                    |
|-------------|------|--------------------------------------------------------------------------------|
| object_name | name | Object name.                                                                   |
| object_id   | oid  | Object OID.                                                                    |
| object_type | name | Object type ( <b>TABLE</b> , <b>INDEX</b> , <b>SEQUENCE</b> , or <b>VIEW</b> ) |
| namespace   | oid  | Namespace that the object belongs to.                                          |

| Name          | Type                     | Description                           |
|---------------|--------------------------|---------------------------------------|
| created       | timestamp with time zone | Creation time of the object.          |
| last_ddl_time | timestamp with time zone | Last modification time of the object. |

**NOTICE**

For details about the value ranges of **created** and **last\_ddl\_time**, see [PG\\_OBJECT](#).

## 12.3.98 MY\_PART\_INDEXES

MY\_PART\_INDEXES displays information about partitioned table indexes accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-187** MY\_PART\_INDEXES columns

| Name                   | Type                  | Description                                                                                                                                                                                       |
|------------------------|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| def_tablespace_name    | name                  | Tablespace name of a partitioned table index                                                                                                                                                      |
| index_owner            | character varying(64) | Owner name of a partitioned table index                                                                                                                                                           |
| index_name             | character varying(64) | Name of a partitioned table index                                                                                                                                                                 |
| partition_count        | bigint                | Number of index partitions of a partitioned table index                                                                                                                                           |
| partitioning_key_count | integer               | Number of partition keys of a partitioned table                                                                                                                                                   |
| partitioning_type      | text                  | Partitioning policy of a partitioned table<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> . |
| schema                 | character varying(64) | Schema of a partitioned table index                                                                                                                                                               |
| table_name             | character varying(64) | Name of the partitioned table to which a partitioned table index belongs                                                                                                                          |

| Name                      | Type    | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subpartitioning_type      | text    | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer | Number of level-2 partition keys of the partitioned table.                                                                                                                                                                                                                                               |

## 12.3.99 MY\_PART\_TABLES

MY\_PART\_TABLES displays information about partitioned tables accessible to the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-188** MY\_PART\_TABLES columns

| Name                   | Type                  | Description                                                                                                                                                                                          |
|------------------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_owner            | character varying(64) | Owner name of a partitioned table.                                                                                                                                                                   |
| table_name             | character varying(64) | Name of the partitioned table.                                                                                                                                                                       |
| partitioning_type      | text                  | Partitioning policy of the partitioned table.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current partitioned table, see <a href="#">CREATE TABLE PARTITION</a> . |
| partition_count        | bigint                | Number of partitions of the partitioned table.                                                                                                                                                       |
| partitioning_key_count | integer               | Number of partition keys of the partitioned table.                                                                                                                                                   |
| def_tablespace_name    | name                  | Tablespace name of the partitioned table.                                                                                                                                                            |
| schema                 | character varying(64) | Schema of the partitioned table.                                                                                                                                                                     |

| Name                      | Type    | Description                                                                                                                                                                                                                                                                                              |
|---------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| subpartitioning_type      | text    | Partitioning policy of a level-2 partitioned table. If the partitioned table is a level-1 partitioned table, <b>NONE</b> is displayed.<br><b>NOTE</b><br>For details about the supported partitioning policies of the current level-2 partitioned table, see <a href="#">CREATE TABLE SUBPARTITION</a> . |
| def_subpartition_count    | integer | Default number of level-2 partitions to be created. The value is <b>1</b> for a level-2 partitioned table and <b>0</b> for a level-1 partitioned table.                                                                                                                                                  |
| subpartitioning_key_count | integer | Number of level-2 partition keys of the partitioned table.                                                                                                                                                                                                                                               |

## 12.3.100 MY\_PROCEEDURES

MY\_PROCEEDURES displays information about stored procedures or functions owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-189** MY\_PROCEEDURES columns

| Name            | Type                  | Description                                       |
|-----------------|-----------------------|---------------------------------------------------|
| owner           | character varying(64) | Owner of a stored procedure or function.          |
| object_name     | character varying(64) | Name of a stored procedure or function.           |
| argument_number | smallint              | Number of input parameters in a stored procedure. |

## 12.3.101 MY\_SEQUENCES

MY\_SEQUENCES displays the sequence information about the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-190** MY\_SEQUENCES columns

| Name           | Type                  | Description          |
|----------------|-----------------------|----------------------|
| sequence_owner | character varying(64) | Owner of a sequence. |
| sequence_name  | character varying(64) | Name of a sequence.  |

## 12.3.102 MY\_SOURCE

MY\_SOURCE displays information about stored procedures or functions owned by the current user and provides columns defined by the stored procedures or the functions. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-191** MY\_SOURCE columns

| Name  | Type                  | Description                                   |
|-------|-----------------------|-----------------------------------------------|
| owner | character varying(64) | Owner of a stored procedure or function.      |
| name  | character varying(64) | Name of a stored procedure or function.       |
| text  | text                  | Definition of a stored procedure or function. |

## 12.3.103 MY\_SYNONYMS

MY\_SYNONYMS displays synonyms accessible to the current user.

**Table 12-192** MY\_SYNONYMS columns

| Name         | Type | Description                                                                                                                                                                                                              |
|--------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| schema_name  | text | Name of the schema to which the synonym belongs.                                                                                                                                                                         |
| synonym_name | text | Synonym name.                                                                                                                                                                                                            |
| table_owner  | text | Owner of the associated object. Although the column is called <b>table_owner</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym. |

| Name              | Type | Description                                                                                                                                                                                                                          |
|-------------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| table_name        | text | Name of the associated object. Although the column is called <b>table_name</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym.               |
| table_schema_name | text | Schema name of the associated object. Although the column is called <b>table_schema_name</b> , its associated object is not necessarily a table. It can be any common database object, such as a view, stored procedure, or synonym. |

### 12.3.104 MY\_TAB\_COLUMNS

MY\_TAB\_COLUMNS displays information about table columns accessible to the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-193** MY\_TAB\_COLUMNS columns

| Name           | Type                   | Description                                                                                                    |
|----------------|------------------------|----------------------------------------------------------------------------------------------------------------|
| owner          | character varying(64)  | Table owner.                                                                                                   |
| table_name     | character varying(64)  | Table name.                                                                                                    |
| column_name    | character varying(64)  | Column name.                                                                                                   |
| data_type      | character varying(128) | Data type of the column.                                                                                       |
| data_length    | integer                | Length of a column, in bytes.                                                                                  |
| data_precision | integer                | Precision of a data type. This column is valid for the numeric data type and <b>NULL</b> for other data types. |
| data_scale     | integer                | Number of decimal places. This column is valid for the numeric data type and <b>0</b> for other data types.    |
| nullable       | bpchar                 | Specifies whether a column can be empty ( <b>n</b> for a primary key constraint or NOT NULL constraint).       |



| Name        | Type    | Description                                                                                                         |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------|
| column_id   | integer | Sequence number of a column when a table is created.                                                                |
| avg_col_len | numeric | Average length of a column, in bytes.                                                                               |
| char_length | numeric | Length of a column, in characters. This parameter is valid only for the varchar, nvarchar2, bpchar, and char types. |
| comments    | text    | Comments.                                                                                                           |

### 12.3.105 MY\_TAB\_COMMENTS

MY\_TAB\_COMMENTS displays comments on all tables and views owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-194** MY\_TAB\_COMMENTS columns

| Name       | Type                  | Description                |
|------------|-----------------------|----------------------------|
| owner      | character varying(64) | Owner of a table or view.  |
| table_name | character varying(64) | Name of the table or view. |
| comments   | text                  | Comments.                  |

### 12.3.106 MY\_TAB\_PARTITIONS

MY\_TAB\_PARTITIONS displays all table partitions accessible to the current user. Each table partition of a partitioned table accessible to the current user has one record in USER\_TAB\_PARTITIONS. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-195** MY\_TAB\_PARTITIONS columns

| Name           | Type                  | Description                        |
|----------------|-----------------------|------------------------------------|
| table_owner    | character varying(64) | Owner name of a partitioned table. |
| table_name     | character varying(64) | Name of a partitioned table.       |
| partition_name | character varying(64) | Partition name.                    |

| Name               | Type                  | Description                                                                                                                                                                                                                                                                                                                                          |
|--------------------|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| high_value         | text                  | Limit of a partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>• For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>• For list partitioning, the value list of each partition is displayed.</li><li>• For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name    | name                  | Tablespace name of a partition.                                                                                                                                                                                                                                                                                                                      |
| schema             | character varying(64) | Schema of a partitioned table.                                                                                                                                                                                                                                                                                                                       |
| subpartition_count | bigint                | Number of level-2 partitions.                                                                                                                                                                                                                                                                                                                        |
| high_value_length  | integer               | Length of the partition boundary value expression.                                                                                                                                                                                                                                                                                                   |

## 12.3.107 MY\_TAB\_SUBPARTITIONS

MY\_TAB\_SUBPARTITIONS displays information about level-2 partitions owned by the current user. This view exists in both the PG\_CATALOG and SYS schemas.

**Table 12-196** MY\_TAB\_SUBPARTITIONS columns

| Name              | Type                  | Description                  |
|-------------------|-----------------------|------------------------------|
| table_owner       | character varying(64) | Role name.                   |
| table_name        | character varying(64) | Relational table name.       |
| partition_name    | character varying(64) | Partition name.              |
| subpartition_name | character varying(64) | Name of a level-2 partition. |

| Name              | Type                  | Description                                                                                                                                                                                                                                                                                                                                            |
|-------------------|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| high_value        | text                  | Limit of a level-2 partition.<br><b>NOTE</b> <ul style="list-style-type: none"><li>For range partitioning and interval partitioning, the upper limit of each partition is displayed.</li><li>For list partitioning, the value list of each partition is displayed.</li><li>For hash partitioning, the number of each partition is displayed.</li></ul> |
| tablespace_name   | name                  | Tablespace name of a level-2 partitioned table.                                                                                                                                                                                                                                                                                                        |
| schema            | character varying(64) | Name of a namespace.                                                                                                                                                                                                                                                                                                                                   |
| high_value_length | integer               | Length of the boundary value expression in level-2 partition.                                                                                                                                                                                                                                                                                          |

### 12.3.108 MY\_TABLES

MY\_TABLES displays information about the tables owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-197** MY\_TABLES columns

| Name            | Type                  | Description                                                                                                                                           |
|-----------------|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| owner           | character varying(64) | Table owner.                                                                                                                                          |
| table_name      | character varying(64) | Table name.                                                                                                                                           |
| tablespace_name | character varying(64) | Tablespace name of the table.                                                                                                                         |
| dropped         | character varying     | Specifies whether the current record is deleted. <ul style="list-style-type: none"><li><b>yes</b>: deleted.</li><li><b>no</b>: not deleted.</li></ul> |
| num_rows        | numeric               | Estimated number of rows in the table.                                                                                                                |
| status          | character varying(8)  | Specifies whether the current record is valid. <ul style="list-style-type: none"><li><b>valid</b>: valid.</li></ul>                                   |

| Name      | Type         | Description                                                                                                                                                                                           |
|-----------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| temporary | character(1) | Specifies whether the table is a temporary table. <ul style="list-style-type: none"><li>• <b>y</b>: The table is a temporary table.</li><li>• <b>n</b>: The table is not a temporary table.</li></ul> |

## 12.3.109 MY\_TRIGGERS

MY\_TRIGGERS displays information about the triggers owned by the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-198** MY\_TRIGGERS columns

| Name         | Type                  | Description                                                                                                                                                                                                                                                                             |
|--------------|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| trigger_name | character varying(64) | Trigger name.                                                                                                                                                                                                                                                                           |
| table_name   | character varying(64) | Relational table name.                                                                                                                                                                                                                                                                  |
| table_owner  | character varying(64) | Role name.                                                                                                                                                                                                                                                                              |
| status       | character varying(64) | Trigger status. <ul style="list-style-type: none"><li>• <b>O</b>: The trigger is enabled in origin or local mode.</li><li>• <b>D</b>: The trigger is disabled.</li><li>• <b>R</b>: The trigger is enabled in replica mode.</li><li>• <b>A</b>: The trigger is always enabled.</li></ul> |

## 12.3.110 MY\_VIEWS

MY\_VIEWS displays information about all views of the current user. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-199** MY\_VIEWS columns

| Name      | Type                  | Description      |
|-----------|-----------------------|------------------|
| owner     | character varying(64) | Owner of a view. |
| view_name | character varying(64) | View name.       |

## 12.3.111 PG\_CURSORS

PG\_CURSORS displays cursors that are currently available.

**Table 12-200** PG\_CURSORS columns

| Name          | Type                     | Description                                                                                                                                                                                                                      |
|---------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name          | text                     | Cursor name.                                                                                                                                                                                                                     |
| statement     | text                     | Query statement when the cursor is declared to change.                                                                                                                                                                           |
| is_holdable   | boolean                  | Specifies whether the cursor is holdable (that is, whether the cursor can be accessed after the transaction that declared the cursor has committed). If it is, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> . |
| is_binary     | boolean                  | Specifies whether the cursor is declared BINARY. If it is, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> .                                                                                                     |
| is_scrollable | boolean                  | Specifies whether the cursor is scrollable (it allows rows to be retrieved in a nonsequential manner). If it is, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> .                                               |
| creation_time | timestamp with time zone | Timestamp at which the cursor is declared.                                                                                                                                                                                       |

## 12.3.112 PG\_COMM\_DELAY

PG\_COMM\_DELAY displays the communications library delay status for a single DN.

**Table 12-201** PG\_COMM\_DELAY columns

| Name        | Type    | Description                                                                  |
|-------------|---------|------------------------------------------------------------------------------|
| node_name   | text    | Node name                                                                    |
| remote_name | text    | Name of the peer node                                                        |
| remote_host | text    | IP address of the peer node                                                  |
| stream_num  | integer | Number of logical stream connections used by the current physical connection |

| Name      | Type    | Description                                                                                                                                                                               |
|-----------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| min_delay | integer | Minimum delay of the current physical connection within 1 minute, in microsecond.<br><b>NOTE</b><br>A negative result is invalid. Wait until the delay status is updated and query again. |
| average   | integer | Average delay of the current physical connection within 1 minute, in microsecond.                                                                                                         |
| max_delay | integer | Maximum delay of the current physical connection within 1 minute, in microsecond.                                                                                                         |

### 12.3.113 PG\_COMM\_RECV\_STREAM

**PG\_COMM\_RECV\_STREAM** displays the receiving stream status of all the communication libraries for a single DN.

**Table 12-202** PG\_COMM\_RECV\_STREAM columns

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name   | text    | Node name                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| local_tid   | bigint  | ID of the thread using this stream                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| remote_name | text    | Name of the peer node                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| remote_tid  | bigint  | Peer thread ID                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| idx         | integer | Peer DN ID in the local DN                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| sid         | integer | Stream ID in the physical connection                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| tcp_sock    | integer | TCP socket used in the stream                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| state       | text    | Status of the stream <ul style="list-style-type: none"><li>• <b>UNKNOWN</b>: The logical connection status is unknown.</li><li>• <b>READY</b>: The logical connection is ready.</li><li>• <b>RUN</b>: The logical connection sends packets normally.</li><li>• <b>HOLD</b>: The logical connection is waiting to send packets.</li><li>• <b>CLOSED</b>: The logical connection is closed.</li><li>• <b>TO_CLOSED</b>: The logical connection will be closed.</li></ul> |
| query_id    | bigint  | <b>debug_query_id</b> corresponding to the stream                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Name       | Type    | Description                                                      |
|------------|---------|------------------------------------------------------------------|
| pn_id      | integer | <b>plan_node_id</b> of the query executed by the stream          |
| send_smp   | integer | <b>smpid</b> of the sender of the query executed by the stream   |
| recv_smp   | integer | <b>smpid</b> of the receiver of the query executed by the stream |
| recv_bytes | bigint  | Total data volume received from the stream, in byte              |
| time       | bigint  | Current lifecycle service duration of the stream, in ms          |
| speed      | bigint  | Average receiving rate of the stream, in byte/s                  |
| quota      | bigint  | Current communication quota value of the stream, in byte         |
| buff_use   | bigint  | Current size of the data cache of the stream, in byte            |

### 12.3.114 PG\_COMM\_SEND\_STREAM

**PG\_COMM\_SEND\_STREAM** displays the sending stream status of all the communication libraries for a single DN.

**Table 12-203** PG\_COMM\_SEND\_STREAM columns

| Name        | Type    | Description                          |
|-------------|---------|--------------------------------------|
| node_name   | text    | Node name                            |
| local_tid   | bigint  | ID of the thread using this stream   |
| remote_name | text    | Name of the peer node                |
| remote_tid  | bigint  | Peer thread ID                       |
| idx         | integer | Peer DN ID in the local DN           |
| sid         | integer | Stream ID in the physical connection |
| tcp_sock    | integer | TCP socket used in the stream        |

| Name       | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state      | text    | Status of the stream <ul style="list-style-type: none"><li>• <b>UNKNOWN</b>: The logical connection status is unknown.</li><li>• <b>READY</b>: The logical connection is ready.</li><li>• <b>RUN</b>: The logical connection sends packets normally.</li><li>• <b>HOLD</b>: The logical connection is waiting to send packets.</li><li>• <b>CLOSED</b>: The logical connection is closed.</li><li>• <b>TO_CLOSED</b>: The logical connection will be closed.</li></ul> |
| query_id   | bigint  | <b>debug_query_id</b> corresponding to the stream                                                                                                                                                                                                                                                                                                                                                                                                                      |
| pn_id      | integer | <b>plan_node_id</b> of the query executed by the stream                                                                                                                                                                                                                                                                                                                                                                                                                |
| send_smp   | integer | <b>smpid</b> of the sender of the query executed by the stream                                                                                                                                                                                                                                                                                                                                                                                                         |
| recv_smp   | integer | <b>smpid</b> of the receiver of the query executed by the stream                                                                                                                                                                                                                                                                                                                                                                                                       |
| send_bytes | bigint  | Total data volume sent by the stream, in byte                                                                                                                                                                                                                                                                                                                                                                                                                          |
| time       | bigint  | Current lifecycle service duration of the stream, in ms                                                                                                                                                                                                                                                                                                                                                                                                                |
| speed      | bigint  | Average sending rate of the stream, in byte/s                                                                                                                                                                                                                                                                                                                                                                                                                          |
| quota      | bigint  | Current communication quota value of the stream, in byte                                                                                                                                                                                                                                                                                                                                                                                                               |
| wait_quota | bigint  | Extra time generated when the stream waits the quota value, in ms                                                                                                                                                                                                                                                                                                                                                                                                      |

## 12.3.115 PG\_COMM\_STATUS

PG\_COMM\_STATUS displays the communications library status for a single DN.

**Table 12-204** PG\_COMM\_STATUS columns

| Name       | Type    | Description                                                          |
|------------|---------|----------------------------------------------------------------------|
| node_name  | text    | Node name.                                                           |
| rxpck_rate | integer | Receiving rate of the communications library on the node, in byte/s. |



| Name             | Type    | Description                                                              |
|------------------|---------|--------------------------------------------------------------------------|
| txpck_rate       | integer | Sending rate of the communications library on the node, in byte/s.       |
| rxkbyte_rate     | bigint  | Receiving rate of the communications library on the node, in kbyte/s.    |
| txkbyte_rate     | bigint  | Sending rate of the communications library on the node, in kbyte/s.      |
| buffer           | bigint  | Size of the buffer of the Cmailbox.                                      |
| memkbyte_libcomm | bigint  | Communication memory size of the <b>libcomm</b> process, in bytes.       |
| memkbyte_libpq   | bigint  | Communication memory size of the <b>libpq</b> process, in bytes.         |
| used_pm          | integer | Real-time usage of the <b>postmaster</b> thread.                         |
| used_sflow       | integer | Real-time usage of the <b>gs_sender_flow_controller</b> thread.          |
| used_rflow       | integer | Real-time usage of the <b>gs_receiver_flow_controller</b> thread.        |
| used_rloop       | integer | Highest real-time usage among multiple <b>gs_receivers_loop</b> threads. |
| stream           | integer | Total number of used logical connections.                                |

### 12.3.116 PG\_CONTROL\_GROUP\_CONFIG

PG\_CONTROL\_GROUP\_CONFIG displays Cgroup configuration information in the system. Only users with the SYSADMIN permission can query this view.

**Table 12-205** PG\_CONTROL\_GROUP\_CONFIG columns

| Name                    | Type | Description                             |
|-------------------------|------|-----------------------------------------|
| pg_control_group_config | text | Configuration information of the Cgroup |

### 12.3.117 PG\_EXT\_STATS

PG\_EXT\_STATS displays extended statistics stored in [PG\\_STATISTIC\\_EXT](#).

**Table 12-206** PG\_EXT\_STATS columns

| Name       | Type       | Reference                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------|------------|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| schemaname | name       | <b>nspname</b> in <a href="#">PG_NAMESPACE</a>    | Name of the schema that contains a table.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| tablename  | name       | <b>relname</b> in <a href="#">PG_CLASS</a>        | Table name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| attname    | int2vector | <b>stakey</b> in <a href="#">PG_STATISTIC_EXT</a> | Columns to be combined for collecting statistics                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| inherited  | boolean    | -                                                 | Table inheritance is not supported currently. The value of this column is <b>false</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| null_frac  | real       | -                                                 | Percentage of column combinations that are null to all records.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| avg_width  | integer    | -                                                 | Average width of column combinations, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| n_distinct | real       | -                                                 | <ul style="list-style-type: none"><li>Estimated number of distinct values in a column combination if the value is greater than 0.</li><li>Negative number obtained by multiplying the result calculated by dividing the number of distinct values by the number of rows by -1 if the value is less than 0. For example, <b>-1</b> indicates that the number of distinct values is the same as the number of rows for a column combination.<ol style="list-style-type: none"><li>The negated form is used when <b>ANALYZE</b> believes that the number of distinct values is likely to increase as the table grows.</li><li>The positive form is used when the column seems to have a fixed number of possible values.</li></ol></li><li>The number of distinct values is unknown if the value is <b>0</b>.</li></ul> |

| Name                  | Type     | Reference | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|----------|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n_dndistinct          | real     | -         | <p>Number of not-null distinct values in the <b>dn1</b> column combination.</p> <ul style="list-style-type: none"> <li>Exact number of distinct values if the value is greater than <b>0</b>.</li> <li>Negative number obtained by multiplying the result calculated by dividing the number of distinct values by the number of rows by -1 if the value is less than 0. For example, if a value in a column combination appears twice in average, <b>n_dndistinct</b> equals <b>-0.5</b>.</li> <li>The number of distinct values is unknown if the value is <b>0</b>.</li> </ul> |
| most_common_vals      | anyarray | -         | <p>List of the most common values in a column combination. If this combination does not have the most common values, this column will be <b>NULL</b>. None of the most common values in the column is <b>NULL</b>.</p>                                                                                                                                                                                                                                                                                                                                                           |
| most_common_freqs     | real[]   | -         | <p>List of the frequencies of the most common values in a column combination. The frequencies are obtained by dividing the number of occurrences of each value by the number of rows. If the value of <b>most_common_vals</b> is <b>NULL</b>, the value of this column is <b>NULL</b>.</p>                                                                                                                                                                                                                                                                                       |
| most_common_vals_null | anyarray | -         | <p>List of the most common values in a column combination. If this combination does not have the most common values, this column will be <b>NULL</b>. At least one of the common values in the column is <b>NULL</b>.</p>                                                                                                                                                                                                                                                                                                                                                        |

| Name                   | Type     | Reference | Description                                                                                                                                                                                                                                                                                |
|------------------------|----------|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| most_common_freqs_null | real[]   | -         | List of the frequencies of the most common values in a column combination. The frequencies are obtained by dividing the number of occurrences of each value by the number of rows. If the value of <b>most_common_vals_null</b> is <b>NULL</b> , the value of this column is <b>NULL</b> . |
| histogram_bounds       | anyarray | -         | Boundary value list of the histogram.                                                                                                                                                                                                                                                      |

### 12.3.118 PG\_GET\_INVALID\_BACKENDS

**PG\_GET\_INVALID\_BACKENDS** provides information about background threads on the primary database node that are connected to the current standby node.

**Table 12-207** PG\_GET\_INVALID\_BACKENDS columns

| Name          | Type                     | Description                                         |
|---------------|--------------------------|-----------------------------------------------------|
| pid           | bigint                   | Thread ID                                           |
| node_name     | text                     | Node information connected to the background thread |
| dbname        | name                     | Name of the connected database                      |
| backend_start | timestamp with time zone | Background thread startup time                      |
| query         | text                     | Query statement executed by the background thread   |

### 12.3.119 PG\_GET\_SENDERS\_CATCHUP\_TIME

**PG\_GET\_SENDERS\_CATCHUP\_TIME** provides catchup information of the currently active primary/standby instance sender thread on the database node.

**Table 12-208** PG\_GET\_SENDERS\_CATCHUP\_TIME columns

| Name | Type   | Description               |
|------|--------|---------------------------|
| pid  | bigint | Current sender thread ID. |

| Name          | Type                     | Description                         |
|---------------|--------------------------|-------------------------------------|
| lwpid         | integer                  | Current sender lwpid.               |
| local_role    | text                     | Local role.                         |
| peer_role     | text                     | Peer role.                          |
| state         | text                     | Current sender's replication status |
| type          | text                     | Current sender type                 |
| catchup_start | timestamp with time zone | Startup time of a catchup task.     |
| catchup_end   | timestamp with time zone | End time of the catchup task.       |

### 12.3.120 PG\_GROUP

PG\_GROUP displays the database role authentication and the relationship between group members.

**Table 12-209** PG\_GROUP columns

| Name     | Type  | Description                                         |
|----------|-------|-----------------------------------------------------|
| groname  | name  | Group name                                          |
| grosysid | oid   | Group ID                                            |
| grolist  | oid[] | Array containing the IDs of all roles in the group. |

### 12.3.121 PG\_GTT\_ATTACHED\_PIDS

PG\_GTT\_ATTACHED\_PIDS is used to check which sessions are using global temporary tables by calling the `pg_gtt_attached_pid()` function.

**Table 12-210** PG\_GTT\_ATTACHED\_PIDS columns

| Name       | Type     | Description                       |
|------------|----------|-----------------------------------|
| schemaname | name     | Schema name.                      |
| tablename  | name     | Name of a global temporary table. |
| relid      | oid      | OID of a global temporary table.  |
| pids       | bigint[] | Thread PID list.                  |

| Name       | Type     | Description      |
|------------|----------|------------------|
| sessionids | bigint[] | Session ID list. |

## 12.3.122 PG\_GTT\_RELSTATS

PG\_GTT\_RELSTATS displays basic information about all global temporary tables of the current session by calling `pg_get_gtt_relstats()`.

**Table 12-211** PG\_GTT\_RELSTATS columns

| Name          | Type    | Description                                                                                                 |
|---------------|---------|-------------------------------------------------------------------------------------------------------------|
| schemaname    | name    | Schema name.                                                                                                |
| tablename     | name    | Name of a global temporary table.                                                                           |
| relfilenode   | oid     | ID of a file object.                                                                                        |
| relpages      | integer | Number of disk pages of a global temporary table.                                                           |
| reltuples     | real    | Number of records in a global temporary table.                                                              |
| relallvisible | integer | Number of pages that are marked as all visible.                                                             |
| relfrozenxid  | xid     | All transaction IDs before this one have been replaced by a permanent (frozen) transaction ID in the table. |
| relminmxid    | xid     | Reserved.                                                                                                   |

## 12.3.123 PG\_GTT\_STATS

PG\_GTT\_STATS displays statistics about a single column in all global temporary tables of the current session by calling `pg_get_gtt_statistics()`.

**Table 12-212** PG\_GTT\_STATS columns

| Name       | Type    | Description                                                                             |
|------------|---------|-----------------------------------------------------------------------------------------|
| schemaname | name    | Schema name.                                                                            |
| tablename  | name    | Name of a global temporary table.                                                       |
| attname    | name    | Attribute name.                                                                         |
| inherited  | Boolean | Specifies whether to collect statistics for objects that have inheritance relationship. |

| Name                      | Type    | Description                                                                                       |
|---------------------------|---------|---------------------------------------------------------------------------------------------------|
| null_frac                 | real    | Percentage of column entries that are null.                                                       |
| avg_width                 | integer | Average stored width, in bytes, of non-null entries.                                              |
| n_distinct                | real    | Number of distinct, non-null data values in the column.                                           |
| most_common_vals          | text[]  | List of the most common values, which is sorted by occurrence frequency.                          |
| most_common_freqs         | real[]  | Frequencies of the most common values.                                                            |
| histogram_bounds          | text[]  | Data distribution (excluding the most common values) in a frequency histogram description column. |
| correlation               | real    | Correlation coefficient.                                                                          |
| most_common_elements      | text[]  | List of the most common element values, which is used for the array type or some other type.      |
| most_common_element_freqs | real[]  | Frequencies of the most common element values.                                                    |
| elem_count_histogram      | real[]  | Array type histogram.                                                                             |

## 12.3.124 PG\_INDEXES

PG\_INDEXES displays information about each index in the database.

**Table 12-213** PG\_INDEXES columns

| Name       | Type | Reference                              | Description                                          |
|------------|------|----------------------------------------|------------------------------------------------------|
| schemaname | name | <b>nspname</b> in <b>PG_NAMESPACE</b>  | Name of the schema that contains tables and indexes. |
| tablename  | name | <b>relname</b> in <b>PG_CLASS</b>      | Name of the table where the index is located.        |
| indexname  | name | <b>relname</b> in <b>PG_CLASS</b>      | Index name                                           |
| tablespace | name | <b>nspname</b> in <b>PG_TABLESPACE</b> | Name of the tablespace that contains the index.      |

| Name     | Type | Reference | Description                                                     |
|----------|------|-----------|-----------------------------------------------------------------|
| indexdef | text | -         | Index definition (a reconstructed <b>CREATE INDEX</b> command). |

## 12.3.125 PG\_LOCKS

PG\_LOCKS displays information about locks held by open transactions.

**Table 12-214** PG\_LOCKS columns

| Name          | Type     | Reference                                 | Description                                                                                                                                                                                                                        |
|---------------|----------|-------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| locktype      | text     | -                                         | Type of the locked object: relation, extend, page, tuple, transactionid, virtualxid, object, userlock, and advisory.                                                                                                               |
| database      | oid      | <b>oid</b> in <a href="#">PG_DATABASE</a> | OID of the database in which the locked target exists. <ul style="list-style-type: none"><li>The OID is <b>0</b> if the target is a shared object.</li><li>The OID is <b>NULL</b> if the locked object is a transaction.</li></ul> |
| relation      | oid      | <b>oid</b> in <a href="#">PG_CLASS</a>    | OID of the relationship targeted by the lock. The value is <b>NULL</b> if the object is not a relationship or part of a relationship.                                                                                              |
| page          | integer  | -                                         | Page number targeted by the lock within the relationship ( <b>NULL</b> if the object is not a relation page or row page).                                                                                                          |
| tuple         | smallint | -                                         | Row number targeted by the lock within the page ( <b>NULL</b> if the object is not a row).                                                                                                                                         |
| bucket        | integer  | -                                         | Bucket number corresponding to the child table. The value is <b>NULL</b> if the target is not a table.                                                                                                                             |
| virtualxid    | text     | -                                         | ID of the virtual transaction. The value is <b>NULL</b> if the object is not a virtual transaction.                                                                                                                                |
| transactionid | xid      | -                                         | ID of the transaction targeted by the lock. The value is <b>NULL</b> if the object is not a transaction ID.                                                                                                                        |



| Name                | Type     | Reference                     | Description                                                                                                                                                               |
|---------------------|----------|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| classid             | oid      | <b>oid</b> in <b>PG_CLASS</b> | OID of the system catalog that contains the object ( <b>NULL</b> if the object is not a general database object).                                                         |
| objid               | oid      | -                             | OID of the locked object within its system catalog. The value is <b>NULL</b> if the object is not a general database object.                                              |
| objsubid            | smallint | -                             | Column number for a column in the table ( <b>0</b> if the object is of other object type and <b>NULL</b> if the object is not a general database object).                 |
| virtualtrans action | text     | -                             | ID of the virtual transaction holding or awaiting this lock.                                                                                                              |
| pid                 | bigint   | -                             | ID of the server logic thread that holds or waits for the lock. ( <b>NULL</b> if the lock is held by a prepared transaction)                                              |
| sessionid           | bigint   | -                             | ID of the session that holds or waits for the lock.                                                                                                                       |
| mode                | text     | -                             | Lock mode held or desired by this thread.                                                                                                                                 |
| granted             | boolean  | -                             | <ul style="list-style-type: none"> <li>The value is <b>TRUE</b> if the lock is a held lock.</li> <li>The value is <b>FALSE</b> if the lock is an awaited lock.</li> </ul> |
| fastpath            | boolean  | -                             | The value is <b>TRUE</b> if the lock is obtained through <b>fast-path</b> , and is <b>FALSE</b> if the lock is obtained through the main lock table.                      |
| locktag             | text     | -                             | Lock information that the session waits for. It can be parsed using the locktag_decode() function.                                                                        |
| global_sessionid    | text     | -                             | Global session ID.                                                                                                                                                        |

### 12.3.126 PG\_NODE\_ENV

PG\_NODE\_ENV displays the environment variable information of the current node.

**Table 12-215** PG\_NODE\_ENV columns

| Name          | Type    | Description                         |
|---------------|---------|-------------------------------------|
| node_name     | text    | Current node name.                  |
| host          | text    | Host name of the node.              |
| process       | integer | Number of the node process.         |
| port          | integer | Port ID of the node.                |
| installpath   | text    | Installation directory of the node. |
| datapath      | text    | Data directory of the node.         |
| log_directory | text    | Log directory of the node.          |

## 12.3.127 PG\_OS\_THREADS

**PG\_OS\_THREADS** provides status information about all the threads under the current node.

**Table 12-216** PG\_OS\_THREADS columns

| Name          | Type                     | Description                                              |
|---------------|--------------------------|----------------------------------------------------------|
| node_name     | text                     | Current node name                                        |
| pid           | bigint                   | PID of the thread running under the current node process |
| lwpid         | integer                  | Lightweight thread ID corresponding to the PID           |
| thread_name   | text                     | Thread name corresponding to the PID                     |
| creation_time | timestamp with time zone | Thread creation time corresponding to the PID            |

## 12.3.128 PG\_PREPARED\_STATEMENTS

**PG\_PREPARED\_STATEMENTS** displays information about all prepared statements that are available in the current session.

**Table 12-217** PG\_PREPARED\_STATEMENTS columns

| Name | Type | Description                           |
|------|------|---------------------------------------|
| name | text | Identifier of the prepared statement. |

| Name            | Type                     | Description                                                                                                                                                                                                                                                                      |
|-----------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| statement       | text                     | Query string for creating this prepared statement. For prepared statements created through SQL, this is the PREPARE statement committed by the client. For prepared statements created through the frontend/backend protocol, this is the text of the prepared statement itself. |
| prepare_time    | timestamp with time zone | Timestamp when the prepared statement is created.                                                                                                                                                                                                                                |
| parameter_types | regtype[]                | Expected parameter types for the prepared statement in the form of an array of <b>regtype</b> . The OID corresponding to an element of this array can be obtained by converting the <b>regtype</b> value to an OID value.                                                        |
| from_sql        | boolean                  | <ul style="list-style-type: none"><li>• <b>True</b> if the prepared statement was created through the PREPARE statement</li><li>• <b>False</b> if the statement was prepared through the frontend/backend protocol</li></ul>                                                     |

## 12.3.129 PG\_PREPARED\_XACTS

PG\_PREPARED\_XACTS displays information about transactions that are currently prepared for two-phase commit.

**Table 12-218** PG\_PREPARED\_XACTS columns

| Name        | Type                     | Reference                            | Description                                          |
|-------------|--------------------------|--------------------------------------|------------------------------------------------------|
| transaction | xid                      | -                                    | Numeric identifier of the prepared transaction.      |
| gid         | text                     | -                                    | Global identifier of the prepared transaction.       |
| prepared    | timestamp with time zone | -                                    | Time at which the transaction is prepared for commit |
| owner       | name                     | <b>rolname</b> in <b>PG_AUTHID</b>   | Name of the user who executes the transaction.       |
| database    | name                     | <b>datname</b> in <b>PG_DATABASE</b> | Name of the database that executes the transaction.  |

## 12.3.130 PG\_REPLICATION\_ORIGIN\_STATUS

PG\_REPLICATION\_ORIGIN\_STATUS displays the replication status of the replication source.

**Table 12-219** PG\_REPLICATION\_ORIGIN\_STATUS columns

| Name        | Type | Description                           |
|-------------|------|---------------------------------------|
| local_id    | oid  | Replication source ID.                |
| external_id | text | Name of the replication source.       |
| remote_lsn  | text | Remote LSN of the replication source. |
| local_lsn   | text | Local LSN of the replication source.  |

## 12.3.131 PG\_REPLICATION\_SLOTS

PG\_REPLICATION\_SLOTS displays replication slot information.

**Table 12-220** PG\_REPLICATION\_SLOTS columns

| Name      | Type    | Description                                                                                                                                                              |
|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| slot_name | text    | Replication slot name.                                                                                                                                                   |
| plugin    | text    | Name of the output plug-in corresponding to the logical replication slot.                                                                                                |
| slot_type | text    | Replication slot type. <ul style="list-style-type: none"><li>• <b>physical</b>: physical replication slot.</li><li>• <b>logical</b>: logical replication slot.</li></ul> |
| datoid    | oid     | OID of the database where the replication slot resides.                                                                                                                  |
| database  | name    | Name of the database where the replication slot resides.                                                                                                                 |
| active    | Boolean | Determines whether the replication slot is activated. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>            |
| xmin      | xid     | XID of the earliest transaction that the database must reserve for the replication slot.                                                                                 |

| Name            | Type    | Description                                                                                                              |
|-----------------|---------|--------------------------------------------------------------------------------------------------------------------------|
| catalog_xmin    | xid     | XID of the earliest system catalog-involved transaction that the database must reserve for the logical replication slot. |
| restart_lsn     | text    | Physical location of the earliest Xlog required by the replication slot.                                                 |
| dummy_standby   | Boolean | Not supported in the current version.                                                                                    |
| confirmed_flush | text    | Dedicated logical replication slot. The client confirms the location of the received log.                                |

## 12.3.132 PG\_RLSPOLICIES

PG\_RLSPOLICIES displays information about row-level security policies. The initial user and users with the SYSADMIN attribute can view all policy information. Other users can view only the policy information in their own tables.

**Table 12-221** PG\_RLSPOLICIES columns

| Name             | Type   | Description                                                                                                                                                                                                                                                                                                         |
|------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| schemaname       | name   | Name of the schema of the table object to which a row-level security policy is applied.                                                                                                                                                                                                                             |
| tablename        | name   | Name of the table object to which the row-level security policy is applied.                                                                                                                                                                                                                                         |
| polycyname       | name   | Name of a row-level security policy.                                                                                                                                                                                                                                                                                |
| policypermissive | text   | Expression concatenation mode of a row-level security policy. Value range: <ul style="list-style-type: none"><li>● <b>PERMISSIVE</b>: permissive policy, which is concatenated using the OR expression.</li><li>● <b>RESTRICTIVE</b>: restrictive policy, which is concatenated using the AND expression.</li></ul> |
| policyroles      | name[] | List of users affected by the row-level security policy. If this parameter is not specified, all users are affected.                                                                                                                                                                                                |
| polycycmd        | text   | SQL operations affected by a row-level security policy.                                                                                                                                                                                                                                                             |
| policyqual       | text   | Expression of a row-level security policy.                                                                                                                                                                                                                                                                          |

## 12.3.133 PG\_ROLES

PG\_ROLES displays information about database roles. The initial user and users with the SYSADMIN or CREATEROLE attribute can view information about all roles. Other users can view only their own information.

**Table 12-222** PG\_ROLES columns

| Name          | Type    | Reference | Description                                                                                                                                                                                                                                                                                                      |
|---------------|---------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rolname       | name    | N/A       | Role name.                                                                                                                                                                                                                                                                                                       |
| rolsuper      | boolean | N/A       | Specifies whether the role is the initial SYSADMIN with the highest permissions. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                         |
| rolinherit    | boolean | N/A       | Specifies whether the role inherits the permissions for this type of roles. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                              |
| rolcreatorole | boolean | N/A       | Specifies whether the role can create other roles. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                       |
| rolcreatedb   | boolean | N/A       | Specifies whether the role can create databases. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                         |
| rolcatupdate  | boolean | N/A       | Specifies whether the role can update system catalogs directly. Only the initial SYSADMIN whose <b>usesysid</b> is set to <b>10</b> has this permission. This permission is unavailable for other users. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| rolcanlogin   | boolean | N/A       | Specifies whether the role can log in to the database. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                                                                                                   |

| Name           | Type                     | Reference                                     | Description                                                                                                                                                 |
|----------------|--------------------------|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rolreplication | boolean                  | N/A                                           | Specifies whether the role can be replicated.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul> |
| rolauditadmin  | boolean                  | N/A                                           | Specifies whether the role is an AUDITADMIN.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>  |
| rolsystemadmin | boolean                  | N/A                                           | Specifies whether the role is a SYSADMIN.<br><ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>     |
| rolconnlimit   | integer                  | N/A                                           | Sets the maximum number of concurrent connections that this role can initiate if this role can log in. The value <b>-1</b> indicates no limit.              |
| rolpassword    | text                     | N/A                                           | Encrypted user password. The value is displayed as <b>*****</b> .                                                                                           |
| rolvalidbegin  | timestamp with time zone | N/A                                           | Start time for account validity ( <b>NULL</b> if start time is not specified).                                                                              |
| rolvaliduntil  | timestamp with time zone | N/A                                           | End time for account validity ( <b>NULL</b> if end time is not specified).                                                                                  |
| rolparentid    | oid                      | <b>rolparentid</b> in <b>PG_AUTHID</b>        | OID of a group user to which the user belongs.                                                                                                              |
| roltabspace    | text                     | N/A                                           | Storage space of the user permanent table, in KB.                                                                                                           |
| rolconfig      | text[]                   | <b>setconfig</b> in <b>PG_DB_ROLE_SETTING</b> | Default value of runtime configuration items.                                                                                                               |
| oid            | oid                      | <b>oid</b> in <b>PG_AUTHID</b>                | Role ID                                                                                                                                                     |

| Name             | Type    | Reference                           | Description                                                                                                                                                                  |
|------------------|---------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| roluseft         | boolean | <b>roluseft</b> in <b>PG_AUTHID</b> | Specifies whether the role can perform operations on foreign tables. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| rolkind          | "char"  | N/A                                 | Role type <ul style="list-style-type: none"><li>• <b>n</b>: common user, that is, non-permanent user.</li><li>• <b>p</b>: permanent user.</li></ul>                          |
| nodegroup        | name    | N/A                                 | Unsupported currently.                                                                                                                                                       |
| roltemp space    | text    | N/A                                 | Storage space of the user temporary table, in KB.                                                                                                                            |
| rolspill space   | text    | N/A                                 | Operator disk flushing space of the user, in KB.                                                                                                                             |
| rolmonitoradmin  | boolean | N/A                                 | Specifies whether the role is a MONADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                            |
| roloperatoradmin | boolean | N/A                                 | Specifies whether the role is an OPRADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                           |
| rolpolicyadmin   | boolean | N/A                                 | Specifies whether the role is a POLADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                            |

## 12.3.134 PG\_RULES

PG\_RULES is used to query useful information about rewriting rules.

**Table 12-223** PG\_RULES columns

| Name       | Type | Description                                  |
|------------|------|----------------------------------------------|
| schemaname | name | Name of the schema that contains a table.    |
| tablename  | name | Name of the table to which the rule applies. |



| Name       | Type | Description                                              |
|------------|------|----------------------------------------------------------|
| rulename   | name | Name of a rule.                                          |
| definition | text | Rule definition (reconstructed by the CREATE statement). |

### 12.3.135 PG\_RUNNING\_XACTS

PG\_RUNNING\_XACTS displays the running transaction information on the current node.

**Table 12-224** PG\_RUNNING\_XACTS columns

| Name        | Type    | Description                                                                                                                                                       |
|-------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| handle      | integer | Slot handle in the transaction manager corresponding to a transaction. The value is fixed at <b>-1</b> .                                                          |
| gxid        | xid     | Transaction ID.                                                                                                                                                   |
| state       | tinyint | Transaction status. <ul style="list-style-type: none"><li>• <b>3</b>: prepared.</li><li>• <b>0</b>: starting.</li></ul>                                           |
| node        | text    | Node name.                                                                                                                                                        |
| xmin        | xid     | Minimum transaction ID on the node.                                                                                                                               |
| vacuum      | Boolean | Specifies whether the current transaction is lazy vacuum. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| timeline    | bigint  | Number of database restarts.                                                                                                                                      |
| prepare_xid | xid     | ID of the transaction in the prepared state; otherwise, the value is <b>0</b> .                                                                                   |
| pid         | bigint  | Thread ID corresponding to the transaction.                                                                                                                       |
| next_xid    | xid     | Transaction ID sent by other nodes to the current node. The value is fixed at <b>0</b> .                                                                          |

### 12.3.136 PG\_SECLABELS

PG\_SECLABELS provides information about security labels.

**Table 12-225** PG\_SECLABELS columns

| Name         | Type    | Reference                                      | Description                                                                                                                                                                               |
|--------------|---------|------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| objoid       | oid     | Any OID column                                 | OID of the object that this security label pertains to.                                                                                                                                   |
| classoid     | oid     | <b>oid</b> in <a href="#">PG_CLASS</a>         | OID of the system catalog where the object to which the security label belongs is located.                                                                                                |
| objsubid     | integer | -                                              | Column number for the security label on a table column ( <b>objoid</b> and <b>classoid</b> refer to the table itself). The value is <b>0</b> for all other object types.                  |
| objtype      | text    | -                                              | Type of the object to which the label belongs, in text format. Example: <ul style="list-style-type: none"><li>• <b>table</b>: table type.</li><li>• <b>column</b>: column type.</li></ul> |
| objnamespace | oid     | <b>oid</b> in <a href="#">PG_NAMESPACE</a>     | OID of the namespace for the object, if applicable; otherwise, <b>NULL</b> .                                                                                                              |
| objname      | text    | -                                              | Name of the object to which the label belongs, in text format.                                                                                                                            |
| provider     | text    | <b>provider</b> in <a href="#">PG_SECLABEL</a> | Provider of the label.                                                                                                                                                                    |
| label        | text    | <b>label</b> in <a href="#">PG_SECLABEL</a>    | Security label name.                                                                                                                                                                      |

## 12.3.137 PG\_SETTINGS

PG\_SETTINGS displays information about parameters of the running database.

**Table 12-226** PG\_SETTINGS columns

| Name     | Type | Description                     |
|----------|------|---------------------------------|
| name     | text | Parameter name.                 |
| setting  | text | Current parameter value.        |
| unit     | text | Unit of the parameter.          |
| category | text | Logical group of the parameter. |

| Name       | Type    | Description                                                                                                                                                                       |
|------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| short_desc | text    | Brief description of the parameter.                                                                                                                                               |
| extra_desc | text    | Detailed description of the parameter.                                                                                                                                            |
| context    | text    | Context required to set the parameter value, including <b>internal</b> , <b>postmaster</b> , <b>sighup</b> , <b>backend</b> , <b>superuser</b> , and <b>user</b> .                |
| vartype    | text    | Parameter type, including <b>bool</b> , <b>enum</b> , <b>integer</b> , <b>real</b> , or <b>string</b> .                                                                           |
| source     | text    | Method of assigning the parameter value.                                                                                                                                          |
| min_val    | text    | Minimum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                                   |
| max_val    | text    | Maximum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                                   |
| enumvals   | text[]  | Valid values of an enum-type parameter. If the parameter type is not enum, the value of this column is <b>null</b> .                                                              |
| boot_val   | text    | Default parameter value used upon the database startup.                                                                                                                           |
| reset_val  | text    | Default parameter value used upon the database reset.                                                                                                                             |
| sourcefile | text    | Configuration file used to set parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> .                   |
| sourceline | integer | Row number of the configuration file for setting parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> . |

### 12.3.138 PG\_SHADOW

PG\_SHADOW displays the attributes of all roles marked with rolcanlogin in **PG\_AUTHID**. Only users with the SYSADMIN permission can access this system view.

The information in this view is basically the same as that of the view in **PG\_USER**. The difference is that passwords are sensitive and displayed as **\*\*\*\*\*** in the latter.

**Table 12-227** PG\_SHADOW columns

| Name        | Type                     | Reference                              | Description                                                                                                                                                                                                                     |
|-------------|--------------------------|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| username    | name                     | <b>rolname</b> in <b>PG_AUTHID</b>     | Username.                                                                                                                                                                                                                       |
| usesysid    | oid                      | <b>oid</b> in <b>PG_AUTHID</b>         | ID of this user.                                                                                                                                                                                                                |
| usecreatedb | boolean                  | -                                      | Specifies whether a user has the permissions to create databases. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                       |
| usesuper    | boolean                  | -                                      | Specifies whether the user is a SYSADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                                                               |
| usecatupd   | boolean                  | -                                      | Specifies whether the user can update a view. Even the SYSADMIN cannot update the view unless this column is <b>t</b> . <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| userepl     | boolean                  | -                                      | Specifies whether the user has the permissions to duplicate data streams. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                                               |
| passwd      | text                     | <b>rolpassword</b> in <b>PG_AUTHID</b> | Password ciphertext. <b>NULL</b> if no password.                                                                                                                                                                                |
| valbegin    | timestamp with time zone | -                                      | Start time for account validity ( <b>NULL</b> if start time is not specified).                                                                                                                                                  |
| valuntil    | timestamp with time zone | -                                      | End time for account validity ( <b>NULL</b> if end time is not specified).                                                                                                                                                      |
| respool     | name                     | -                                      | Resource pool where the user is in.                                                                                                                                                                                             |
| parent      | oid                      | -                                      | Parent user OID.                                                                                                                                                                                                                |
| spacelimit  | text                     | -                                      | Storage space of the permanent table, in KB.                                                                                                                                                                                    |

| Name                 | Type    | Reference                                              | Description                                                                                                                                        |
|----------------------|---------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| useconfig            | text[]  | <b>setconfig</b> in <a href="#">PG_DB_ROLE_SETTING</a> | Default value of runtime configuration items.                                                                                                      |
| tempspaceli<br>mit   | text    | -                                                      | Storage space of the temporary table, in KB.                                                                                                       |
| spillspaceli<br>mit  | text    | -                                                      | Operator disk flushing space, in KB.                                                                                                               |
| usemonitor<br>admin  | Boolean | -                                                      | Specifies whether this user is a MONADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| useoperator<br>admin | Boolean | -                                                      | Specifies whether the user is an OPRADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| usepolicyad<br>min   | Boolean | -                                                      | Specifies whether the user is a POLADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>  |

## 12.3.139 PG\_STATS

PG\_STATS displays single-column statistics stored in the pg\_statistic table. The **autovacuum\_naptime** parameter specifies the interval for updating statistics recorded in the view.

**Table 12-228** PG\_STATS columns

| Name       | Type    | Reference                                      | Description                                                                              |
|------------|---------|------------------------------------------------|------------------------------------------------------------------------------------------|
| schemaname | name    | <b>nspname</b> in <a href="#">PG_NAMESPACE</a> | Name of the schema that the table is in.                                                 |
| tablename  | name    | <b>relname</b> in <a href="#">PG_CLASS</a>     | Table name.                                                                              |
| attname    | name    | <b>attname</b> in <a href="#">PG_ATTRIBUTE</a> | Column name.                                                                             |
| inherited  | boolean | -                                              | Table inheritance is not supported currently. The value of this column is <b>false</b> . |

| Name         | Type    | Reference | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------|---------|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| null_frac    | real    | -         | Percentage of column entries that are null.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| avg_width    | integer | -         | Average width in bytes of column's entries.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| n_distinct   | real    | -         | <ul style="list-style-type: none"> <li>Estimated number of distinct values in the column if the value is greater than <b>0</b>.</li> <li>Negative number obtained by multiplying the result calculated by dividing the number of distinct values by the number of rows by <b>-1</b> if the value is less than <b>0</b>. For example, <b>-1</b> indicates that the number of distinct values is the same as the number of rows for a column combination. <ol style="list-style-type: none"> <li>The negated form is used when ANALYZE believes that the number of distinct values is likely to increase as the table grows.</li> <li>The positive form is used when the column seems to have a fixed number of possible values.</li> </ol> </li> <li>The number of distinct values is unknown if the value is <b>0</b>.</li> </ul> |
| n_dndistinct | real    | -         | <p>Number of not-null distinct values in the <b>dn1</b> column.</p> <ul style="list-style-type: none"> <li>Exact number of distinct values if the value is greater than <b>0</b>.</li> <li>Negative number obtained by multiplying the result calculated by dividing the number of distinct values by the number of rows by <b>-1</b> if the value is less than <b>0</b>. For example, if the value of a column appears twice in average, set <b>n_dndistinct</b> to <b>-0.5</b>.</li> <li>The number of distinct values is unknown if the value is <b>0</b>.</li> </ul>                                                                                                                                                                                                                                                          |

| Name                   | Type     | Reference | Description                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------|----------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| most_common_vals       | anyarray | -         | List of the most common values in a column. If this column does not have the most common value, the value is <b>NULL</b> .                                                                                                                                                                                                                                                   |
| most_common_freqs      | real[]   | -         | List of the frequencies of the most common values in a column. The frequencies are obtained by dividing the number of occurrences of each value by the number of rows. The value is <b>NULL</b> if the value of <b>most_common_vals</b> is <b>NULL</b> .                                                                                                                     |
| histogram_bounds       | anyarray | -         | Frequency histogram consisting of values excluding null values and MCVs. If a value appears in the value of <b>most_common_vals</b> , it does not appear in the histogram. If the column data type does not have the < operator or the list specified by <b>most_common_vals</b> contains all values of the column, the histogram information of the column is <b>NULL</b> . |
| correlation            | real     | -         | Correlation between the physical row sequence and logical row sequence of a column value. The value ranges from -1 to +1. When the value is close to -1 or +1, the index scan overhead is less than that when the value is close to 0 because random access to the disk is reduced. This column is <b>NULL</b> if the column data type does not have a < operator.           |
| most_common_elems      | anyarray | -         | A list of non-null element values most often appearing.                                                                                                                                                                                                                                                                                                                      |
| most_common_elem_freqs | real[]   | -         | List that records the frequency of the most commonly used non-null elements.                                                                                                                                                                                                                                                                                                 |
| elem_count_histogram   | real[]   | -         | A histogram of the counts of distinct non-null element values.                                                                                                                                                                                                                                                                                                               |

## 12.3.140 PG\_STAT\_ACTIVITY

PG\_STAT\_ACTIVITY shows information about the current user's queries. The columns save information about the last query.

**Table 12-229** PG\_STAT\_ACTIVITY columns

| Name             | Type                     | Description                                                                                                                                                                                                                           |
|------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid            | oid                      | OID of the database that the user session connects to in the backend.                                                                                                                                                                 |
| datname          | name                     | Name of the database that the user session connects to in the backend.                                                                                                                                                                |
| pid              | bigint                   | Backend thread ID.                                                                                                                                                                                                                    |
| sessionid        | bigint                   | Session ID.                                                                                                                                                                                                                           |
| usesysid         | oid                      | OID of the user logged in to the backend.                                                                                                                                                                                             |
| username         | name                     | Name of the user logged in to the backend.                                                                                                                                                                                            |
| application_name | text                     | Name of the application connected to the backend.                                                                                                                                                                                     |
| client_addr      | inet                     | IP address of the client connected to the backend. If this column is <b>null</b> , it indicates either the client is connected through a Unix socket on the server or this is an internal process, such as AUTOVACUUM.                |
| client_hostname  | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                                     |
| client_port      | integer                  | TCP port number that the client uses for communication with this backend ( <b>-1</b> if a Unix socket is used).                                                                                                                       |
| backend_start    | timestamp with time zone | Time when this session was started, that is, when the client connected to the server.                                                                                                                                                 |
| xact_start       | timestamp with time zone | Time when the current active transaction was started ( <b>null</b> if no transaction is active). If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column. |



| Name         | Type                     | Description                                                                                                                                                                                                                                                                                             |
|--------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| query_start  | timestamp with time zone | Time when the current active query was started, or time when the last query was started if the value of <b>state</b> is not <b>active</b> . For a stored procedure, function, or package, the first query time is displayed and does not change with the running of statements in the stored procedure. |
| state_change | timestamp with time zone | Time when <b>state</b> was last modified.                                                                                                                                                                                                                                                               |
| waiting      | Boolean                  | Specifies whether the backend is currently waiting for a lock. If the backend is currently waiting for a lock, the value is <b>true</b> .                                                                                                                                                               |
| enqueue      | text                     | Unsupported currently.                                                                                                                                                                                                                                                                                  |

| Name          | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state         | text   | <p>Overall status of the backend. The value must be one of the following:</p> <ul style="list-style-type: none"><li>• <b>active</b>: The backend is executing a query.</li><li>• <b>idle</b>: The backend is waiting for a new client command.</li><li>• <b>idle in transaction</b>: The backend is in a transaction, but there is no statement being executed in the transaction.</li><li>• <b>idle in transaction (aborted)</b>: The backend is in a transaction, but there are statements failed in the transaction.</li><li>• <b>fastpath function call</b>: The backend is executing a fast-path function.</li><li>• <b>disabled</b>: This state is reported if <b>track_activities</b> is disabled in this backend.</li></ul> <p><b>NOTE</b><br/>Common users can view their own session status only. The state information of other accounts is empty. For example, after user <b>judy</b> is connected to the database, the state information of user <b>joe</b> and the initial user omm in <b>pg_stat_activity</b> is empty.</p> <pre>SELECT datname, username, usesysid, state,pid FROM pg_stat_activity;  datname   username   usesysid   state    pid -----+-----+-----+-----+-----  testdb   omm             10          139968752121616  testdb   omm             10          139968903116560  db_tpcc   judy       16398   active   139968391403280  testdb   omm             10          139968643069712  testdb   omm             10          139968680818448  testdb   joe        16390          139968563377936 (6 rows)</pre> |
| resource_pool | name   | Resource pool used by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| query_id      | bigint | Query ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

| Name            | Type   | Description                                                                                                                                                                       |
|-----------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| query           | text   | Latest query at the backend. If the value of <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed. |
| connection_info | text   | A string in JSON format recording the driver type, driver version, driver deployment path, and process owner of the connected database.                                           |
| unique_sql_id   | bigint | Unique SQL statement ID.                                                                                                                                                          |
| trace_id        | text   | Driver-specific trace ID, which is associated with an application request.                                                                                                        |

## 12.3.141 PG\_STAT\_ACTIVITY\_NG

PG\_STAT\_ACTIVITY\_NG displays information about all queries in the logical database instance of the current user.

**Table 12-230** PG\_STAT\_ACTIVITY\_NG columns

| Name             | Type   | Description                                                                                                                                                                                                            |
|------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid            | oid    | OID of the database that the user session connects to in the backend.                                                                                                                                                  |
| datname          | name   | Name of the database that the user session connects to in the backend.                                                                                                                                                 |
| pid              | bigint | Backend thread ID.                                                                                                                                                                                                     |
| sessionid        | bigint | Session ID.                                                                                                                                                                                                            |
| usesysid         | oid    | OID of the user logged in to the backend.                                                                                                                                                                              |
| username         | name   | Name of the user logged in to the backend.                                                                                                                                                                             |
| application_name | text   | Name of the application connected to the backend.                                                                                                                                                                      |
| client_addr      | inet   | IP address of the client connected to the backend. If this column is null, it indicates either that the client is connected via a Unix socket on the server machine or this is an internal process such as autovacuum. |

| Name            | Type                     | Description                                                                                                                                                                                                                                                                                |
|-----------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| client_hostname | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                                                                                          |
| client_port     | integer                  | TCP port number that the client uses for communication with the backend (-1 if a Unix socket is used)                                                                                                                                                                                      |
| backend_start   | timestamp with time zone | Time when this process was started, that is, when the client connected to the server                                                                                                                                                                                                       |
| xact_start      | timestamp with time zone | Time when current transaction was started (null if no transaction is active). If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column.                                                                         |
| query_start     | timestamp with time zone | Time when the currently active query was started, or time when the last query was started if <b>state</b> is not <b>active</b> . For a stored procedure, function, or package, the first query time is queried and does not change with the running of statements in the stored procedure. |
| state_change    | timestamp with time zone | Time when <b>state</b> was last modified.                                                                                                                                                                                                                                                  |
| waiting         | Boolean                  | Specifies whether the backend is currently waiting for a lock. If the backend is currently waiting for a lock, the value is <b>true</b> . Otherwise, the value is <b>false</b> .                                                                                                           |
| enqueue         | text                     | Queuing status of a statement. Possible values are: <ul style="list-style-type: none"> <li><b>waiting in queue</b>: The statement is in the queue.</li> <li><b>Empty</b>: The statement is running.</li> </ul>                                                                             |

| Name          | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state         | text   | <p>Overall status of the backend. Its value can be:</p> <ul style="list-style-type: none"> <li>• <b>active</b>: The backend is executing a query.</li> <li>• <b>idle</b>: The backend is waiting for a new client command.</li> <li>• <b>idle in transaction</b>: The backend is in a transaction, but there is no statement being executed in the transaction.</li> <li>• <b>idle in transaction (aborted)</b>: The backend is in a transaction, but there are statements failed in the transaction.</li> <li>• <b>fastpath function call</b>: The backend is executing a fast-path function.</li> <li>• <b>disabled</b>: This state is reported if <b>track_activities</b> is disabled in this backend.</li> </ul> <p><b>NOTE</b><br/>Common users can view their own session status only. The state information of other accounts is empty. For example, after user <b>judy</b> is connected to the database, the state information of user <b>joe</b> and the initial user <b>omm</b> in <b>pg_stat_activity</b> is empty.</p> <pre>SELECT datname, username, usesysid, state,pid FROM pg_stat_activity_ng;  datname   username   usesysid   state    pid -----+-----+-----+-----+-----  testdb   omm             10          139968752121616  testdb   omm             10          139968903116560  db_tpcds   judy          16398   active   139968391403280  testdb   omm             10          139968643069712  testdb   omm             10          139968680818448  testdb   joe           16390          139968563377936 (6 rows)</pre> |
| resource_pool | name   | Resource pool used by the user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| query_id      | bigint | Query ID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| Name       | Type | Description                                                                                                                                                                       |
|------------|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| query      | text | Latest query at the backend. If the value of <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed. |
| node_group | text | Logical database instance of the user to which the statement belongs.                                                                                                             |

### 12.3.142 PG\_STAT\_ALL\_INDEXES

PG\_STAT\_ALL\_INDEXES is used to query statistics about accesses to a specific index by querying each index row in the current database.

Indexes can be used via either simple index scans or bitmap index scans. In a bitmap scan the output of several indexes can be combined via AND or OR rules, so it is difficult to associate individual heap row fetches with specific indexes when a bitmap scan is used. Therefore, each bitmap scan increments the **pg\_stat\_all\_indexes.idx\_tup\_read** count(s) for the index(es) it uses, and it increments the **pg\_stat\_all\_tables.idx\_tup\_fetch** count for the table, but it does not affect **pg\_stat\_all\_indexes.idx\_tup\_fetch**.

**Table 12-231** PG\_STAT\_ALL\_INDEXES columns

| Name          | Type   | Description                                                              |
|---------------|--------|--------------------------------------------------------------------------|
| relid         | oid    | OID of the table that the index is created for.                          |
| indexrelid    | oid    | OID of the index.                                                        |
| schemaname    | name   | Name of the schema that the index is in.                                 |
| relname       | name   | Name of the table for the index.                                         |
| indexrelname  | name   | Index name.                                                              |
| idx_scan      | bigint | Number of index scans initiated on the index.                            |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index. |

## 12.3.143 PG\_STAT\_ALL\_TABLES

PG\_STAT\_ALL\_TABLES is used to query one row for each table in the current database (including TOAST tables), showing statistics about accesses to that specific table.

**Table 12-232** PG\_STAT\_ALL\_TABLES columns

| Name             | Type                     | Description                                                          |
|------------------|--------------------------|----------------------------------------------------------------------|
| relid            | oid                      | OID of a table.                                                      |
| schemaname       | name                     | Name of the schema that the table is in.                             |
| relname          | name                     | Table name.                                                          |
| seq_scan         | bigint                   | Number of sequential scans initiated on the table.                   |
| seq_tup_read     | bigint                   | Number of live rows fetched by sequential scans.                     |
| idx_scan         | bigint                   | Number of index scans initiated on the table.                        |
| idx_tup_fetch    | bigint                   | Number of live rows fetched by index scans.                          |
| n_tup_ins        | bigint                   | Number of rows inserted.                                             |
| n_tup_upd        | bigint                   | Number of rows updated.                                              |
| n_tup_del        | bigint                   | Number of rows deleted.                                              |
| n_tup_hot_upd    | bigint                   | Number of rows HOT updated (with no separate index update required). |
| n_live_tup       | bigint                   | Estimated number of live rows.                                       |
| n_dead_tup       | bigint                   | Estimated number of inactive rows.                                   |
| last_vacuum      | timestamp with time zone | Last time when the table was cleared.                                |
| last_autovacuum  | timestamp with time zone | Last time when the table was vacuumed by the AUTOVACUUM daemon.      |
| last_analyze     | timestamp with time zone | Last time when the table was analyzed.                               |
| last_autoanalyze | timestamp with time zone | Last time when the table was analyzed by the AUTOVACUUM daemon.      |

| Name              | Type                     | Description                                                                                                                                                              |
|-------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| vacuum_count      | bigint                   | Number of times the table is cleared.                                                                                                                                    |
| autovacuum_count  | bigint                   | Number of times the table has been vacuumed by the autovacuum daemon.                                                                                                    |
| analyze_count     | bigint                   | Number of times the table is analyzed.                                                                                                                                   |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon.                                                                                                    |
| last_data_changed | timestamp with time zone | Last time at which the table was updated (by INSERT/UPDATE/DELETE or EXCHANGE/TRUNCATE/DROP PARTITION). This column is recorded only on the local primary database node. |

### 12.3.144 PG\_STAT\_BAD\_BLOCK

PG\_STAT\_BAD\_BLOCK shows statistics about Page or CU verification failures after a node is started.

**Table 12-233** PG\_STAT\_BAD\_BLOCK columns

| Name         | Type     | Description                                                                                                                                                                    |
|--------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| nodename     | text     | Node name.                                                                                                                                                                     |
| databaseid   | integer  | OID of a database.                                                                                                                                                             |
| tablespaceid | integer  | Tablespace OID.                                                                                                                                                                |
| relfilenode  | integer  | File object ID.                                                                                                                                                                |
| bucketid     | smallint | ID of the bucket for consistent hashing.                                                                                                                                       |
| forknum      | integer  | File type. The values are as follows: <ul style="list-style-type: none"><li>• 0: main data file.</li><li>• 1: FSM file.</li><li>• 2: VM file.</li><li>• 3: BCM file.</li></ul> |
| error_count  | integer  | Number of verification failures.                                                                                                                                               |



| Name       | Type                     | Description                              |
|------------|--------------------------|------------------------------------------|
| first_time | timestamp with time zone | Time of the first verification failure.  |
| last_time  | timestamp with time zone | Time of the latest verification failure. |

## 12.3.145 PG\_STAT\_BGWRITER

PG\_STAT\_BGWRITER displays statistics about the activity of the background writer process.

**Table 12-234** PG\_STAT\_BGWRITER columns

| Name                  | Type             | Description                                                                                                |
|-----------------------|------------------|------------------------------------------------------------------------------------------------------------|
| checkpoints_timed     | bigint           | Number of scheduled checkpoints that have been performed.                                                  |
| checkpoints_req       | bigint           | Number of checkpoints that have been performed actively.                                                   |
| checkpoint_write_time | double precision | Total time spent in the checkpoint processing portion when writing files to disk, in milliseconds.         |
| checkpoint_sync_time  | double precision | Total time spent in the checkpoint processing portion when synchronizing files to disk, in milliseconds.   |
| buffers_checkpoint    | bigint           | Number of buffers written by checkpoints.                                                                  |
| buffers_clean         | bigint           | Number of buffers written by the background writer process.                                                |
| maxwritten_clean      | bigint           | Number of times that cleanup scanning stops because the background writer process writes too many buffers. |
| buffers_backend       | bigint           | Number of buffers written directly by the backend.                                                         |

| Name                  | Type                     | Description                                                                                                                                                   |
|-----------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| buffers_backend_fsync | bigint                   | Number of times that the backend calls fsync (usually, even if the backend executes these write actions, the background writer process processes them again). |
| buffers_alloc         | bigint                   | Number of buffers allocated.                                                                                                                                  |
| stats_reset           | timestamp with time zone | Time when these statistics were last reset.                                                                                                                   |

## 12.3.146 PG\_STAT\_DATABASE

PG\_STAT\_DATABASE contains database statistics for each database in GaussDB.

**Table 12-235** PG\_STAT\_DATABASE columns

| Name          | Type    | Description                                                                                                                                                                                                                             |
|---------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid         | oid     | OID of a database                                                                                                                                                                                                                       |
| datname       | name    | Database name.                                                                                                                                                                                                                          |
| numbackends   | integer | Number of backends currently connected to this database. This is the only column in the view that returns the status value. Other columns return the accumulated value since the last reset.                                            |
| xact_commit   | bigint  | Number of transactions in this database that have been committed.                                                                                                                                                                       |
| xact_rollback | bigint  | Number of transactions in this database that have been rolled back.                                                                                                                                                                     |
| blks_read     | bigint  | Number of disk blocks read in this database.                                                                                                                                                                                            |
| blks_hit      | bigint  | Number of times that disk blocks have been found in the buffer cache and therefore the disk blocks do not need to be read (only those found in the buffer cache are counted, excluding those found in the file system cache of the OS). |
| tup_returned  | bigint  | Number of rows returned by queries in this database                                                                                                                                                                                     |
| tup_fetched   | bigint  | Number of rows fetched by queries in this database                                                                                                                                                                                      |

| Name           | Type                     | Description                                                                                                                                                                                                                          |
|----------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tup_inserted   | bigint                   | Number of rows inserted by queries in this database                                                                                                                                                                                  |
| tup_updated    | bigint                   | Number of rows updated by queries in this database                                                                                                                                                                                   |
| tup_deleted    | bigint                   | Number of rows deleted by queries in this database                                                                                                                                                                                   |
| conflicts      | bigint                   | Number of queries canceled due to database recovery conflicts (conflicts occurring only on the standby node). For details, see <a href="#">PG_STAT_DATABASE_CONFLICTS</a> .                                                          |
| temp_files     | bigint                   | Number of temporary files created by queries in this database. Calculates all temporary files, regardless of why the temporary files are created (such as sorting or hashing) or the setting of the <b>log_temp_files</b> parameter. |
| temp_bytes     | bigint                   | Total amount of data written to temporary files by queries in this database. All temporary files are calculated regardless of why they are created or the setting of the <b>log_temp_files</b> parameter.                            |
| deadlocks      | bigint                   | Number of all deadlocks detected in the database since the last reset of statistics.                                                                                                                                                 |
| blk_read_time  | double precision         | Time spent in reading data file blocks by backends in this database (unit: ms).                                                                                                                                                      |
| blk_write_time | double precision         | Time spent in writing data file blocks by backends in this database (unit: ms).                                                                                                                                                      |
| stats_reset    | timestamp with time zone | Time when the status statistics are reset.                                                                                                                                                                                           |

## 12.3.147 PG\_STAT\_DATABASE\_CONFLICTS

**PG\_STAT\_DATABASE\_CONFLICTS** displays statistics about database conflicts.

**Table 12-236** PG\_STAT\_DATABASE\_CONFLICTS columns

| Name  | Type | Description |
|-------|------|-------------|
| datid | oid  | Database ID |

| Name             | Type   | Description                       |
|------------------|--------|-----------------------------------|
| datname          | name   | Database name                     |
| confl_tablespace | bigint | Number of conflicting tablespaces |
| confl_lock       | bigint | Number of conflicting locks       |
| confl_snapshot   | bigint | Number conflicting snapshots      |
| confl_bufferpin  | bigint | Number of conflicting buffers     |
| confl_deadlock   | bigint | Number of conflicting deadlocks   |

## 12.3.148 PG\_STAT\_USER\_FUNCTIONS

**PG\_STAT\_USER\_FUNCTIONS** shows user-defined function status information in the namespace. (The language of the function is non-internal language.)

**Table 12-237** PG\_STAT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                                     |
|------------|------------------|---------------------------------------------------------------------------------|
| funcid     | oid              | OID of a function                                                               |
| schemaname | name             | Schema name                                                                     |
| funcname   | name             | Function name                                                                   |
| calls      | bigint           | Number of times the function has been called                                    |
| total_time | double precision | Total time spent in the function and all other functions called by it           |
| self_time  | double precision | Total time spent in the function itself, excluding other functions called by it |

## 12.3.149 PG\_STAT\_USER\_INDEXES

**PG\_STAT\_USER\_INDEXES** displays information about the index status of user-defined ordinary tables and TOAST tables.

**Table 12-238** PG\_STAT\_USER\_INDEXES columns

| Name       | Type | Description                    |
|------------|------|--------------------------------|
| relid      | oid  | OID of the table for an index. |
| indexrelid | oid  | OID of the index.              |

| Name          | Type   | Description                                                                                         |
|---------------|--------|-----------------------------------------------------------------------------------------------------|
| schemaname    | name   | Name of the schema that the index is in.                                                            |
| relname       | name   | Name of the table for the index.                                                                    |
| indexrelname  | name   | Index name.                                                                                         |
| idx_scan      | bigint | Number of index scans initiated on the index.                                                       |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index.                                             |
| idx_tup_fetch | bigint | Number of live table rows fetched in the original table by a simple index scan that uses the index. |

## 12.3.150 PG\_STAT\_USER\_TABLES

PG\_STAT\_USER\_TABLES displays status information about user-defined ordinary tables and TOAST tables in the namespaces.

**Table 12-239** PG\_STAT\_USER\_TABLES columns

| Name          | Type   | Description                                                          |
|---------------|--------|----------------------------------------------------------------------|
| relid         | oid    | OID of a table.                                                      |
| schemaname    | name   | Name of the schema that this table is in.                            |
| relname       | name   | Table name.                                                          |
| seq_scan      | bigint | Number of sequential scans initiated on this table.                  |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.                     |
| idx_scan      | bigint | Number of index scans initiated on this table.                       |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.                          |
| n_tup_ins     | bigint | Number of rows inserted.                                             |
| n_tup_upd     | bigint | Number of rows updated.                                              |
| n_tup_del     | bigint | Number of rows deleted.                                              |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required). |

| Name              | Type                     | Description                                                                               |
|-------------------|--------------------------|-------------------------------------------------------------------------------------------|
| n_live_tup        | bigint                   | Estimated number of live rows.                                                            |
| n_dead_tup        | bigint                   | Estimated number of inactive rows.                                                        |
| last_vacuum       | timestamp with time zone | Last time when this table was manually vacuumed (excluding <b>VACUUM FULL</b> ).          |
| last_autovacuum   | timestamp with time zone | Last time when the table was vacuumed by the AUTOVACUUM daemon.                           |
| last_analyze      | timestamp with time zone | Last time when the table was manually analyzed.                                           |
| last_autoanalyze  | timestamp with time zone | Last time when the table was analyzed by the AUTOVACUUM daemon.                           |
| vacuum_count      | bigint                   | Number of times this table has been manually vacuumed (not counting <b>VACUUM FULL</b> ). |
| autovacuum_count  | bigint                   | Number of times this table has been vacuumed by the autovacuum daemon.                    |
| analyze_count     | bigint                   | Number of times this table has been manually analyzed.                                    |
| autoanalyze_count | bigint                   | Number of times this table has been analyzed by the autovacuum daemon.                    |
| last_data_changed | timestamp with time zone | Last modification time of the table data.                                                 |

## 12.3.151 PG\_STAT\_REPLICATION

PG\_STAT\_REPLICATION displays information about log synchronization status, such as the locations of the sender sending logs and the receiver receiving logs.

**Table 12-240** PG\_STAT\_REPLICATION columns

| Name     | Type   | Description        |
|----------|--------|--------------------|
| pid      | bigint | PID of the thread. |
| usesysid | oid    | User system ID.    |
| username | name   | Username.          |

| Name                     | Type                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| application_name         | text                     | Program name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| client_addr              | inet                     | Client address.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| client_hostname          | text                     | Client name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| client_port              | integer                  | Port of the client.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| backend_start            | timestamp with time zone | Start time of the program.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| state                    | text                     | Status of the log synchronization thread: <ul style="list-style-type: none"><li>• <b>startup</b>: The thread is being started.</li><li>• <b>catchup</b>: The thread is establishing a connection between the standby node and the primary node.</li><li>• <b>streaming</b>: The thread has established a connection between the standby node and the primary node and is replicating data streams.</li><li>• <b>backup</b>: The thread is sending a backup.</li><li>• <b>stopping</b>: The thread is being stopped.</li></ul> |
| sender_sent_location     | text                     | Location where the sender sends logs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| receiver_write_location  | text                     | Location where the receiver writes logs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| receiver_flush_location  | text                     | Location where the receive end flushes logs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| receiver_replay_location | text                     | Location where the receive end replays logs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| sync_priority            | integer                  | Priority of synchronous duplication ( <b>0</b> indicates asynchronization).                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

| Name       | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sync_state | text | <p>Synchronization state:</p> <ul style="list-style-type: none"><li>• <b>async</b>: asynchronous replication.</li><li>• <b>sync</b>: synchronous replication.</li><li>• <b>potential</b>: The standby node is asynchronous currently, but if a current synchronization node fails, the standby node becomes synchronous.</li><li>• <b>quorum</b>: switches between the synchronous and asynchronous states to ensure that there are more than a certain number of synchronous standby nodes. Generally, the number of synchronous standby nodes is <math>(n+1)/2-1</math>, where <math>n</math> indicates the total number of copies. Whether the standby node is synchronous depends on whether logs are received first. For details, see the description of the <b>synchronous_standby_names</b> parameter.</li></ul> |

### 12.3.152 PG\_STAT\_SYS\_INDEXES

PG\_STAT\_SYS\_INDEXES displays index status information about all the system catalogs in pg\_catalog and information\_schema.



**Table 12-241** PG\_STAT\_SYS\_INDEXES columns

| Name          | Type   | Description                                                                                         |
|---------------|--------|-----------------------------------------------------------------------------------------------------|
| relid         | oid    | OID of the table for this index.                                                                    |
| indexrelid    | oid    | OID of the index.                                                                                   |
| schemaname    | name   | Name of the schema that the index is in.                                                            |
| relname       | name   | Name of the table for the index.                                                                    |
| indexrelname  | name   | Index name.                                                                                         |
| idx_scan      | bigint | Number of index scans initiated on the index.                                                       |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index.                                             |
| idx_tup_fetch | bigint | Number of live table rows fetched in the original table by a simple index scan that uses the index. |

### 12.3.153 PG\_STAT\_SYS\_TABLES

PG\_STAT\_SYS\_TABLES displays statistics about the system catalogs of all the namespaces in pg\_catalog and information\_schema.

**Table 12-242** PG\_STAT\_SYS\_TABLES columns

| Name          | Type   | Description                                        |
|---------------|--------|----------------------------------------------------|
| relid         | oid    | OID of a table.                                    |
| schemaname    | name   | Name of the schema that the table is in.           |
| relname       | name   | Table name.                                        |
| seq_scan      | bigint | Number of sequential scans initiated on the table. |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.   |
| idx_scan      | bigint | Number of index scans initiated on the table.      |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.        |
| n_tup_ins     | bigint | Number of rows inserted.                           |
| n_tup_upd     | bigint | Number of rows updated.                            |
| n_tup_del     | bigint | Number of rows deleted.                            |

| Name              | Type                     | Description                                                                      |
|-------------------|--------------------------|----------------------------------------------------------------------------------|
| n_tup_hot_upd     | bigint                   | Number of rows HOT updated (with no separate index update required).             |
| n_live_tup        | bigint                   | Estimated number of live rows.                                                   |
| n_dead_tup        | bigint                   | Estimated number of inactive rows.                                               |
| last_vacuum       | timestamp with time zone | Last time when this table was manually vacuumed (excluding VACUUM FULL).         |
| last_autovacuum   | timestamp with time zone | Last time when the table was vacuumed by the autovacuum daemon.                  |
| last_analyze      | timestamp with time zone | Last time when the table was manually analyzed.                                  |
| last_autoanalyze  | timestamp with time zone | Last time when the table was analyzed by the autovacuum daemon.                  |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting VACUUM FULL). |
| autovacuum_count  | bigint                   | Number of times the table has been vacuumed by the autovacuum daemon.            |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed.                            |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon.            |
| last_data_changed | timestamp with time zone | Last modification time of the table data.                                        |

### 12.3.154 PG\_STAT\_XACT\_ALL\_TABLES

**PG\_STAT\_XACT\_ALL\_TABLES** displays transaction status information about all ordinary tables and TOAST tables in the namespaces.

**Table 12-243** PG\_STAT\_XACT\_ALL\_TABLES columns

| Name  | Type | Description      |
|-------|------|------------------|
| relid | oid  | OID of the table |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| schemaname    | name   | Name of the schema that the table is in                             |
| relname       | name   | Table name                                                          |
| seq_scan      | bigint | Number of sequential scans initiated on the table                   |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans                     |
| idx_scan      | bigint | Number of index scans initiated on the table                        |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans                          |
| n_tup_ins     | bigint | Number of rows inserted                                             |
| n_tup_upd     | bigint | Number of rows updated                                              |
| n_tup_del     | bigint | Number of rows deleted                                              |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 12.3.155 PG\_STAT\_XACT\_SYS\_TABLES

**PG\_STAT\_XACT\_SYS\_TABLES** displays transaction status information of the system catalog in the namespace.

**Table 12-244** PG\_STAT\_XACT\_SYS\_TABLES columns

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| relid         | oid    | OID of the table                                                    |
| schemaname    | name   | Name of the schema that the table is in                             |
| relname       | name   | Table name                                                          |
| seq_scan      | bigint | Number of sequential scans initiated on the table                   |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans                     |
| idx_scan      | bigint | Number of index scans initiated on the table                        |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans                          |
| n_tup_ins     | bigint | Number of rows inserted                                             |
| n_tup_upd     | bigint | Number of rows updated                                              |
| n_tup_del     | bigint | Number of rows deleted                                              |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

## 12.3.156 PG\_STAT\_XACT\_USER\_FUNCTIONS

**PG\_STAT\_XACT\_USER\_FUNCTIONS** contains statistics on the execution of each function.

**Table 12-245** PG\_STAT\_XACT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                                     |
|------------|------------------|---------------------------------------------------------------------------------|
| funcid     | oid              | OID of a function                                                               |
| schemaname | name             | Schema name                                                                     |
| funcname   | name             | Function name                                                                   |
| calls      | bigint           | Number of times the function has been called                                    |
| total_time | double precision | Total time spent in the function and all other functions called by it           |
| self_time  | double precision | Total time spent in the function itself, excluding other functions called by it |

## 12.3.157 PG\_STAT\_XACT\_USER\_TABLES

**PG\_STAT\_XACT\_USER\_TABLES** displays transaction status information of the user table in the namespace.

**Table 12-246** PG\_STAT\_XACT\_USER\_TABLES columns

| Name          | Type   | Description                                       |
|---------------|--------|---------------------------------------------------|
| relid         | oid    | OID of the table                                  |
| schemaname    | name   | Name of the schema that the table is in           |
| relname       | name   | Table name                                        |
| seq_scan      | bigint | Number of sequential scans initiated on the table |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans   |
| idx_scan      | bigint | Number of index scans initiated on the table      |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans        |
| n_tup_ins     | bigint | Number of rows inserted                           |
| n_tup_upd     | bigint | Number of rows updated                            |
| n_tup_del     | bigint | Number of rows deleted                            |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

## 12.3.158 PG\_STATIO\_ALL\_INDEXES

PG\_STATIO\_ALL\_INDEXES is used to query information about one row for each index in the current database, showing statistics about I/O on that specific index.

**Table 12-247** PG\_STATIO\_ALL\_INDEXES columns

| Name          | Type   | Description                                      |
|---------------|--------|--------------------------------------------------|
| relid         | oid    | OID of the table for the index.                  |
| indexrelid    | oid    | OID of the index.                                |
| schemaname    | name   | Name of the schema that the index is in.         |
| relname       | name   | Name of the table that the index is created for. |
| indexrelname  | name   | Index name.                                      |
| idx_blks_read | bigint | Number of disk blocks read from the index.       |
| idx_blks_hit  | bigint | Number of cache hits in the index.               |

## 12.3.159 PG\_STATIO\_ALL\_SEQUENCES

PG\_STATIO\_ALL\_SEQUENCES displays the I/O statistics of each sequence in the current database.

**Table 12-248** PG\_STATIO\_ALL\_SEQUENCES columns

| Name       | Type   | Description                                   |
|------------|--------|-----------------------------------------------|
| relid      | oid    | OID of this sequence.                         |
| schemaname | name   | Name of the schema that this sequence is in.  |
| relname    | name   | Name of the sequence.                         |
| blks_read  | bigint | Number of disk blocks read from the sequence. |

| Name     | Type   | Description                            |
|----------|--------|----------------------------------------|
| blks_hit | bigint | Number of cache hits in this sequence. |

### 12.3.160 PG\_STATIO\_ALL\_TABLES

PG\_STATIO\_ALL\_TABLES is used to query I/O statistics of each table (including TOAST tables) in the current database.

**Table 12-249** PG\_STATIO\_ALL\_TABLES columns

| Name            | Type   | Description                                                                    |
|-----------------|--------|--------------------------------------------------------------------------------|
| relid           | oid    | OID of a table.                                                                |
| schemaname      | name   | Name of the schema that this table is in.                                      |
| relname         | name   | Table name.                                                                    |
| heap_blks_read  | bigint | Number of disk blocks read from this table.                                    |
| heap_blks_hit   | bigint | Number of cache hits in this table.                                            |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in this table.                     |
| idx_blks_hit    | bigint | Number of cache hits of all indexes in this table.                             |
| toast_blks_read | bigint | Number of disk blocks read from TOAST tables (if any) in the table.            |
| toast_blks_hit  | bigint | Number of cache hits (if any) in the TOAST table of this table.                |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table indexes (if any) in the table. |
| tidx_blks_hit   | bigint | Number of cache hits (if any) in the TOAST table indexes of this table.        |

### 12.3.161 PG\_STATIO\_SYS\_INDEXES

PG\_STATIO\_SYS\_INDEXES displays I/O status information about all system catalog indexes in the namespaces.

**Table 12-250** PG\_STATIO\_SYS\_INDEXES columns

| Name          | Type   | Description                                      |
|---------------|--------|--------------------------------------------------|
| relid         | oid    | OID of the table for the index.                  |
| indexrelid    | oid    | OID of the index.                                |
| schemaname    | name   | Name of the schema that the index is in.         |
| relname       | name   | Name of the table that the index is created for. |
| indexrelname  | name   | Index name.                                      |
| idx_blks_read | bigint | Number of disk blocks read from the index.       |
| idx_blks_hit  | bigint | Number of cache hits in the index.               |

## 12.3.162 PG\_STATIO\_SYS\_SEQUENCES

PG\_STATIO\_SYS\_SEQUENCES displays I/O status information about all the sequences in the namespaces.

**Table 12-251** PG\_STATIO\_SYS\_SEQUENCES columns

| Name       | Type   | Description                                    |
|------------|--------|------------------------------------------------|
| relid      | oid    | OID of this sequence.                          |
| schemaname | name   | Name of the schema that this sequence is in.   |
| relname    | name   | Name of this sequence.                         |
| blks_read  | bigint | Number of disk blocks read from this sequence. |
| blks_hit   | bigint | Number of cache hits in this sequence.         |

## 12.3.163 PG\_STATIO\_SYS\_TABLES

PG\_STATIO\_SYS\_TABLES shows I/O status information about all the system catalogs in the namespaces.

**Table 12-252** PG\_STATIO\_SYS\_TABLES columns

| Name  | Type | Description     |
|-------|------|-----------------|
| relid | oid  | OID of a table. |

| Name            | Type   | Description                                                                    |
|-----------------|--------|--------------------------------------------------------------------------------|
| schemaname      | name   | Name of the schema that this table is in.                                      |
| relname         | name   | Table name.                                                                    |
| heap_blks_read  | bigint | Number of disk blocks read from this table.                                    |
| heap_blks_hit   | bigint | Number of cache hits in this table.                                            |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in this table.                     |
| idx_blks_hit    | bigint | Number of cache hits of all indexes in this table.                             |
| toast_blks_read | bigint | Number of disk blocks read from TOAST tables (if any) in the table.            |
| toast_blks_hit  | bigint | Number of cache hits (if any) in the TOAST table of this table.                |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table indexes (if any) in the table. |
| tidx_blks_hit   | bigint | Number of cache hits (if any) in the TOAST table indexes of this table.        |

## 12.3.164 PG\_STATIO\_USER\_INDEXES

PG\_STATIO\_USER\_INDEXES displays I/O status information about all the user relationship table indexes in the namespaces.

**Table 12-253** PG\_STATIO\_USER\_INDEXES columns

| Name          | Type   | Description                                |
|---------------|--------|--------------------------------------------|
| relid         | oid    | OID of the table for the index.            |
| indexrelid    | oid    | OID of the index.                          |
| schemaname    | name   | Name of the schema that the index is in.   |
| relname       | name   | Name of the table for the index.           |
| indexrelname  | name   | Index name.                                |
| idx_blks_read | bigint | Number of disk blocks read from the index. |
| idx_blks_hit  | bigint | Number of cache hits in the index.         |



## 12.3.165 PG\_STATIO\_USER\_SEQUENCES

PG\_STATIO\_USER\_SEQUENCES shows I/O status information about all the user relationship table sequences in the namespaces.

**Table 12-254** PG\_STATIO\_USER\_SEQUENCES columns

| Name       | Type   | Description                                   |
|------------|--------|-----------------------------------------------|
| relid      | oid    | OID of this sequence.                         |
| schemaname | name   | Name of the schema where the sequence is in.  |
| relname    | name   | Name of the sequence.                         |
| blks_read  | bigint | Number of disk blocks read from the sequence. |
| blks_hit   | bigint | Number of cache hits in this sequence.        |

## 12.3.166 PG\_STATIO\_USER\_TABLES

PG\_STATIO\_USER\_TABLES displays I/O status information about all the user relationship tables in the namespaces.

**Table 12-255** PG\_STATIO\_USER\_TABLES columns

| Name            | Type   | Description                                                                    |
|-----------------|--------|--------------------------------------------------------------------------------|
| relid           | oid    | OID of a table.                                                                |
| schemaname      | name   | Name of the schema that this table is in.                                      |
| relname         | name   | Table name.                                                                    |
| heap_blks_read  | bigint | Number of disk blocks read from this table.                                    |
| heap_blks_hit   | bigint | Number of cache hits in this table.                                            |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in this table.                     |
| idx_blks_hit    | bigint | Number of cache hits of all indexes in this table.                             |
| toast_blks_read | bigint | Number of disk blocks read from TOAST tables (if any) in the table.            |
| toast_blks_hit  | bigint | Number of cache hits (if any) in the TOAST table of this table.                |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table indexes (if any) in the table. |

| Name          | Type   | Description                                                             |
|---------------|--------|-------------------------------------------------------------------------|
| tidx_blks_hit | bigint | Number of cache hits (if any) in the TOAST table indexes of this table. |

## 12.3.167 PG\_TABLES

PG\_TABLES is used to query useful information about each table in a database.

**Table 12-256** PG\_TABLES columns

| Name         | Type                     | Reference                                                    | Description                                                                                                                               |
|--------------|--------------------------|--------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| schemaname   | name                     | <b>nspname</b> in <a href="#">PG_NAMESPACE</a>               | Name of the schema that contains a table.                                                                                                 |
| tablename    | name                     | <b>relname</b> in <a href="#">PG_CLASS</a>                   | Table name.                                                                                                                               |
| tableowner   | name                     | <b>pg_get_userbyid(reowner)</b> in <a href="#">PG_CLASS</a>  | Table owner.                                                                                                                              |
| tablespace   | name                     | <b>spcname</b> in <a href="#">PG_TABLESPACE</a>              | Tablespace that contains the table (default value: <b>NULL</b> ).                                                                         |
| hasindexes   | boolean                  | <b>relhasindex</b> in <a href="#">PG_CLASS</a>               | Specifies whether the table has (or recently had) an index. If it does, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> . |
| hasrules     | boolean                  | <b>relhasrules</b> in <a href="#">PG_CLASS</a>               | Specifies whether the table has rules. If it does, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> .                      |
| hastriggers  | boolean                  | <b>RELHASTRIGGERS</b> in <a href="#">PG_CLASS</a>            | Specifies whether the table has triggers. If it does, the value is <b>TRUE</b> . Otherwise, the value is <b>FALSE</b> .                   |
| tablecreator | name                     | <b>pg_get_userbyid(creator)</b> in <a href="#">PG_OBJECT</a> | Table creator.                                                                                                                            |
| created      | timestamp with time zone | <b>ctime</b> in <a href="#">PG_OBJECT</a>                    | Time when the table is created.                                                                                                           |

| Name          | Type                     | Reference                                 | Description                                                 |
|---------------|--------------------------|-------------------------------------------|-------------------------------------------------------------|
| last_ddl_time | timestamp with time zone | <b>mtime</b> in <a href="#">PG_OBJECT</a> | Time when the last DDL operation is performed on the table. |

## 12.3.168 PG\_TDE\_INFO

PG\_TDE\_INFO displays encryption information of the entire database.

**Table 12-257** PG\_TDE\_INFO columns

| Name       | Type    | Description                                                                                                                                                                  |
|------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| is_encrypt | Boolean | Encryption status of the database. <ul style="list-style-type: none"><li>• <b>f</b>: The database is not encrypted.</li><li>• <b>t</b>: The database is encrypted.</li></ul> |
| g_tde_algo | text    | Encryption algorithm. <ul style="list-style-type: none"><li>• <b>SM4-CTR-128</b></li><li>• <b>AES-CTR-128</b></li></ul>                                                      |
| remain     | text    | Reserved column.                                                                                                                                                             |

## 12.3.169 PG\_TIMEZONE\_ABBREVS

PG\_TIMEZONE\_ABBREVS displays information about all available time zones.

**Table 12-258** PG\_TIMEZONE\_ABBREVS columns

| Name       | Type     | Description                                                                                                                          |
|------------|----------|--------------------------------------------------------------------------------------------------------------------------------------|
| abbrev     | text     | Abbreviation of the time zone name.                                                                                                  |
| utc_offset | interval | Offset from UTC.                                                                                                                     |
| is_dst     | Boolean  | Specifies whether the DST is being used. If the DST is being used, the value is <b>TRUE</b> ; otherwise, the value is <b>FALSE</b> . |

## 12.3.170 PG\_TIMEZONE\_NAMES

PG\_TIMEZONE\_NAMES displays all time zone names that can be recognized by SET TIMEZONE, along with their abbreviations, UTC offsets, and daylight saving time (DST) statuses.

**Table 12-259** PG\_TIMEZONE\_NAMES columns

| Name       | Type     | Description                                                                                                                          |
|------------|----------|--------------------------------------------------------------------------------------------------------------------------------------|
| name       | text     | Name of the time zone.                                                                                                               |
| abbrev     | text     | Abbreviation of the time zone name.                                                                                                  |
| utc_offset | interval | Offset from UTC.                                                                                                                     |
| is_dst     | boolean  | Specifies whether the DST is being used. If the DST is being used, the value is <b>TRUE</b> ; otherwise, the value is <b>FALSE</b> . |

## 12.3.171 PG\_TOTAL\_MEMORY\_DETAIL

PG\_TOTAL\_MEMORY\_DETAIL displays memory usage of a node in the database.

**Table 12-260** PG\_TOTAL\_MEMORY\_DETAIL columns

| Name        | Type    | Description                                |
|-------------|---------|--------------------------------------------|
| nodename    | text    | Specifies the node name.                   |
| memorytype  | text    | Memory name                                |
| memorybytes | integer | Size of the used memory in the unit of MB. |

## 12.3.172 PG\_TOTAL\_USER\_RESOURCE\_INFO

PG\_TOTAL\_USER\_RESOURCE\_INFO displays resource usage of all users. Only administrators can query this view. This view is valid only when the GUC parameter **use\_workload\_manager** is set to **on**. I/O monitoring items are valid only when **enable\_logical\_io\_statistics** is set to **on**.

**Table 12-261** PG\_TOTAL\_USER\_RESOURCE\_INFO columns

| Name        | Type    | Description         |
|-------------|---------|---------------------|
| username    | name    | Username.           |
| used_memory | integer | Used memory, in MB. |

| Name              | Type             | Description                                                                                                                                                                                                                                               |
|-------------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| total_memory      | integer          | Available memory, in MB. The value 0 indicates that the available memory is not limited and depends on the maximum memory available in the database.                                                                                                      |
| used_cpu          | double precision | Number of CPU cores in use. CPU usage data is collected only in complex jobs, and the value is the CPU usage of the related Cgroup.                                                                                                                       |
| total_cpu         | integer          | Total number of CPU cores of the Cgroup associated with the user on the node.                                                                                                                                                                             |
| used_space        | bigint           | Used permanent table storage space, in KB.                                                                                                                                                                                                                |
| total_space       | bigint           | Available permanent table storage space, in KB (-1 if the storage space is not limited).                                                                                                                                                                  |
| used_temp_space   | bigint           | Used temporary storage space, in KB.                                                                                                                                                                                                                      |
| total_temp_space  | bigint           | Total available temporary space, in KB (-1 if the temporary space is not limited).                                                                                                                                                                        |
| used_spill_space  | bigint           | Size of the used operator-level data flushing space, in KB.                                                                                                                                                                                               |
| total_spill_space | bigint           | Total size of the available operator-level data flushing space, in KB (-1 if the space is not limited).                                                                                                                                                   |
| read_kbytes       | bigint           | Primary database node: total bytes read by the user's complex jobs on all database nodes in the last 5 seconds, in KB.<br><br>Database node: total bytes read by the user's complex jobs from the instance startup time to the current time, in KB.       |
| write_kbytes      | bigint           | Primary database node: total bytes written by the user's complex jobs on all database nodes in the last 5 seconds, in KB.<br><br>Database node: total bytes written by the user's complex jobs from the instance startup time to the current time, in KB. |
| read_counts       | bigint           | Primary database node: total number of read times of the user's complex jobs on all database nodes in the last 5 seconds.<br><br>Database node: total number of read times of the user's complex jobs from the instance startup time to the current time. |

| Name         | Type             | Description                                                                                                                                                                                                                                             |
|--------------|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| write_counts | bigint           | Primary database node: total number of write times of the user's complex jobs on all database nodes in the last 5 seconds.<br>Database node: total number of write times of the user's complex jobs from the instance startup time to the current time. |
| read_speed   | double precision | Primary database node: average read rate of the user's complex jobs on a single database node in the last 5 seconds, in KB/s.<br>Database node: average read rate of the user's complex jobs on the database node in the last 5 seconds, in KB/s.       |
| write_speed  | double precision | Primary database node: average write rate of the user's complex jobs on a single database node in the last 5 seconds, in KB/s.<br>Database node: average write rate of the user's complex jobs on the database node in the last 5 seconds, in KB/s.     |

### 12.3.173 PG\_TOTAL\_USER\_RESOURCE\_INFO\_OID

PG\_TOTAL\_USER\_RESOURCE\_INFO\_OID displays resource usage of all users. Only administrators can query this view. This view is valid only when the GUC parameter **use\_workload\_manager** is set to **on**.

**Table 12-262** PG\_TOTAL\_USER\_RESOURCE\_INFO\_OID columns

| Name         | Type             | Description                                                                                                                                                    |
|--------------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| userid       | oid              | User ID.                                                                                                                                                       |
| used_memory  | integer          | Size of the memory being used, in MB.                                                                                                                          |
| total_memory | integer          | Available memory (unit: MB) The value <b>0</b> indicates that the available memory is not limited and depends on the maximum memory available in the database. |
| used_cpu     | double precision | Number of CPU cores in use.                                                                                                                                    |
| total_cpu    | integer          | Total number of CPU cores of the Cgroup associated with the user on the node.                                                                                  |
| used_space   | bigint           | Used storage space, in KB.                                                                                                                                     |
| total_space  | bigint           | Available storage space, in KB (-1 if the space is not limited).                                                                                               |

| Name              | Type             | Description                                                                        |
|-------------------|------------------|------------------------------------------------------------------------------------|
| used_temp_space   | bigint           | Used temporary storage space, in KB.                                               |
| total_temp_space  | bigint           | Total available temporary space, in KB (-1 if the temporary space is not limited). |
| used_spill_space  | bigint           | Used disk space for spilling, in KB.                                               |
| total_spill_space | bigint           | Total available disk space for spilling, in KB (-1 if the space is not limited).   |
| read_kbytes       | bigint           | Amount of data read from the disk, in KB.                                          |
| write_kbytes      | bigint           | Amount of data written to the disk, in KB.                                         |
| read_counts       | bigint           | Number of disk read times.                                                         |
| write_counts      | bigint           | Number of disk write times.                                                        |
| read_speed        | double precision | Disk read rate, in B/ms.                                                           |
| write_speed       | double precision | Disk write rate, in B/ms.                                                          |

### 12.3.174 PG\_USER

PG\_USER displays information about database users. By default, only the initial user and users with the SYSADMIN attribute can view the information. Other users can view the information only after being granted permissions.

**Table 12-263** PG\_USER columns

| Name        | Type    | Description                                                                                                                                                                              |
|-------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| username    | name    | Username.                                                                                                                                                                                |
| usesysid    | oid     | ID of this user.                                                                                                                                                                         |
| usecreatedb | boolean | Specifies whether a user has the permissions to create databases. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>                |
| usesuper    | boolean | Specifies whether the user is the initial SYSADMIN with the highest permissions. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |

| Name            | Type                     | Description                                                                                                                                                                                                                                                                                                         |
|-----------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| usecatupd       | boolean                  | Specifies whether the user can directly update system catalogs. Only the initial SYSADMIN whose <b>usesysid</b> is set to <b>10</b> has this permission. This permission is unavailable for other users. <ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul> |
| userepl         | boolean                  | Specifies whether the user has the permissions to duplicate data streams. <ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                |
| passwd          | text                     | Encrypted user password. The value is displayed as <b>*****</b> .                                                                                                                                                                                                                                                   |
| valbegin        | timestamp with time zone | Start time for account validity ( <b>NULL</b> if start time is not specified).                                                                                                                                                                                                                                      |
| valuntil        | timestamp with time zone | End time for account validity ( <b>NULL</b> if end time is not specified).                                                                                                                                                                                                                                          |
| respool         | name                     | Resource pool where the user is in.                                                                                                                                                                                                                                                                                 |
| parent          | oid                      | Parent user OID.                                                                                                                                                                                                                                                                                                    |
| spacelimit      | text                     | Storage space of the permanent table, in KB.                                                                                                                                                                                                                                                                        |
| useconfig       | text[]                   | Default value of runtime configuration items. For details, see <b>setconfig</b> in <a href="#">PG_DB_ROLE_SETTING</a> .                                                                                                                                                                                             |
| nodegroup       | name                     | Name of the logical database associated with the user. If the user does not manage the logical database, this column is left blank.                                                                                                                                                                                 |
| tempspacelimit  | text                     | Storage space of the temporary table, in KB.                                                                                                                                                                                                                                                                        |
| spillspacelimit | text                     | Operator disk flushing space, in KB.                                                                                                                                                                                                                                                                                |
| usemonitoradmin | Boolean                  | Specifies whether the user is a MONADMIN. <ul style="list-style-type: none"> <li>• <b>t</b> (true): yes.</li> <li>• <b>f</b> (false): no.</li> </ul>                                                                                                                                                                |



| Name             | Type    | Description                                                                                                                                        |
|------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| useoperatoradmin | Boolean | Specifies whether the user is an OPRADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul> |
| usepolicyadmin   | Boolean | Specifies whether the user is a POLADMIN. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes.</li><li>• <b>f</b> (false): no.</li></ul>  |

### 12.3.175 PG\_USER\_MAPPINGS

PG\_USER\_MAPPINGS displays user mapping information.

This is essentially a publicly readable view of [PG\\_USER\\_MAPPING](#) that leaves out the options column if the user has no rights to use the system catalog and query the view. Common users must be authorized to access this view.

**Table 12-264** PG\_USER\_MAPPINGS columns

| Name      | Type   | Reference                                           | Description                                                                                                                                                            |
|-----------|--------|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| umid      | oid    | <b>oid</b> in <a href="#">PG_USER_MAPPING</a>       | OID of the user mapping.                                                                                                                                               |
| srvid     | oid    | <b>oid</b> in <a href="#">PG_FOREIGN_SERVER</a>     | OID of the foreign server that contains the mapping.                                                                                                                   |
| srvname   | name   | <b>srvname</b> in <a href="#">PG_FOREIGN_SERVER</a> | Name of the foreign server.                                                                                                                                            |
| umuser    | oid    | <b>oid</b> in <a href="#">PG_AUTHID</a>             | OID of the local role being mapped ( <b>0</b> if the user mapping is public)                                                                                           |
| username  | name   | -                                                   | Name of the local user to be mapped.                                                                                                                                   |
| umoptions | text[] | -                                                   | User mapping specific options. If the current user is the owner of the foreign server, it is a string in the format of "keyword=value". Otherwise, it is <b>NULL</b> . |

### 12.3.176 PG\_VARIABLE\_INFO

PG\_VARIABLE\_INFO records information about transaction IDs and OIDs of the current node in the database.

**Table 12-265** PG\_VARIABLE\_INFO columns

| Name                         | Type | Description                                                                |
|------------------------------|------|----------------------------------------------------------------------------|
| node_name                    | text | Node name                                                                  |
| next_oid                     | oid  | OID generated next time for the node                                       |
| next_xid                     | xid  | Transaction ID generated next time for the node                            |
| oldest_xid                   | xid  | Oldest transaction ID on the node                                          |
| xid_vac_limit                | xid  | Critical point (transaction ID) that triggers forcible autovacuum          |
| oldest_xid_db                | oid  | OID of the database that has the minimum datafrozensxid on the node        |
| last_extend_cs<br>n_logpage  | xid  | Number of the last extended cslog page                                     |
| start_extend_cs<br>n_logpage | xid  | Number of the page from which cslog extending starts                       |
| next_commit_s<br>eqno        | xid  | CSN generated next time for the node                                       |
| latest_comple<br>ted_xid     | xid  | Latest transaction ID on the node after the transaction commit or rollback |
| startup_max_xi<br>d          | xid  | Last transaction ID before the node is powered off                         |

## 12.3.177 PG\_VIEWS

PG\_VIEWS displays useful information about each view in the database.

**Table 12-266** PG\_VIEWS columns

| Name       | Type | Reference                                      | Description            |
|------------|------|------------------------------------------------|------------------------|
| schemaname | name | <b>nspname</b> in <a href="#">PG_NAMESPACE</a> | Schema name of a view. |
| viewname   | name | <b>relname</b> in <a href="#">PG_CLASS</a>     | View name              |
| viewowner  | name | <b>Erolname</b> in <a href="#">PG_AUTHID</a>   | Owner of the view      |
| definition | text | -                                              | Definition of the view |

## 12.3.178 PGXC\_PREPARED\_XACTS

PGXC\_PREPARED\_XACTS displays two-phase transactions in the prepared phase. Only users with the SYSADMIN or MONADMIN permissions can access this view.

**Table 12-267** PGXC\_PREPARED\_XACTS columns

| Name               | Type | Description                                                                |
|--------------------|------|----------------------------------------------------------------------------|
| pgxc_prepared_xact | text | Displays the two-phase transaction currently in the <b>prepared</b> phase. |

## 12.3.179 PGXC\_THREAD\_WAIT\_STATUS

This view is not supported in centralized scenarios.

## 12.3.180 PLAN\_TABLE

PLAN\_TABLE displays plan information collected by **EXPLAIN PLAN**. Plan information is in a session-level lifecycle. After a session exits, the data will be deleted. Data is isolated between sessions and between users.

**Table 12-268** PLAN\_TABLE columns

| Name         | Type           | Description                                                                                                                       |
|--------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------|
| statement_id | varchar2(30)   | Query tag specified by a user.                                                                                                    |
| plan_id      | bigint         | Query ID.                                                                                                                         |
| id           | int            | ID of each operator in a generated plan.                                                                                          |
| operation    | varchar2(30)   | Operation description of an operator in a plan.                                                                                   |
| options      | varchar2(255)  | Operation option.                                                                                                                 |
| object_name  | name           | Object name corresponding to the operation, which is not the object alias used in the query. The object name is defined by users. |
| object_type  | varchar2(30)   | Object type.                                                                                                                      |
| object_owner | name           | Schema to which the object belongs. It is defined by users.                                                                       |
| projection   | varchar2(4000) | Returned column information.                                                                                                      |
| cost         | float8         | Execution cost estimated by the optimizer for an operator.                                                                        |
| cardinality  | float8         | Number of rows estimated by the optimizer for an operator.                                                                        |

 NOTE

- A valid **object\_type** value consists of a relkind type defined in **PG\_CLASS** (**TABLE**, **INDEX**, **SEQUENCE**, **VIEW**, **COMPOSITE TYPE**, or **TOASTVALUE TOAST**) and the rtekind type used in the plan (**SUBQUERY**, **JOIN**, **FUNCTION**, **VALUES**, **CTE**, or **REMOTE\_QUERY**).
- For RangeTableEntry (RTE), **object\_owner** is the object description used in the plan. Non-user-defined objects do not have **object\_owner**.
- Information in the **statement\_id**, **object\_name**, **object\_owner**, and **projection** columns is stored in letter cases specified by users and information in other columns is stored in uppercase.
- **PLAN\_TABLE** supports only **SELECT** and **DELETE** and does not support other DML operations.

## 12.3.181 SYS\_DUMMY

SYS\_DUMMY is automatically created by the database based on the data dictionary. It has only one text column in only one row for storing expression calculation results. This view is accessible to all users. This view exists in both PG\_CATALOG and SYS schemas.

**Table 12-269** SYS\_DUMMY columns

| Name  | Type | Description                   |
|-------|------|-------------------------------|
| dummy | text | Expression calculation result |

# 13 Schemas

The following table describes the schemas supported in GaussDB.

**Table 13-1** Schemas supported in GaussDB

| Schema         | Description                                                                                                                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dbe_perf       | Diagnoses performance issues and is also the data source of WDR snapshots. After a database is installed, only the initial user and users with the MONADMIN permission can access views and functions in the dbe_perf schema by default.                                                                                         |
| dbe_pldebugger | Debugs PL/SQL functions and stored procedures.                                                                                                                                                                                                                                                                                   |
| snapshot       | Manages data related to WDR snapshots. By default, the initial user or users with the MONADMIN permission can access the data.                                                                                                                                                                                                   |
| sys            | Provides the system information view APIs.                                                                                                                                                                                                                                                                                       |
| pg_catalog     | Maintains system catalog information, including system catalogs and all built-in data types, functions, and operators.                                                                                                                                                                                                           |
| pg_toast       | Stores large objects (for internal use).                                                                                                                                                                                                                                                                                         |
| public         | Public schema, which is used to store public objects. If <b>search_path</b> is not specified and a schema with the same name exists, new tables (and other objects) are created in the schema by default. If no schema with the same name exists, new tables (and other objects) are automatically placed in this public schema. |
| pkg_service    | Manages information about the package service.                                                                                                                                                                                                                                                                                   |
| pkg_util       | Manages information about the package tool.                                                                                                                                                                                                                                                                                      |
| dbe_raw        | Advanced function package dbe_raw, which is used to convert raw data, obtain substrings, and calculate the length.                                                                                                                                                                                                               |

| Schema               | Description                                                                                                                                                                                                                                                                     |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dbe_session          | Advanced function package dbe_session, which is used to set the value of a specified attribute and support user query and verification.                                                                                                                                         |
| dbe_lob              | Advanced function package dbe_lob, which is used to read, write, and copy large files (CLOB/BLOB).                                                                                                                                                                              |
| dbe_match            | Advanced function package dbe_match, which is used to compare character string similarity.                                                                                                                                                                                      |
| dbe_task             | Advanced function package dbe_task, which is used to schedule job tasks, including committing tasks, canceling tasks, synchronizing task status, and updating task information, so that the database can periodically execute specific tasks.                                   |
| dbe_sql              | Advanced function package dbe_sql, which is used to execute dynamic SQL statements and construct query and other commands during application running.                                                                                                                           |
| dbe_file             | Advanced function package <b>dbe_file</b> , which is used to read, copy, write, delete, and rename external database files.                                                                                                                                                     |
| dbe_output           | Advanced function package dbe_output, which is used to print output information.                                                                                                                                                                                                |
| dbe_random           | Advanced function package dbe_random, which is used to generate random seeds and random numbers.                                                                                                                                                                                |
| dbe_application_info | Advanced function package dbe_application_info, which is used for recording client information.                                                                                                                                                                                 |
| dbe_utility          | Advanced function package <b>dbe_utility</b> , which is used to call the debugging tool in a stored procedure, for example, to print error stacks.                                                                                                                              |
| dbe_scheduler        | Advanced function package <b>dbe_scheduler</b> , which is used to create scheduled jobs and enable the database to periodically execute specified tasks through programs and schedules. You can also perform external database tasks by authorizing and providing certificates. |
| information_schema   | Stores information about objects defined in the current database.                                                                                                                                                                                                               |
| dbe_pldeveloper      | Compiles and debugs user stored procedures.                                                                                                                                                                                                                                     |
| dbe_sql_util         | SQL O&M function, including the O&M API of SQL patches.                                                                                                                                                                                                                         |

## 13.1 Information Schema

An information schema named INFORMATION\_SCHEMA automatically exists in all databases. An information schema consists of a group of views that contain information about objects defined in the current database. The owner of this schema is the initial database user. However, all users have only the permission to use this schema and do not have the permission to create objects such as tables and functions.

Information schemas are compatible with PG, including: constraint\_table\_usage, domain\_constraints, domain\_udt\_usage, domains, enabled\_roles, key\_column\_usage, parameters, referential\_constraints, applicable\_roles, administrable\_role\_authorizations, attributes, character\_sets, check\_constraint\_routine\_usage, check\_constraints, collations, collation\_character\_set\_applicability, column\_domain\_usage, column\_privileges, column\_udt\_usage, columns, constraint\_column\_usage, role\_column\_grants, routine\_privileges, role\_routine\_grants, routines, schemata, sequences, table\_constraints, table\_privileges, role\_table\_grants, tables, triggered\_update\_columns, triggers, udt\_privileges, role\_udt\_grants, usage\_privileges, role\_usage\_grants, user\_defined\_types, view\_column\_usage, view\_routine\_usage, view\_table\_usage, views, data\_type\_privileges, element\_types, column\_options, foreign\_data\_wrapper\_options, foreign\_data\_wrappers, foreign\_server\_options, foreign\_servers, foreign\_table\_options, foreign\_tables, user\_mapping\_options, user\_mappings, sql\_features, sql\_implementation\_info, sql\_languages, sql\_packages, sql\_parts, sql\_sizing, and sql\_sizing\_profiles.

The following sections display only the views that are not listed in the preceding description.

### 13.1.1 \_PG\_FOREIGN\_DATA\_WRAPPERS

**\_PG\_FOREIGN\_DATA\_WRAPPERS** displays information about a foreign-data wrapper. Only users with the SYSADMIN permission can query this view.

**Table 13-2** \_PG\_FOREIGN\_DATA\_WRAPPERS columns

| Name                         | Type                              | Description                                                                                       |
|------------------------------|-----------------------------------|---------------------------------------------------------------------------------------------------|
| oid                          | oid                               | OID of the foreign data wrapper.                                                                  |
| fdwowner                     | oid                               | OID of the owner of the foreign data wrapper.                                                     |
| fdwoptions                   | text[]                            | Foreign-data wrapper specific option, expressed in a string in the format of <i>keyword=value</i> |
| foreign_data_wrapper_catalog | information_schema.sql_identifier | Name of the database where the foreign-data wrapper is located (always the current database)      |

|                               |                                   |                                                  |
|-------------------------------|-----------------------------------|--------------------------------------------------|
| foreign_data_wrapper_name     | information_schema.sql_identifier | Name of the foreign-data wrapper                 |
| authorization_identifier      | information_schema.sql_identifier | Role of the owner of the foreign-data wrapper    |
| foreign_data_wrapper_language | information_schema.character_data | Programming language of the foreign-data wrapper |

## 13.1.2 \_PG\_FOREIGN\_SERVERS

**\_PG\_FOREIGN\_SERVERS** displays information about a foreign server. Only users with the SYSADMIN permission can query this view.

**Table 13-3** \_PG\_FOREIGN\_SERVERS columns

| Name                         | Type                              | Description                                                                                  |
|------------------------------|-----------------------------------|----------------------------------------------------------------------------------------------|
| oid                          | oid                               | OID of the foreign server.                                                                   |
| srvoptions                   | text[]                            | Foreign server specific options, expressed in a string in the format of <i>keyword=value</i> |
| foreign_server_catalog       | information_schema.sql_identifier | Name of the database where the foreign server is located (always the current database)       |
| foreign_server_name          | information_schema.sql_identifier | Name of the foreign server                                                                   |
| foreign_data_wrapper_catalog | information_schema.sql_identifier | Name of the database where the foreign-data wrapper is located (always the current database) |
| foreign_data_wrapper_name    | information_schema.sql_identifier | Name of the foreign-data wrapper                                                             |
| foreign_server_type          | information_schema.character_data | Type of the foreign server                                                                   |
| foreign_server_version       | information_schema.character_data | Version of the foreign server                                                                |
| authorization_identifier     | information_schema.sql_identifier | Role of the owner of the foreign server                                                      |



### 13.1.3 \_PG\_FOREIGN\_TABLE\_COLUMNS

**\_PG\_FOREIGN\_TABLE\_COLUMNS** displays column information about a foreign table. Only users with the SYSADMIN permission can query this view.

**Table 13-4** \_PG\_FOREIGN\_TABLE\_COLUMNS columns

| Name          | Type   | Description                                                                                               |
|---------------|--------|-----------------------------------------------------------------------------------------------------------|
| nspname       | name   | Schema name                                                                                               |
| relname       | name   | Table name                                                                                                |
| attname       | name   | Column name                                                                                               |
| attfdwoptions | text[] | Attribute-level foreign data wrapper options, expressed in a string in the format of <i>keyword=value</i> |

### 13.1.4 \_PG\_FOREIGN\_TABLES

**\_PG\_FOREIGN\_TABLES** stores information about all foreign tables defined in the current database, but it only displays the tables that the current user has access to. Only users with the SYSADMIN permission can query this view.

**Table 13-5** \_PG\_FOREIGN\_TABLES columns

| Name                     | Type                              | Description                                                                            |
|--------------------------|-----------------------------------|----------------------------------------------------------------------------------------|
| foreign_table_catalog    | information_schema.sql_identifier | Name of the database where the foreign table is located (always the current database)  |
| foreign_table_schema     | name                              | Name of the schema that the foreign table is in                                        |
| foreign_table_name       | name                              | Name of the foreign table                                                              |
| ftoptions                | text[]                            | Foreign table options                                                                  |
| foreign_server_catalog   | information_schema.sql_identifier | Name of the database where the foreign server is located (always the current database) |
| foreign_server_name      | information_schema.sql_identifier | Name of the foreign server                                                             |
| authorization_identifier | information_schema.sql_identifier | Role of the owner                                                                      |

### 13.1.5 \_PG\_USER\_MAPPINGS

Stores mappings from local users to remote users. Only users with the SYSADMIN permission can query this view.

**Table 13-6** \_PG\_USER\_MAPPINGS columns

| Name                     | Type                              | Description                                                                                  |
|--------------------------|-----------------------------------|----------------------------------------------------------------------------------------------|
| oid                      | oid                               | OID of the mapping from the local user to a remote user.                                     |
| umoptions                | text[]                            | User mapping specific options, expressed in a string in the format of <i>keyword=value</i> . |
| umuser                   | oid                               | OID of the local user being mapped ( <b>0</b> if the user mapping is public).                |
| authorization_identifier | information_schema.sql_identifier | Role of the local user.                                                                      |
| foreign_server_catalog   | information_schema.sql_identifier | Name of the database where the foreign server is defined in.                                 |
| foreign_server_name      | information_schema.sql_identifier | Name of the foreign server.                                                                  |
| srvowner                 | information_schema.sql_identifier | Owner of the foreign server.                                                                 |

### 13.1.6 INFORMATION\_SCHEMA\_CATALOG\_NAME

Displays the name of the current database.

**Table 13-7** INFORMATION\_SCHEMA\_CATALOG\_NAME columns

| Name         | Type                              | Description            |
|--------------|-----------------------------------|------------------------|
| catalog_name | information_schema.sql_identifier | Current database name. |

## 13.2 DBE\_PERF Schema

In the **DBE\_PERF** schema, views are used to diagnose performance issues and are also the data source of WDR snapshots. After the database is installed, only the initial user has the permission for the **DBE\_PERF** schema by default. If the database is upgraded from an earlier version, permissions for the **DBE\_PERF**

schema are the same as those of the earlier version to ensure forward compatibility. Organization views are divided based on multiple dimensions, such as OS, instance, and memory. These views comply with the following naming rules:

- Views starting with **GLOBAL\_**: Request data from database nodes and return the data without processing them.
- Views starting with **SUMMARY\_**: Summarize data in the database . In most cases, data from database nodes (sometimes only the primary database node) is processed and collected.
- Views that do not start with **GLOBAL\_** or **SUMMARY\_**: Local views that do not request data from other database nodes.

## 13.2.1 OS

### 13.2.1.1 OS\_RUNTIME

**OS\_RUNTIME** displays the running status of the current OS.

**Table 13-8** OS\_RUNTIME columns

| Name       | Type    | Description                                                         |
|------------|---------|---------------------------------------------------------------------|
| id         | integer | ID                                                                  |
| name       | text    | Name of the OS running status                                       |
| value      | numeric | Value of the OS running status                                      |
| comments   | text    | Remarks of the OS running status                                    |
| cumulative | boolean | Specifies whether the value of the OS running status is cumulative. |

### 13.2.1.2 GLOBAL\_OS\_RUNTIME

**GLOBAL\_OS\_RUNTIME** provides OS running status information about all normal nodes in the database.

**Table 13-9** GLOBAL\_OS\_RUNTIME columns

| Name      | Type    | Description                      |
|-----------|---------|----------------------------------|
| node_name | name    | Node name.                       |
| id        | integer | ID.                              |
| name      | text    | Name of the OS running status.   |
| value     | numeric | Value of the OS running status.  |
| comments  | text    | Remarks of the OS running status |

| Name       | Type    | Description                                                         |
|------------|---------|---------------------------------------------------------------------|
| cumulative | boolean | Specifies whether the value of the OS running status is cumulative. |

### 13.2.1.3 OS\_THREADS

Provides information about the status of all threads under the current node.

**Table 13-10** OS\_THREADS columns

| Name          | Type                     | Description                                                 |
|---------------|--------------------------|-------------------------------------------------------------|
| node_name     | text                     | Database process name                                       |
| pid           | bigint                   | ID of the thread running under the current database process |
| lwpid         | integer                  | ID of the lightweight thread corresponding to <b>pid</b> .  |
| thread_name   | text                     | Name of the thread corresponding to <b>pid</b> .            |
| creation_time | timestamp with time zone | Creation time of the thread corresponding to <b>pid</b> .   |

### 13.2.1.4 GLOBAL\_OS\_THREADS

GLOBAL\_OS\_THREADS provides status information about threads on all normal nodes in the database.

**Table 13-11** GLOBAL\_OS\_THREADS columns

| Name          | Type                     | Description                                             |
|---------------|--------------------------|---------------------------------------------------------|
| node_name     | text                     | Node name                                               |
| pid           | bigint                   | ID of the thread running under the current node process |
| lwpid         | integer                  | Lightweight thread ID corresponding to <b>pid</b>       |
| thread_name   | text                     | Name of the thread corresponding to <b>pid</b>          |
| creation_time | timestamp with time zone | Creation time of the thread corresponding to <b>pid</b> |

### 13.2.1.5 NODE\_NAME

Provide the names of all normal nodes in the database.

**Table 13-12** NODE\_NAME columns

| Name      | Type | Description |
|-----------|------|-------------|
| node_name | name | Node name   |

## 13.2.2 Instance

### 13.2.2.1 INSTANCE\_TIME

**INSTANCE\_TIME** records time consumption information on the current database node. The time consumption information is classified into the following types:

- DB\_TIME: effective time spent by jobs in multi-core scenarios
- CPU\_TIME: CPU time cost.
- EXECUTION\_TIME: time spent in the executor.
- PARSE\_TIME: time spent on parsing SQL statements
- PLAN\_TIME: time spent on generating plans
- REWRITE\_TIME: time spent on rewriting SQL statements
- PL\_EXECUTION\_TIME: execution time of the PL/SQL stored procedure.
- PL\_COMPILATION\_TIME: compilation time of the PL/SQL stored procedure.
- NET\_SEND\_TIME: time spent on the network
- DATA\_IO\_TIME: time spent on I/Os.

**Table 13-13** INSTANCE\_TIME columns

| Name      | Type    | Description                |
|-----------|---------|----------------------------|
| stat_id   | integer | Statistics ID              |
| stat_name | text    | Type name                  |
| value     | bigint  | Time value (unit: $\mu$ s) |

### 13.2.2.2 GLOBAL\_INSTANCE\_TIME

**GLOBAL\_INSTANCE\_TIME** provides time consumption information about all normal nodes in the database. For details about the time types, see the **INSTANCE\_TIME** view.

**Table 13-14** GLOBAL\_INSTANCE\_TIME columns

| Name      | Type    | Description              |
|-----------|---------|--------------------------|
| node_name | name    | Node name                |
| stat_id   | integer | Statistics ID            |
| stat_name | text    | Type name                |
| value     | bigint  | Duration (unit: $\mu$ s) |

## 13.2.3 Memory

### 13.2.3.1 GS\_SHARED\_MEMORY\_DETAIL

Queries the usage information about shared memory contexts on the current node.

**Table 13-15** GS\_SHARED\_MEMORY\_DETAIL columns

| Name        | Type     | Description                                       |
|-------------|----------|---------------------------------------------------|
| contextname | text     | Name of the memory context.                       |
| level       | smallint | Level of the memory context.                      |
| parent      | text     | Name of the parent memory context.                |
| totalsize   | bigint   | Total size of the shared memory (unit: byte).     |
| freesize    | bigint   | Remaining size of the shared memory (unit: byte). |
| usedsize    | bigint   | Used size of the shared memory (unit: byte).      |

### 13.2.3.2 GLOBAL\_MEMORY\_NODE\_DETAIL

GLOBAL\_MEMORY\_NODE\_DETAIL displays the memory usage of all normal nodes in the database.

**Table 13-16** GLOBAL\_MEMORY\_NODE\_DETAIL columns

| Name     | Type | Description |
|----------|------|-------------|
| nodename | text | Node name.  |

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| memorytype  | text    | <p>Memory name.</p> <ul style="list-style-type: none"><li>• <b>max_process_memory</b>: maximum available memory of the database node.</li><li>• <b>process_used_memory</b>: memory occupied by a process.</li><li>• <b>max_dynamic_memory</b>: maximum dynamic memory.</li><li>• <b>dynamic_used_memory</b>: used dynamic memory.</li><li>• <b>dynamic_peak_memory</b>: dynamic memory peak.</li><li>• <b>dynamic_used_shrctx</b>: context of the used dynamic shared memory.</li><li>• <b>dynamic_peak_shrctx</b>: dynamic peak value of the shared memory context.</li><li>• <b>max_shared_memory</b>: maximum shared memory.</li><li>• <b>shared_used_memory</b>: used shared memory.</li><li>• <b>max_sctpcomm_memory</b>: maximum memory allowed for TCP proxy communications.</li><li>• <b>sctpcomm_used_memory</b>: used memory for TCP proxy communications.</li><li>• <b>sctpcomm_peak_memory</b>: memory peak of TCP proxy communications.</li><li>• <b>other_used_memory</b>: other used memory.</li><li>• <b>gpu_max_dynamic_memory</b>: maximum dynamic GPU memory.</li><li>• <b>gpu_dynamic_used_memory</b>: used dynamic GPU memory.</li><li>• <b>gpu_dynamic_peak_memory</b>: dynamic GPU memory peak.</li><li>• <b>pooler_conn_memory</b>: memory allocated to the connection pool.</li><li>• <b>pooler_freeconn_memory</b>: memory occupied by idle connections in the connection pool.</li><li>• <b>storage_compress_memory</b>: memory used by the storage module for compression.</li><li>• <b>udf_reserved_memory</b>: reserved memory for the user-defined functions.</li></ul> |
| memorybytes | integer | Size of the used memory (unit: MB).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

### 13.2.3.3 GLOBAL\_SHARED\_MEMORY\_DETAIL

GLOBAL\_SHARED\_MEMORY\_DETAIL is used to query the usage of shared memory contexts on all normal nodes in the database.

**Table 13-17** GLOBAL\_SHARED\_MEMORY\_DETAIL columns

| Name        | Type     | Description                                      |
|-------------|----------|--------------------------------------------------|
| node_name   | name     | Node name                                        |
| contextname | text     | Name of the memory context                       |
| level       | smallint | Level of the memory context                      |
| parent      | text     | Name of the parent memory context                |
| totalsize   | bigint   | Total size of the shared memory (unit: byte)     |
| freesize    | bigint   | Remaining size of the shared memory (unit: byte) |
| usedsize    | bigint   | Used size of the shared memory (unit: byte)      |

### 13.2.3.4 MEMORY\_NODE\_DETAIL

MEMORY\_NODE\_DETAIL displays memory usage of a node in the database.

**Table 13-18** MEMORY\_NODE\_DETAIL columns

| Name     | Type | Description |
|----------|------|-------------|
| nodename | text | Node name.  |



| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| memorytype  | text    | <p>Memory name</p> <ul style="list-style-type: none"><li>• <b>max_process_memory</b>: memory occupied by the GaussDB instance.</li><li>• <b>process_used_memory</b>: memory occupied by a process.</li><li>• <b>max_dynamic_memory</b>: maximum dynamic memory.</li><li>• <b>dynamic_used_memory</b>: used dynamic memory.</li><li>• <b>dynamic_peak_memory</b>: dynamic memory peak.</li><li>• <b>dynamic_used_shrctx</b>: maximum dynamic shared memory context.</li><li>• <b>dynamic_peak_shrctx</b>: dynamic peak value of the shared memory context.</li><li>• <b>max_shared_memory</b>: maximum shared memory.</li><li>• <b>shared_used_memory</b>: used shared memory.</li><li>• <b>max_sctpcomm_memory</b>: maximum memory allowed for TCP proxy communications.</li><li>• <b>sctpcomm_used_memory</b>: used memory for TCP proxy communications.</li><li>• <b>sctpcomm_peak_memory</b>: memory peak of TCP proxy communications.</li><li>• <b>other_used_memory</b>: other used memory.</li><li>• <b>gpu_max_dynamic_memory</b>: maximum dynamic GPU memory.</li><li>• <b>gpu_dynamic_used_memory</b>: used dynamic GPU memory.</li><li>• <b>gpu_dynamic_peak_memory</b>: dynamic GPU memory peak.</li><li>• <b>pooler_conn_memory</b>: memory allocated to the connection pool.</li><li>• <b>pooler_freeconn_memory</b>: memory occupied by idle connections in the connection pool.</li><li>• <b>storage_compress_memory</b>: memory used by the storage module for compression.</li><li>• <b>udf_reserved_memory</b>: reserved memory for the user-defined functions.</li></ul> |
| memorybytes | integer | Size of the used memory (unit: MB).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

### 13.2.3.5 SHARED\_MEMORY\_DETAIL

**SHARED\_MEMORY\_DETAIL** queries the usage information about shared memory contexts on the current node.

**Table 13-19 Table 1** SHARED\_MEMORY\_DETAIL columns

| Name        | Type     | Description                                      |
|-------------|----------|--------------------------------------------------|
| contextname | text     | Name of the memory context                       |
| level       | smallint | Level of the memory context                      |
| parent      | text     | Name of the parent memory context                |
| totalsize   | bigint   | Total size of the shared memory (unit: byte)     |
| freesize    | bigint   | Remaining size of the shared memory (unit: byte) |
| usedsize    | bigint   | Used size of the shared memory (unit: byte)      |

## 13.2.4 File

### 13.2.4.1 FILE\_IOSTAT

Records statistics about data file I/Os to indicate I/O performance and detect performance problems such as abnormal I/O operations.

**Table 13-20** FILE\_IOSTAT columns

| Name      | Type   | Description                                          |
|-----------|--------|------------------------------------------------------|
| filenum   | oid    | File identifier                                      |
| dbid      | oid    | Database ID                                          |
| spcid     | oid    | Tablespace ID                                        |
| phyrds    | bigint | Number of times of reading physical files            |
| phywrts   | bigint | Number of times of writing into physical files       |
| phyblkrd  | bigint | Number of times of reading physical file blocks      |
| phyblkwrt | bigint | Number of times of writing into physical file blocks |

| Name     | Type   | Description                                             |
|----------|--------|---------------------------------------------------------|
| readtim  | bigint | Total duration of reading (unit: $\mu$ s)               |
| writetim | bigint | Total duration of writing (unit: $\mu$ s)               |
| avgiotim | bigint | Average duration of reading and writing (unit: $\mu$ s) |
| lstiotim | bigint | Duration of the last file reading (unit: $\mu$ s)       |
| miniotim | bigint | Minimum duration of reading and writing (unit: $\mu$ s) |
| maxiowtm | bigint | Maximum duration of reading and writing (unit: $\mu$ s) |

### 13.2.4.2 SUMMARY\_FILE\_IOSTAT

Records statistics about data file I/Os in the database to indicate I/O performance and detect performance problems such as abnormal I/O operations.

**Table 13-21** SUMMARY\_FILE\_IOSTAT columns

| Name      | Type    | Description                                              |
|-----------|---------|----------------------------------------------------------|
| filenum   | oid     | File ID.                                                 |
| dbid      | oid     | Database ID.                                             |
| spcid     | oid     | Tablespace ID.                                           |
| phyrds    | numeric | Number of times of reading physical files.               |
| phywrts   | numeric | Number of times of writing into physical files.          |
| phyblkrd  | numeric | Number of times of reading physical file blocks.         |
| phyblkwrt | numeric | Number of times of writing into physical file blocks.    |
| readtim   | numeric | Total duration of reading (unit: $\mu$ s).               |
| writetim  | numeric | Total duration of writing (unit: $\mu$ s).               |
| avgiotim  | bigint  | Average duration of reading and writing (unit: $\mu$ s). |
| lstiotim  | bigint  | Duration of the last file reading (unit: $\mu$ s).       |

| Name     | Type   | Description                                              |
|----------|--------|----------------------------------------------------------|
| miniotim | bigint | Minimum duration of reading and writing (unit: $\mu$ s). |
| maxiowtm | bigint | Maximum duration of reading and writing (unit: $\mu$ s). |

### 13.2.4.3 GLOBAL\_FILE\_IOSTAT

GLOBAL\_FILE\_IOSTAT displays statistics about data file I/Os on all nodes.

**Table 13-22** GLOBAL\_FILE\_IOSTAT columns

| Name      | Type   | Description                                              |
|-----------|--------|----------------------------------------------------------|
| node_name | name   | Node name.                                               |
| filenum   | oid    | File ID.                                                 |
| dbid      | oid    | Database ID.                                             |
| spcid     | oid    | Tablespace ID.                                           |
| phyrds    | bigint | Number of times of reading physical files.               |
| phywrts   | bigint | Number of times of writing into physical files.          |
| phyblkrd  | bigint | Number of times of reading physical file blocks.         |
| phyblkwrt | bigint | Number of times of writing into physical file blocks.    |
| readtim   | bigint | Total duration of reading (unit: $\mu$ s).               |
| writetim  | bigint | Total duration of writing (unit: $\mu$ s).               |
| avgiotim  | bigint | Average duration of reading and writing (unit: $\mu$ s). |
| lstiotim  | bigint | Duration of the last file reading (unit: $\mu$ s).       |
| miniotim  | bigint | Minimum duration of reading and writing (unit: $\mu$ s). |
| maxiowtm  | bigint | Maximum duration of reading and writing (unit: $\mu$ s). |

### 13.2.4.4 FILE\_REDO\_IOSTAT

**FILE\_REDO\_IOSTAT** records statistics about redo logs (WALs) on the current node.

**Table 13-23** FILE\_REDO\_IOSTAT columns

| Name      | Type   | Description                                                                                                     |
|-----------|--------|-----------------------------------------------------------------------------------------------------------------|
| phywrts   | bigint | Number of times writing into the WAL buffer                                                                     |
| phyblkwrt | bigint | Number of blocks written into the WAL buffer                                                                    |
| writetim  | bigint | Duration of writing into Xlog files (unit: $\mu$ s)                                                             |
| avgiotim  | bigint | Average duration of writing into Xlog files (unit: $\mu$ s). <b>avgiotim</b> = <b>writetim</b> / <b>phywrts</b> |
| lstiotim  | bigint | Duration of the last writing into Xlog files (unit: $\mu$ s)                                                    |
| miniotim  | bigint | Minimum duration of writing into Xlog files (unit: $\mu$ s)                                                     |
| maxiowtm  | bigint | Maximum duration of writing into Xlog files (unit: $\mu$ s)                                                     |

### 13.2.4.5 SUMMARY\_FILE\_REDO\_IOSTAT

**SUMMARY\_FILE\_REDO\_IOSTAT** records statistics about all redo logs and WALs in the database.

**Table 13-24** SUMMARY\_FILE\_REDO\_IOSTAT columns

| Name      | Type    | Description                                                                                                           |
|-----------|---------|-----------------------------------------------------------------------------------------------------------------------|
| phywrts   | numeric | Number of times writing into the WAL buffer.                                                                          |
| phyblkwrt | numeric | Number of blocks written into the WAL buffer.                                                                         |
| writetim  | numeric | Duration of writing into Xlog files (unit: $\mu$ s).                                                                  |
| avgiotim  | bigint  | Average duration of writing into Xlog files (unit: $\mu$ s). The value is <b>writetim</b> divided by <b>phywrts</b> . |
| lstiotim  | bigint  | Duration of the last writing into Xlog files (unit: $\mu$ s).                                                         |

| Name     | Type   | Description                                                  |
|----------|--------|--------------------------------------------------------------|
| miniotim | bigint | Minimum duration of writing into Xlog files (unit: $\mu$ s). |
| maxiowtm | bigint | Maximum duration of writing into Xlog files (unit: $\mu$ s). |

#### 13.2.4.6 GLOBAL\_FILE\_REDO\_IOSTAT

**GLOBAL\_FILE\_REDO\_IOSTAT** displays statistics about redo logs (WALs) on nodes in the database.

**Table 13-25** GLOBAL\_FILE\_REDO\_IOSTAT columns

| Name      | Type   | Description                                                                                     |
|-----------|--------|-------------------------------------------------------------------------------------------------|
| node_name | name   | Node name                                                                                       |
| phywrts   | bigint | Number of times writing into the WAL buffer                                                     |
| phyblkwrt | bigint | Number of blocks written into the WAL buffer                                                    |
| writetim  | bigint | Duration of writing into Xlog files (unit: $\mu$ s)                                             |
| avgiotim  | bigint | Average duration of writing into Xlog files (unit: $\mu$ s). <b>avgiotim = writetim/phywrts</b> |
| lstiotim  | bigint | Duration of the last writing into Xlog files (unit: $\mu$ s)                                    |
| miniotim  | bigint | Minimum duration of writing into Xlog files (unit: $\mu$ s)                                     |
| maxiowtm  | bigint | Maximum duration of writing into Xlog files (unit: $\mu$ s)                                     |

#### 13.2.4.7 LOCAL\_REL\_IOSTAT

Displays the accumulated I/O status of all data files on the current node.

**Table 13-26** LOCAL\_REL\_IOSTAT columns

| Name   | Type   | Description                               |
|--------|--------|-------------------------------------------|
| phyrds | bigint | Number of times of reading physical files |

| Name      | Type   | Description                                          |
|-----------|--------|------------------------------------------------------|
| phywrts   | bigint | Number of times of writing into physical files       |
| phyblkrd  | bigint | Number of times of reading physical file blocks      |
| phyblkwrt | bigint | Number of times of writing into physical file blocks |

#### 13.2.4.8 GLOBAL\_REL\_IOSTAT

Displays statistics about data file I/Os on all nodes.

**Table 13-27** GLOBAL\_REL\_IOSTAT columns

| Name      | Type   | Description                                          |
|-----------|--------|------------------------------------------------------|
| node_name | name   | Node name                                            |
| phyrds    | bigint | Number of times of reading physical files            |
| phywrts   | bigint | Number of times of writing into physical files       |
| phyblkrd  | bigint | Number of times of reading physical file blocks      |
| phyblkwrt | bigint | Number of times of writing into physical file blocks |

#### 13.2.4.9 SUMMARY\_REL\_IOSTAT

Obtains statistics about data file I/Os on all nodes.

**Table 13-28** SUMMARY\_REL\_IOSTAT columns

| Name      | Type    | Description                                           |
|-----------|---------|-------------------------------------------------------|
| phyrds    | numeric | Number of times of reading physical files.            |
| phywrts   | numeric | Number of times of writing into physical files.       |
| phyblkrd  | numeric | Number of times of reading physical file blocks.      |
| phyblkwrt | numeric | Number of times of writing into physical file blocks. |

## 13.2.5 Object

### 13.2.5.1 STAT\_USER\_TABLES

**STAT\_USER\_TABLES** displays the status information about user-defined ordinary tables in all namespaces on the current node.

**Table 13-29** STAT\_USER\_TABLES columns

| Name            | Type                     | Description                                                                            |
|-----------------|--------------------------|----------------------------------------------------------------------------------------|
| relid           | oid                      | Table OID                                                                              |
| schemaname      | name                     | Name of the schema that the table is in                                                |
| relname         | name                     | Table name                                                                             |
| seq_scan        | bigint                   | Number of sequential scans initiated on the table                                      |
| seq_tup_read    | bigint                   | Number of live rows fetched by sequential scans                                        |
| idx_scan        | bigint                   | Number of index scans initiated on the table                                           |
| idx_tup_fetch   | bigint                   | Number of live rows fetched by index scans                                             |
| n_tup_ins       | bigint                   | Number of rows inserted                                                                |
| n_tup_upd       | bigint                   | Number of rows updated                                                                 |
| n_tup_del       | bigint                   | Number of rows deleted                                                                 |
| n_tup_hot_upd   | bigint                   | Number of rows HOT updated (with no separate index update required)                    |
| n_live_tup      | bigint                   | Estimated number of live rows                                                          |
| n_dead_tup      | bigint                   | Estimated number of dead rows                                                          |
| last_vacuum     | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| last_autovacuum | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                    |
| last_analyze    | timestamp with time zone | Last time at which the table was manually analyzed                                     |



| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times this table has been vacuumed by the autovacuum daemon                   |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon                    |

### 13.2.5.2 SUMMARY\_STAT\_USER\_TABLES

SUMMARY\_STAT\_USER\_TABLES displays the status information about user-defined ordinary tables in all namespaces in the database.

**Table 13-30** SUMMARY\_STAT\_USER\_TABLES columns

| Name          | Type    | Description                                                         |
|---------------|---------|---------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that contains the table                          |
| relname       | name    | Table name                                                          |
| seq_scan      | numeric | Number of sequential scans initiated on this table                  |
| seq_tup_read  | numeric | Number of live rows fetched by sequential scans                     |
| idx_scan      | numeric | Number of index scans initiated on the table                        |
| idx_tup_fetch | numeric | Number of live rows fetched by index scans                          |
| n_tup_ins     | numeric | Number of rows inserted                                             |
| n_tup_upd     | numeric | Number of rows updated                                              |
| n_tup_del     | numeric | Number of rows deleted                                              |
| n_tup_hot_upd | numeric | Number of rows HOT updated (with no separate index update required) |
| n_live_tup    | numeric | Estimated number of live rows                                       |

| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| n_dead_tup        | numeric                  | Estimated number of dead rows                                                           |
| last_vacuum       | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )  |
| last_autovacuum   | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                     |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | numeric                  | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | numeric                  | Number of times the table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | numeric                  | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | numeric                  | Number of times the table has been analyzed by the autovacuum daemon                    |

### 13.2.5.3 GLOBAL\_STAT\_USER\_TABLES

GLOBAL\_STAT\_USER\_TABLES displays the status information about user-defined ordinary tables in all namespaces on each node.

**Table 13-31** GLOBAL\_STAT\_USER\_TABLES columns

| Name       | Type   | Description                                        |
|------------|--------|----------------------------------------------------|
| node_name  | name   | Node name                                          |
| relid      | oid    | OID of the table                                   |
| schemaname | name   | Name of the schema that this table is in.          |
| relname    | name   | Table name                                         |
| seq_scan   | bigint | Number of sequential scans initiated on this table |

| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| seq_tup_read      | bigint                   | Number of live rows fetched by sequential scans                                         |
| idx_scan          | bigint                   | Number of index scans initiated on the table                                            |
| idx_tup_fetch     | bigint                   | Number of live rows fetched by index scans                                              |
| n_tup_ins         | bigint                   | Number of rows inserted                                                                 |
| n_tup_upd         | bigint                   | Number of rows updated                                                                  |
| n_tup_del         | bigint                   | Number of rows deleted                                                                  |
| n_tup_hot_upd     | bigint                   | Number of rows HOT updated (with no separate index update required).                    |
| n_live_tup        | bigint                   | Estimated number of live rows                                                           |
| n_dead_tup        | bigint                   | Estimated number of dead rows.                                                          |
| last_vacuum       | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )  |
| last_autovacuum   | timestamp with time zone | Last time when the table was vacuumed by the AUTOVACUUM daemon.                         |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze  | timestamp with time zone | Last time when the table was analyzed by the AUTOVACUUM daemon.                         |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times the table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the AUTOVACUUM daemon.                   |

### 13.2.5.4 STAT\_USER\_INDEXES

**STAT\_USER\_INDEXES** displays the status information about the index of user-defined ordinary tables in the current database.

**Table 13-32** STAT\_USER\_INDEXES columns

| Name          | Type   | Description                                                             |
|---------------|--------|-------------------------------------------------------------------------|
| relid         | oid    | OID of the table for this index                                         |
| indexrelid    | oid    | OID of the index                                                        |
| schemaname    | name   | Name of the schema that the index is in                                 |
| relname       | name   | Name of the table for the index                                         |
| indexrelname  | name   | Index name                                                              |
| idx_scan      | bigint | Number of index scans initiated on the index                            |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index                  |
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index |

### 13.2.5.5 SUMMARY\_STAT\_USER\_INDEXES

**SUMMARY\_STAT\_USER\_INDEXES** displays the status information about the index of user-defined ordinary tables in all databases in the database.

**Table 13-33** SUMMARY\_STAT\_USER\_INDEXES columns

| Name          | Type    | Description                                                              |
|---------------|---------|--------------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in.                                 |
| relname       | name    | Name of the table for the index                                          |
| indexrelname  | name    | Index name                                                               |
| idx_scan      | numeric | Number of index scans initiated on the index.                            |
| idx_tup_read  | numeric | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | numeric | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.6 GLOBAL\_STAT\_USER\_INDEXES

GLOBAL\_STAT\_USER\_INDEXES displays the status information about the index of user-defined ordinary tables in all databases of each node.

**Table 13-34** GLOBAL\_STAT\_USER\_INDEXES columns

| Name          | Type   | Description                                                              |
|---------------|--------|--------------------------------------------------------------------------|
| node_name     | name   | Node name.                                                               |
| relid         | oid    | OID of the table for this index.                                         |
| indexrelid    | oid    | OID of the index.                                                        |
| schemaname    | name   | Name of the schema that the index is in.                                 |
| relname       | name   | Name of the table for the index.                                         |
| indexrelname  | name   | Index name.                                                              |
| idx_scan      | bigint | Number of index scans initiated on the index.                            |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.7 STAT\_SYS\_TABLES

STAT\_SYS\_TABLES displays statistics about all the system catalogs in the **pg\_catalog**, **information\_schema**, and **pg\_toast** schemas on a single node.

**Table 13-35** STAT\_SYS\_TABLES columns

| Name         | Type   | Description                                       |
|--------------|--------|---------------------------------------------------|
| relid        | oid    | Table OID                                         |
| schemaname   | name   | Name of the schema that the table is in           |
| relname      | name   | Table name                                        |
| seq_scan     | bigint | Number of sequential scans initiated on the table |
| seq_tup_read | bigint | Number of live rows fetched by sequential scans   |
| idx_scan     | bigint | Number of index scans initiated on the table      |

| Name              | Type                     | Description                                                                              |
|-------------------|--------------------------|------------------------------------------------------------------------------------------|
| idx_tup_fetch     | bigint                   | Number of live rows fetched by index scans                                               |
| n_tup_ins         | bigint                   | Number of rows inserted                                                                  |
| n_tup_upd         | bigint                   | Number of rows updated                                                                   |
| n_tup_del         | bigint                   | Number of rows deleted                                                                   |
| n_tup_hot_upd     | bigint                   | Number of rows HOT updated (with no separate index update required)                      |
| n_live_tup        | bigint                   | Estimated number of live rows                                                            |
| n_dead_tup        | bigint                   | Estimated number of dead rows                                                            |
| last_vacuum       | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )   |
| last_autovacuum   | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                      |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                      |
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                      |
| vacuum_count      | bigint                   | Number of times this table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times this table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | bigint                   | Number of times this table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times this table has been analyzed by the autovacuum daemon                    |

### 13.2.5.8 SUMMARY\_STAT\_SYS\_TABLES

SUMMARY\_STAT\_SYS\_TABLES displays statistics about all the system catalogs in the pg\_catalog, information\_schema, and pg\_toast schemas in the database.

**Table 13-36** SUMMARY\_STAT\_SYS\_TABLES columns

| Name             | Type                     | Description                                                                             |
|------------------|--------------------------|-----------------------------------------------------------------------------------------|
| schemaname       | name                     | Name of the schema that contains the table                                              |
| relname          | name                     | Table name                                                                              |
| seq_scan         | numeric                  | Number of sequential scans initiated on this table                                      |
| seq_tup_read     | numeric                  | Number of live rows fetched by sequential scans                                         |
| idx_scan         | numeric                  | Number of index scans initiated on the table                                            |
| idx_tup_fetch    | numeric                  | Number of live rows fetched by index scans                                              |
| n_tup_ins        | numeric                  | Number of rows inserted                                                                 |
| n_tup_upd        | numeric                  | Number of rows updated                                                                  |
| n_tup_del        | numeric                  | Number of rows deleted                                                                  |
| n_tup_hot_upd    | numeric                  | Number of rows HOT updated (with no separate index update required)                     |
| n_live_tup       | numeric                  | Estimated number of live rows                                                           |
| n_dead_tup       | numeric                  | Estimated number of dead rows                                                           |
| last_vacuum      | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )  |
| last_autovacuum  | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                     |
| last_analyze     | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count     | numeric                  | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count | numeric                  | Number of times the table has been vacuumed by the autovacuum daemon                    |

| Name              | Type    | Description                                                          |
|-------------------|---------|----------------------------------------------------------------------|
| analyze_count     | numeric | Number of times the table has been manually analyzed                 |
| autoanalyze_count | numeric | Number of times the table has been analyzed by the autovacuum daemon |

### 13.2.5.9 GLOBAL\_STAT\_SYS\_TABLES

**GLOBAL\_STAT\_SYS\_TABLES** displays statistics about all the system catalogs in the **pg\_catalog**, **information\_schema**, and **pg\_toast** schemas on each node.

**Table 13-37** GLOBAL\_STAT\_SYS\_TABLES columns

| Name          | Type                     | Description                                                                            |
|---------------|--------------------------|----------------------------------------------------------------------------------------|
| node_name     | name                     | Node name                                                                              |
| relid         | oid                      | Table OID                                                                              |
| schemaname    | name                     | Name of the schema that contains the table                                             |
| relname       | name                     | Table name                                                                             |
| seq_scan      | bigint                   | Number of sequential scans initiated on this table                                     |
| seq_tup_read  | bigint                   | Number of live rows fetched by sequential scans                                        |
| idx_scan      | bigint                   | Number of index scans initiated on the table                                           |
| idx_tup_fetch | bigint                   | Number of live rows fetched by index scans                                             |
| n_tup_ins     | bigint                   | Number of rows inserted                                                                |
| n_tup_upd     | bigint                   | Number of rows updated                                                                 |
| n_tup_del     | bigint                   | Number of rows deleted                                                                 |
| n_tup_hot_upd | bigint                   | Number of rows HOT updated (with no separate index update required)                    |
| n_live_tup    | bigint                   | Estimated number of live rows                                                          |
| n_dead_tup    | bigint                   | Estimated number of dead rows                                                          |
| last_vacuum   | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> ) |



| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| last_autovacuum   | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                     |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times the table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon                    |

### 13.2.5.10 STAT\_SYS\_INDEXES

**STAT\_SYS\_INDEXES** displays index status information about all the system catalogs in the **pg\_catalog**, **information\_schema**, and **pg\_toast** schemas.

**Table 13-38** STAT\_SYS\_INDEXES columns

| Name          | Type   | Description                                                             |
|---------------|--------|-------------------------------------------------------------------------|
| relid         | oid    | OID of the table for this index                                         |
| indexrelid    | oid    | OID of the index                                                        |
| schemaname    | name   | Name of the schema that the index is in                                 |
| relname       | name   | Name of the table that the index is created for                         |
| indexrelname  | name   | Index name                                                              |
| idx_scan      | bigint | Number of index scans initiated on the index                            |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index                  |
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index |

### 13.2.5.11 SUMMARY\_STAT\_SYS\_INDEXES

SUMMARY\_STAT\_SYS\_INDEXES displays index status information about all the system catalogs in the pg\_catalog, information\_schema, and pg\_toast schemas in the database.

**Table 13-39** SUMMARY\_STAT\_SYS\_INDEXES columns

| Name          | Type    | Description                                                              |
|---------------|---------|--------------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in.                                 |
| relname       | name    | Name of the table for the index                                          |
| indexrelname  | name    | Index name                                                               |
| idx_scan      | numeric | Number of index scans initiated on the index.                            |
| idx_tup_read  | numeric | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | numeric | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.12 GLOBAL\_STAT\_SYS\_INDEXES

GLOBAL\_STAT\_SYS\_INDEXES displays index status information about all the system catalogs in the pg\_catalog, information\_schema, and pg\_toast schemas on each node.

**Table 13-40** GLOBAL\_STAT\_SYS\_INDEXES columns

| Name         | Type   | Description                                             |
|--------------|--------|---------------------------------------------------------|
| node_name    | name   | Node name.                                              |
| relid        | oid    | OID of the table for this index.                        |
| indexrelid   | oid    | OID of the index.                                       |
| schemaname   | name   | Name of the schema that the index is in.                |
| relname      | name   | Name of the table for the index.                        |
| indexrelname | name   | Index name.                                             |
| idx_scan     | bigint | Number of index scans initiated on the index.           |
| idx_tup_read | bigint | Number of index entries returned by scans on the index. |

| Name          | Type   | Description                                                              |
|---------------|--------|--------------------------------------------------------------------------|
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.13 STAT\_ALL\_TABLES

**STAT\_ALL\_TABLES** displays statistics about one row for each table (including TOAST tables) in databases on this node.

**Table 13-41** STAT\_ALL\_TABLES columns

| Name            | Type                     | Description                                                                            |
|-----------------|--------------------------|----------------------------------------------------------------------------------------|
| relid           | oid                      | Table OID                                                                              |
| schemaname      | name                     | Name of the schema that the table is in                                                |
| relname         | name                     | Table name                                                                             |
| seq_scan        | bigint                   | Number of sequential scans initiated on the table                                      |
| seq_tup_read    | bigint                   | Number of live rows fetched by sequential scans                                        |
| idx_scan        | bigint                   | Number of index scans initiated on the table                                           |
| idx_tup_fetch   | bigint                   | Number of live rows fetched by index scans                                             |
| n_tup_ins       | bigint                   | Number of rows inserted                                                                |
| n_tup_upd       | bigint                   | Number of rows updated                                                                 |
| n_tup_del       | bigint                   | Number of rows deleted                                                                 |
| n_tup_hot_upd   | bigint                   | Number of rows HOT updated (with no separate index update required)                    |
| n_live_tup      | bigint                   | Estimated number of live rows                                                          |
| n_dead_tup      | bigint                   | Estimated number of dead rows                                                          |
| last_vacuum     | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| last_autovacuum | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                    |
| last_analyze    | timestamp with time zone | Last time at which the table was manually analyzed                                     |

| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times this table has been vacuumed by the autovacuum daemon                   |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon                    |

#### 13.2.5.14 SUMMARY\_STAT\_ALL\_TABLES

SUMMARY\_STAT\_ALL\_TABLES displays statistics about one row for each table (including TOAST tables) in a database in the entire database system.

**Table 13-42** SUMMARY\_STAT\_ALL\_TABLES columns

| Name          | Type    | Description                                                         |
|---------------|---------|---------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that contains the table                          |
| relname       | name    | Table name                                                          |
| seq_scan      | numeric | Number of sequential scans initiated on this table                  |
| seq_tup_read  | numeric | Number of live rows fetched by sequential scans                     |
| idx_scan      | numeric | Number of index scans initiated on the table                        |
| idx_tup_fetch | numeric | Number of live rows fetched by index scans                          |
| n_tup_ins     | numeric | Number of rows inserted                                             |
| n_tup_upd     | numeric | Number of rows updated                                              |
| n_tup_del     | numeric | Number of rows deleted                                              |
| n_tup_hot_upd | numeric | Number of rows HOT updated (with no separate index update required) |
| n_live_tup    | numeric | Estimated number of live rows                                       |

| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| n_dead_tup        | numeric                  | Estimated number of dead rows                                                           |
| last_vacuum       | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )  |
| last_autovacuum   | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                     |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | numeric                  | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | numeric                  | Number of times the table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | numeric                  | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | numeric                  | Number of times the table has been analyzed by the autovacuum daemon                    |

### 13.2.5.15 GLOBAL\_STAT\_ALL\_TABLES

**GLOBAL\_STAT\_ALL\_TABLES** displays statistics about one row for each table (including TOAST tables) in databases on each node.

**Table 13-43** GLOBAL\_STAT\_ALL\_TABLES columns

| Name         | Type   | Description                                        |
|--------------|--------|----------------------------------------------------|
| node_name    | name   | Node name                                          |
| relid        | oid    | OID of the table                                   |
| schemaname   | name   | Name of the schema that contains the table         |
| relname      | name   | Table name                                         |
| seq_scan     | bigint | Number of sequential scans initiated on this table |
| seq_tup_read | bigint | Number of live rows fetched by sequential scans    |

| Name              | Type                     | Description                                                                             |
|-------------------|--------------------------|-----------------------------------------------------------------------------------------|
| idx_scan          | bigint                   | Number of index scans initiated on the table                                            |
| idx_tup_fetch     | bigint                   | Number of live rows fetched by index scans                                              |
| n_tup_ins         | bigint                   | Number of rows inserted                                                                 |
| n_tup_upd         | bigint                   | Number of rows updated                                                                  |
| n_tup_del         | bigint                   | Number of rows deleted                                                                  |
| n_tup_hot_upd     | bigint                   | Number of rows HOT updated (with no separate index update required)                     |
| n_live_tup        | bigint                   | Estimated number of live rows                                                           |
| n_dead_tup        | bigint                   | Estimated number of dead rows                                                           |
| last_vacuum       | timestamp with time zone | Last time at which this table was manually vacuumed (not counting <b>VACUUM FULL</b> )  |
| last_autovacuum   | timestamp with time zone | Last time at which this table was vacuumed by the autovacuum daemon                     |
| last_analyze      | timestamp with time zone | Last time at which this table was manually analyzed                                     |
| last_autoanalyze  | timestamp with time zone | Last time at which this table was analyzed by the autovacuum daemon                     |
| vacuum_count      | bigint                   | Number of times the table has been manually vacuumed (not counting <b>VACUUM FULL</b> ) |
| autovacuum_count  | bigint                   | Number of times the table has been vacuumed by the autovacuum daemon                    |
| analyze_count     | bigint                   | Number of times the table has been manually analyzed                                    |
| autoanalyze_count | bigint                   | Number of times the table has been analyzed by the autovacuum daemon                    |

### 13.2.5.16 STAT\_ALL\_INDEXES

STAT\_ALL\_INDEXES contains every row of each index in databases on the current node, showing statistics about accesses to specific indexes.

**Table 13-44** STAT\_ALL\_INDEXES columns

| Name          | Type   | Description                                                              |
|---------------|--------|--------------------------------------------------------------------------|
| relid         | oid    | OID of the table for this index.                                         |
| indexrelid    | oid    | OID of the index.                                                        |
| schemaname    | name   | Name of the schema that the index is in.                                 |
| relname       | name   | Name of the table for the index.                                         |
| indexrelname  | name   | Index name.                                                              |
| idx_scan      | bigint | Number of index scans initiated on the index.                            |
| idx_tup_read  | bigint | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | bigint | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.17 SUMMARY\_STAT\_ALL\_INDEXES

SUMMARY\_STAT\_ALL\_INDEXES displays access statistics about each index in the GaussDB database.

**Table 13-45** SUMMARY\_STAT\_ALL\_INDEXES columns

| Name          | Type    | Description                                                              |
|---------------|---------|--------------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in.                                 |
| relname       | name    | Name of the table for the index.                                         |
| indexrelname  | name    | Index name.                                                              |
| idx_scan      | numeric | Number of index scans initiated on the index.                            |
| idx_tup_read  | numeric | Number of index entries returned by scans on the index.                  |
| idx_tup_fetch | numeric | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.18 GLOBAL\_STAT\_ALL\_INDEXES

Contains every row of each index in databases on each node, showing statistics about accesses to specific indexes.

**Table 13-46** GLOBAL\_STAT\_ALL\_INDEXES columns

| Name            | Type   | Description                                                              |
|-----------------|--------|--------------------------------------------------------------------------|
| node_name       | name   | Node name                                                                |
| relid           | oid    | OID of the table for this index                                          |
| indexrelid      | oid    | OID of the index                                                         |
| schemaname      | name   | Name of the schema that the index is in.                                 |
| relname         | name   | Name of the table for the index                                          |
| indexrelname    | name   | Index name                                                               |
| idx_scan        | bigint | Number of index scans initiated on the index.                            |
| idx_tup_read    | bigint | Number of index entries returned by scans on the index.                  |
| idx_tup_fetched | bigint | Number of live table rows fetched by simple index scans using the index. |

### 13.2.5.19 STAT\_DATABASE

**STAT\_DATABASE** contains statistics for each database on this node.

**Table 13-47** STAT\_DATABASE columns

| Name          | Type    | Description                                                                                                                                                                                                            |
|---------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid         | oid     | OID of a database                                                                                                                                                                                                      |
| datname       | name    | Database name                                                                                                                                                                                                          |
| numbackends   | integer | Number of backends currently connected to this database. This is the only column in this view that returns a value reflecting the current state; all other columns return the accumulated values since the last reset. |
| xact_commit   | bigint  | Number of transactions in this database that have been committed                                                                                                                                                       |
| xact_rollback | bigint  | Number of transactions in this database that have been rolled back                                                                                                                                                     |
| blks_read     | bigint  | Number of disk blocks read in this database                                                                                                                                                                            |
| blks_hit      | bigint  | Number of times disk blocks were found in the buffer cache (unnecessary as the number includes only hits in the database buffer cache)                                                                                 |



| Name           | Type                     | Description                                                                                                                                                                                                                              |
|----------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tup_returned   | bigint                   | Number of rows returned by queries in this database                                                                                                                                                                                      |
| tup_fetched    | bigint                   | Number of rows fetched by queries in this database                                                                                                                                                                                       |
| tup_inserted   | bigint                   | Number of rows inserted by queries in this database                                                                                                                                                                                      |
| tup_updated    | bigint                   | Number of rows updated by queries in this database                                                                                                                                                                                       |
| tup_deleted    | bigint                   | Number of rows deleted by queries in this database                                                                                                                                                                                       |
| conflicts      | bigint                   | Number of queries canceled due to database recovery conflicts (conflicts occurring only on the standby node). For details, see <a href="#">STAT_DATABASE_CONFLICTS</a> .                                                                 |
| temp_files     | bigint                   | Number of temporary files created by queries in this database. All temporary files are counted, regardless of why the temporary file was created (for example, sorting or hashing), and regardless of the <b>log_temp_files</b> setting. |
| temp_bytes     | bigint                   | Total amount of data written to temporary files by queries in this database. All temporary files are counted, regardless of why the temporary file was created, and regardless of the <b>log_temp_files</b> setting.                     |
| deadlocks      | bigint                   | Number of deadlocks detected in this database                                                                                                                                                                                            |
| blk_read_time  | double precision         | Time spent reading data file blocks by backends in this database (unit: ms)                                                                                                                                                              |
| blk_write_time | double precision         | Time spent writing data file blocks by backends in this database (unit: ms)                                                                                                                                                              |
| stats_reset    | timestamp with time zone | Time at which the current statistics were reset                                                                                                                                                                                          |

### 13.2.5.20 SUMMARY\_STAT\_DATABASE

SUMMARY\_STAT\_DATABASE contains every row of each database in the database, showing database-wide statistics.

**Table 13-48** SUMMARY\_STAT\_DATABASE columns

| Name          | Type    | Description                                                                                                                                                                                                                              |
|---------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datname       | name    | Database name.                                                                                                                                                                                                                           |
| numbackends   | bigint  | Number of backends currently connected to this database. This is the only column in this view that returns a value reflecting the current state; all other columns return the accumulated values since the last reset.                   |
| xact_commit   | numeric | Number of transactions in this database that have been committed.                                                                                                                                                                        |
| xact_rollback | numeric | Number of transactions in this database that have been rolled back.                                                                                                                                                                      |
| blks_read     | numeric | Number of disk blocks read in this database.                                                                                                                                                                                             |
| blks_hit      | numeric | Number of times disk blocks were found already in the cache, so that a read was not necessary (this only includes hits in the GaussDB cache, not the OS's file system cache).                                                            |
| tup_returned  | numeric | Number of rows returned by queries in this database.                                                                                                                                                                                     |
| tup_fetched   | numeric | Number of rows fetched by queries in this database.                                                                                                                                                                                      |
| tup_inserted  | bigint  | Number of rows inserted by queries in this database.                                                                                                                                                                                     |
| tup_updated   | bigint  | Number of rows updated by queries in this database.                                                                                                                                                                                      |
| tup_deleted   | bigint  | Number of rows deleted by queries in this database.                                                                                                                                                                                      |
| conflicts     | bigint  | Number of queries canceled due to database recovery conflicts (conflicts occurring only on the standby node). For details, see <a href="#">STAT_DATABASE_CONFLICTS</a> .                                                                 |
| temp_files    | numeric | Number of temporary files created by queries in this database. All temporary files are counted, regardless of why the temporary file was created (for example, sorting or hashing), and regardless of the <b>log_temp_files</b> setting. |

| Name           | Type                     | Description                                                                                                                                                                                                          |
|----------------|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| temp_bytes     | numeric                  | Total amount of data written to temporary files by queries in this database. All temporary files are counted, regardless of why the temporary file was created, and regardless of the <b>log_temp_files</b> setting. |
| deadlocks      | bigint                   | Number of deadlocks detected in this database.                                                                                                                                                                       |
| blk_read_time  | double precision         | Time spent in reading data file blocks by backends in this database (unit: ms).                                                                                                                                      |
| blk_write_time | double precision         | Time spent in writing data file blocks by backends in this database (unit: ms).                                                                                                                                      |
| stats_reset    | timestamp with time zone | Time at which the current statistics were reset.                                                                                                                                                                     |

### 13.2.5.21 GLOBAL\_STAT\_DATABASE

GLOBAL\_STAT\_DATABASE contains every row of each database on each node in the entire database system, showing database-wide statistics.

**Table 13-49** GLOBAL\_STAT\_DATABASE columns

| Name          | Type    | Description                                                                                                                                                                                                            |
|---------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name     | name    | Node name.                                                                                                                                                                                                             |
| datid         | oid     | Database OID.                                                                                                                                                                                                          |
| datname       | name    | Name of the database                                                                                                                                                                                                   |
| numbackends   | integer | Number of backends currently connected to this database. This is the only column in this view that returns a value reflecting the current state; all other columns return the accumulated values since the last reset. |
| xact_commit   | bigint  | Number of transactions in this database that have been committed.                                                                                                                                                      |
| xact_rollback | bigint  | Number of transactions in this database that have been rolled back.                                                                                                                                                    |
| blks_read     | bigint  | Number of disk blocks read in this database.                                                                                                                                                                           |

| Name           | Type                     | Description                                                                                                                                                                                                                              |
|----------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| blks_hit       | bigint                   | Number of times disk blocks were found in the buffer cache (unnecessary as the number includes only hits in the database kernel buffer cache)                                                                                            |
| tup_returned   | bigint                   | Number of rows returned by queries in this database                                                                                                                                                                                      |
| tup_fetched    | bigint                   | Number of rows fetched by queries in this database                                                                                                                                                                                       |
| tup_inserted   | bigint                   | Number of rows inserted by queries in this database                                                                                                                                                                                      |
| tup_updated    | bigint                   | Number of rows updated by queries in this database                                                                                                                                                                                       |
| tup_deleted    | bigint                   | Number of rows deleted by queries in this database                                                                                                                                                                                       |
| conflicts      | bigint                   | Number of queries canceled due to database recovery conflicts (conflicts occurring only on the standby node). For details, see <a href="#">STAT_DATABASE_CONFLICTS</a> .                                                                 |
| temp_files     | bigint                   | Number of temporary files created by queries in this database. All temporary files are counted, regardless of why the temporary file was created (for example, sorting or hashing), and regardless of the <b>log_temp_files</b> setting. |
| temp_bytes     | bigint                   | Total amount of data written to temporary files by queries in this database. All temporary files are counted, regardless of why the temporary file was created, and regardless of the <b>log_temp_files</b> setting.                     |
| deadlocks      | bigint                   | Number of deadlocks detected in this database.                                                                                                                                                                                           |
| blk_read_time  | double precision         | Time spent in reading data file blocks by backends in this database (unit: ms).                                                                                                                                                          |
| blk_write_time | double precision         | Time spent in writing data file blocks by backends in this database (unit: ms).                                                                                                                                                          |
| stats_reset    | timestamp with time zone | Time at which the current statistics were reset.                                                                                                                                                                                         |

### 13.2.5.22 STAT\_DATABASE\_CONFLICTS

**STAT\_DATABASE\_CONFLICTS** displays statistics about database conflicts on the current node.

**Table 13-50** STAT\_DATABASE\_CONFLICTS columns

| Name             | Type   | Description                       |
|------------------|--------|-----------------------------------|
| datid            | oid    | Database ID                       |
| datname          | name   | Database name                     |
| confl_tablespace | bigint | Number of conflicting tablespaces |
| confl_lock       | bigint | Number of conflicting locks       |
| confl_snapshot   | bigint | Number of conflicting snapshots   |
| confl_bufferpin  | bigint | Number of conflicting buffers     |
| confl_deadlock   | bigint | Number of conflicting deadlocks   |

### 13.2.5.23 SUMMARY\_STAT\_DATABASE\_CONFLICTS

**SUMMARY\_STAT\_DATABASE\_CONFLICTS** displays statistics about database conflicts in the database.

**Table 13-51** SUMMARY\_STAT\_DATABASE\_CONFLICTS columns

| Name             | Type   | Description                       |
|------------------|--------|-----------------------------------|
| datname          | name   | Database name                     |
| confl_tablespace | bigint | Number of conflicting tablespaces |
| confl_lock       | bigint | Number of conflicting locks       |
| confl_snapshot   | bigint | Number of conflicting snapshots   |
| confl_bufferpin  | bigint | Number of conflicting buffers     |
| confl_deadlock   | bigint | Number of conflicting deadlocks   |

### 13.2.5.24 GLOBAL\_STAT\_DATABASE\_CONFLICTS

**GLOBAL\_STAT\_DATABASE\_CONFLICTS** displays statistics about database conflicts on each node.

**Table 13-52** GLOBAL\_STAT\_DATABASE\_CONFLICTS columns

| Name             | Type   | Description                       |
|------------------|--------|-----------------------------------|
| node_name        | name   | Node name                         |
| datid            | oid    | Database ID                       |
| datname          | name   | Database name                     |
| confl_tablespace | bigint | Number of conflicting tablespaces |
| confl_lock       | bigint | Number of conflicting locks       |
| confl_snapshot   | bigint | Number of conflicting snapshots   |
| confl_bufferpin  | bigint | Number of conflicting buffers     |
| confl_deadlock   | bigint | Number of conflicting deadlocks   |

### 13.2.5.25 STAT\_XACT\_ALL\_TABLES

**STAT\_XACT\_ALL\_TABLES** displays transaction status information about all ordinary tables and TOAST tables in the current namespace.

**Table 13-53** STAT\_XACT\_ALL\_TABLES columns

| Name            | Type   | Description                                       |
|-----------------|--------|---------------------------------------------------|
| relid           | oid    | Table OID                                         |
| schemaname      | name   | Name of the schema that the table is in           |
| relname         | name   | Table name                                        |
| seq_scan        | bigint | Number of sequential scans initiated on the table |
| seq_tup_read    | bigint | Number of live rows fetched by sequential scans   |
| idx_scan        | bigint | Number of index scans initiated on the table      |
| idx_tup_fetched | bigint | Number of live rows fetched by index scans        |
| n_tup_ins       | bigint | Number of rows inserted                           |
| n_tup_upd       | bigint | Number of rows updated                            |
| n_tup_del       | bigint | Number of rows deleted                            |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.26 SUMMARY\_STAT\_XACT\_ALL\_TABLES

SUMMARY\_STAT\_XACT\_ALL\_TABLES displays transaction status information about all ordinary tables and TOAST tables in all namespaces in the database.

**Table 13-54** SUMMARY\_STAT\_XACT\_ALL\_TABLES columns

| Name          | Type    | Description                                                         |
|---------------|---------|---------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that contains the table                          |
| relname       | name    | Table name.                                                         |
| seq_scan      | numeric | Number of sequential scans initiated on this table.                 |
| seq_tup_read  | numeric | Number of live rows fetched by sequential scans.                    |
| idx_scan      | numeric | Number of index scans initiated on the table                        |
| idx_tup_fetch | numeric | Number of live rows fetched by index scans.                         |
| n_tup_ins     | numeric | Number of rows inserted.                                            |
| n_tup_upd     | numeric | Number of rows updated.                                             |
| n_tup_del     | numeric | Number of rows deleted.                                             |
| n_tup_hot_upd | numeric | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.27 GLOBAL\_STAT\_XACT\_ALL\_TABLES

GLOBAL\_STAT\_XACT\_ALL\_TABLES displays transaction status information about all ordinary tables and TOAST tables in namespaces on each node.

**Table 13-55** GLOBAL\_STAT\_XACT\_ALL\_TABLES columns

| Name       | Type | Description                                |
|------------|------|--------------------------------------------|
| node_name  | name | Node name                                  |
| relid      | oid  | OID of the table                           |
| schemaname | name | Name of the schema that contains the table |
| relname    | name | Table name.                                |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| seq_scan      | bigint | Number of sequential scans initiated on this table.                 |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.                    |
| idx_scan      | bigint | Number of index scans initiated on the table                        |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.                         |
| n_tup_ins     | bigint | Number of rows inserted.                                            |
| n_tup_upd     | bigint | Number of rows updated.                                             |
| n_tup_del     | bigint | Number of rows deleted.                                             |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.28 STAT\_XACT\_SYS\_TABLES

**STAT\_XACT\_SYS\_TABLES** displays transaction status information about the system catalogs in namespaces on the current node.

**Table 13-56** STAT\_XACT\_SYS\_TABLES columns

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| relid         | oid    | Table OID                                                           |
| schemaname    | name   | Name of the schema that the table is in                             |
| relname       | name   | Table name.                                                         |
| seq_scan      | bigint | Number of sequential scans initiated on the table.                  |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.                    |
| idx_scan      | bigint | Number of index scans initiated on the table.                       |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.                         |
| n_tup_ins     | bigint | Number of rows inserted.                                            |
| n_tup_upd     | bigint | Number of rows updated.                                             |
| n_tup_del     | bigint | Number of rows deleted.                                             |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |



### 13.2.5.29 SUMMARY\_STAT\_XACT\_SYS\_TABLES

SUMMARY\_STAT\_XACT\_SYS\_TABLES displays transaction status information about the system catalogs in namespaces in the database.

**Table 13-57** SUMMARY\_STAT\_XACT\_SYS\_TABLES columns

| Name          | Type    | Description                                                         |
|---------------|---------|---------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that contains the table                          |
| relname       | name    | Table name.                                                         |
| seq_scan      | numeric | Number of sequential scans initiated on this table.                 |
| seq_tup_read  | numeric | Number of live rows fetched by sequential scans.                    |
| idx_scan      | numeric | Number of index scans initiated on the table                        |
| idx_tup_fetch | numeric | Number of live rows fetched by index scans.                         |
| n_tup_ins     | numeric | Number of rows inserted.                                            |
| n_tup_upd     | numeric | Number of rows updated.                                             |
| n_tup_del     | numeric | Number of rows deleted.                                             |
| n_tup_hot_upd | numeric | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.30 GLOBAL\_STAT\_XACT\_SYS\_TABLES

GLOBAL\_STAT\_XACT\_SYS\_TABLES displays transaction status information about the system catalogs in namespaces on each node.

**Table 13-58** GLOBAL\_STAT\_XACT\_SYS\_TABLES columns

| Name         | Type   | Description                                         |
|--------------|--------|-----------------------------------------------------|
| node_name    | name   | Node name.                                          |
| relid        | oid    | OID of the table                                    |
| schemaname   | name   | Name of the schema that contains the table          |
| relname      | name   | Table name.                                         |
| seq_scan     | bigint | Number of sequential scans initiated on this table. |
| seq_tup_read | bigint | Number of live rows fetched by sequential scans.    |
| idx_scan     | bigint | Number of index scans initiated on the table        |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.                         |
| n_tup_ins     | bigint | Number of rows inserted.                                            |
| n_tup_upd     | bigint | Number of rows updated.                                             |
| n_tup_del     | bigint | Number of rows deleted.                                             |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.31 STAT\_XACT\_USER\_TABLES

**STAT\_XACT\_USER\_TABLES** displays transaction status information about the user tables in namespaces on the current node.

**Table 13-59** STAT\_XACT\_USER\_TABLES columns

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| relid         | oid    | Table OID                                                           |
| schemaname    | name   | Name of the schema that the table is in                             |
| relname       | name   | Table name.                                                         |
| seq_scan      | bigint | Number of sequential scans initiated on the table.                  |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.                    |
| idx_scan      | bigint | Number of index scans initiated on the table.                       |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.                         |
| n_tup_ins     | bigint | Number of rows inserted.                                            |
| n_tup_upd     | bigint | Number of rows updated.                                             |
| n_tup_del     | bigint | Number of rows deleted.                                             |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.32 SUMMARY\_STAT\_XACT\_USER\_TABLES

**SUMMARY\_STAT\_XACT\_USER\_TABLES** displays the transaction status information about the user tables in namespaces in the database.

**Table 13-60** SUMMARY\_STAT\_XACT\_USER\_TABLES columns

| Name          | Type    | Description                                                         |
|---------------|---------|---------------------------------------------------------------------|
| schemaname    | name    | Name of the schema that contains the table                          |
| relname       | name    | Table name.                                                         |
| seq_scan      | numeric | Number of sequential scans initiated on this table.                 |
| seq_tup_read  | numeric | Number of live rows fetched by sequential scans.                    |
| idx_scan      | numeric | Number of index scans initiated on the table                        |
| idx_tup_fetch | numeric | Number of live rows fetched by index scans.                         |
| n_tup_ins     | numeric | Number of rows inserted.                                            |
| n_tup_upd     | numeric | Number of rows updated.                                             |
| n_tup_del     | numeric | Number of rows deleted.                                             |
| n_tup_hot_upd | numeric | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.33 GLOBAL\_STAT\_XACT\_USER\_TABLES

**GLOBAL\_STAT\_XACT\_USER\_TABLES** displays transaction status information about the user tables in namespaces on each node.

**Table 13-61** GLOBAL\_STAT\_XACT\_USER\_TABLES columns

| Name          | Type   | Description                                         |
|---------------|--------|-----------------------------------------------------|
| node_name     | name   | Node name.                                          |
| relid         | oid    | OID of the table                                    |
| schemaname    | name   | Name of the schema that contains the table          |
| relname       | name   | Table name.                                         |
| seq_scan      | bigint | Number of sequential scans initiated on this table. |
| seq_tup_read  | bigint | Number of live rows fetched by sequential scans.    |
| idx_scan      | bigint | Number of index scans initiated on the table.       |
| idx_tup_fetch | bigint | Number of live rows fetched by index scans.         |
| n_tup_ins     | bigint | Number of rows inserted.                            |
| n_tup_upd     | bigint | Number of rows updated.                             |

| Name          | Type   | Description                                                         |
|---------------|--------|---------------------------------------------------------------------|
| n_tup_del     | bigint | Number of rows deleted.                                             |
| n_tup_hot_upd | bigint | Number of rows HOT updated (with no separate index update required) |

### 13.2.5.34 STAT\_XACT\_USER\_FUNCTIONS

**STAT\_XACT\_USER\_FUNCTIONS** displays statistics about function executions in the current transaction.

**Table 13-62** STAT\_XACT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                               |
|------------|------------------|---------------------------------------------------------------------------|
| funcid     | oid              | OID of a function                                                         |
| schemaname | name             | Schema name                                                               |
| funcname   | name             | Function name.                                                            |
| calls      | bigint           | Number of times the function has been called.                             |
| total_time | double precision | Total time spent in the function and all other functions called by it     |
| self_time  | double precision | Total time spent in this function, excluding other functions called by it |

### 13.2.5.35 SUMMARY\_STAT\_XACT\_USER\_FUNCTIONS

**SUMMARY\_STAT\_XACT\_USER\_FUNCTIONS** displays statistics about function executions in transactions in the database.

**Table 13-63** SUMMARY\_STAT\_XACT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                           |
|------------|------------------|-----------------------------------------------------------------------|
| schemaname | name             | Schema name                                                           |
| funcname   | name             | Function name.                                                        |
| calls      | numeric          | Number of times that the function has been called.                    |
| total_time | double precision | Total time spent in the function and all other functions called by it |

| Name      | Type             | Description                                                                     |
|-----------|------------------|---------------------------------------------------------------------------------|
| self_time | double precision | Total time spent in the function itself, excluding other functions called by it |

### 13.2.5.36 GLOBAL\_STAT\_XACT\_USER\_FUNCTIONS

**GLOBAL\_STAT\_XACT\_USER\_FUNCTIONS** displays statistics about function executions in transactions on each node.

**Table 13-64** GLOBAL\_STAT\_XACT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                                |
|------------|------------------|----------------------------------------------------------------------------|
| node_name  | name             | Node name.                                                                 |
| funcid     | oid              | OID of a function.                                                         |
| schemaname | name             | Schema name                                                                |
| funcname   | name             | Function name.                                                             |
| calls      | bigint           | Number of times that the function has been called.                         |
| total_time | double precision | Total time spent in this function and all other functions called by it.    |
| self_time  | double precision | Total time spent in this function, excluding other functions called by it. |

### 13.2.5.37 STAT\_BAD\_BLOCK

**STAT\_BAD\_BLOCK** displays information about table and index read failures on the current node.

**Table 13-65** STAT\_BAD\_BLOCK columns

| Name         | Type     | Description                             |
|--------------|----------|-----------------------------------------|
| nodename     | text     | Node name.                              |
| databaseid   | integer  | OID of the database                     |
| tablespaceid | integer  | OID of the tablespace                   |
| relfilenode  | integer  | File node of this relation              |
| bucketid     | smallint | ID of the bucket for consistent hashing |

| Name        | Type                     | Description                                          |
|-------------|--------------------------|------------------------------------------------------|
| forknum     | integer                  | Fork number                                          |
| error_count | integer                  | Number of errors                                     |
| first_time  | timestamp with time zone | Time when the page damage occurs for the first time. |
| last_time   | timestamp with time zone | Time when the page damage occurs for the last time.  |

### 13.2.5.38 SUMMARY\_STAT\_BAD\_BLOCK

Displays information about table and index read failures in the database.

**Table 13-66** SUMMARY\_STAT\_BAD\_BLOCK columns

| Name         | Type                     | Description                                            |
|--------------|--------------------------|--------------------------------------------------------|
| databaseid   | integer                  | OID of the database.                                   |
| tablespaceid | integer                  | OID of the tablespace.                                 |
| relfilenode  | integer                  | File node of this relation.                            |
| forknum      | bigint                   | Fork number.                                           |
| error_count  | bigint                   | Number of errors.                                      |
| first_time   | timestamp with time zone | Time when the page damage occurred for the first time. |
| last_time    | timestamp with time zone | Time when the page damage occurred for the last time.  |

### 13.2.5.39 GLOBAL\_STAT\_BAD\_BLOCK

Displays information about table and index read failures on each node.

**Table 13-67** GLOBAL\_STAT\_BAD\_BLOCK columns

| Name         | Type    | Description                 |
|--------------|---------|-----------------------------|
| node_name    | text    | Node name.                  |
| databaseid   | integer | OID of the database.        |
| tablespaceid | integer | OID of the tablespace.      |
| relfilenode  | integer | File node of this relation. |
| forknum      | integer | Fork number.                |

| Name        | Type                     | Description                                            |
|-------------|--------------------------|--------------------------------------------------------|
| error_count | integer                  | Number of errors.                                      |
| first_time  | timestamp with time zone | Time when the page damage occurred for the first time. |
| last_time   | timestamp with time zone | Time when the page damage occurred for the last time.  |

#### 13.2.5.40 STAT\_USER\_FUNCTIONS

**STAT\_USER\_FUNCTIONS** displays user-defined function status information in the current namespace. (The language of the function is non-internal language.)

**Table 13-68** STAT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                                           |
|------------|------------------|---------------------------------------------------------------------------------------|
| funcid     | oid              | Function ID.                                                                          |
| schemaname | name             | Schema name                                                                           |
| funcname   | name             | UDF name.                                                                             |
| calls      | bigint           | Number of times the function has been called.                                         |
| total_time | double precision | Total time spent in this function, including other functions called by it (unit: ms). |
| self_time  | double precision | Time spent in this function, excluding other functions called by it (unit: ms)        |

#### 13.2.5.41 SUMMARY\_STAT\_USER\_FUNCTIONS

**SUMMARY\_STAT\_USER\_FUNCTIONS** displays statistics about user-defined functions on all database nodes.

**Table 13-69** SUMMARY\_STAT\_USER\_FUNCTIONS columns

| Name       | Type    | Description                                   |
|------------|---------|-----------------------------------------------|
| schemaname | name    | Schema name.                                  |
| funcname   | name    | Function name.                                |
| calls      | numeric | Number of times the function has been called. |

| Name       | Type             | Description                                                                        |
|------------|------------------|------------------------------------------------------------------------------------|
| total_time | double precision | Total time spent in this function and all other functions called by it (unit: ms). |
| self_time  | double precision | Time spent in this function, excluding other functions called by it (unit: ms).    |

### 13.2.5.42 GLOBAL\_STAT\_USER\_FUNCTIONS

GLOBAL\_STAT\_USER\_FUNCTIONS displays statistics about user-defined functions on each node in the database.

**Table 13-70** GLOBAL\_STAT\_USER\_FUNCTIONS columns

| Name       | Type             | Description                                                                          |
|------------|------------------|--------------------------------------------------------------------------------------|
| node_name  | name             | Node name.                                                                           |
| funcid     | oid              | ID of the function.                                                                  |
| schemaname | name             | Schema name                                                                          |
| funcname   | name             | Function name                                                                        |
| calls      | bigint           | Number of times the function has been called.                                        |
| total_time | double precision | Total time spent in this function and all other functions called by it (unit: ms)    |
| self_time  | double precision | Total time spent in this function, excluding other functions called by it (unit: ms) |

## 13.2.6 Workload

### 13.2.6.1 WORKLOAD\_SQL\_COUNT

Displays the distribution of SQL statements in workloads on the current node. Common users can view only the distribution of SQL statements executed by themselves in workloads, whereas the initial user can view the overall load status of workloads.

**Table 13-71** WORKLOAD\_SQL\_COUNT columns

| Name     | Type | Description    |
|----------|------|----------------|
| workload | name | Workload name. |



| Name         | Type   | Description                  |
|--------------|--------|------------------------------|
| select_count | bigint | Number of SELECT statements. |
| update_count | bigint | Number of UPDATE statements. |
| insert_count | bigint | Number of INSERT statements. |
| delete_count | bigint | Number of DELETE statements. |
| ddl_count    | bigint | Number of DDL statements.    |
| dml_count    | bigint | Number of DML statements.    |
| dcl_count    | bigint | Number of DCL statements.    |

### 13.2.6.2 SUMMARY\_WORKLOAD\_SQL\_COUNT

SUMMARY\_WORKLOAD\_SQL\_COUNT displays the distribution of SQL statements in workloads on the primary database node in the database.

**Table 13-72** SUMMARY\_WORKLOAD\_SQL\_COUNT columns

| Name         | Type   | Description                  |
|--------------|--------|------------------------------|
| node_name    | name   | Node name.                   |
| workload     | name   | Workload name.               |
| select_count | bigint | Number of SELECT statements. |
| update_count | bigint | Number of UPDATE statements. |
| insert_count | bigint | Number of INSERT statements. |
| delete_count | bigint | Number of DELETE statements. |
| ddl_count    | bigint | Number of DDL statements.    |
| dml_count    | bigint | Number of DML statements.    |
| dcl_count    | bigint | Number of DCL statements.    |

### 13.2.6.3 WORKLOAD\_TRANSACTION

**WORKLOAD\_TRANSACTION** displays information about transactions loaded on the current node.

**Table 13-73** WORKLOAD\_TRANSACTION columns

| Name                | Type   | Description                                                      |
|---------------------|--------|------------------------------------------------------------------|
| workload            | name   | Workload name                                                    |
| commit_counter      | bigint | Number of user transactions committed                            |
| rollback_counter    | bigint | Number of user transactions rolled back                          |
| resp_min            | bigint | Minimum response time of user transactions (unit: $\mu$ s)       |
| resp_max            | bigint | Maximum response time of user transactions (unit: $\mu$ s)       |
| resp_avg            | bigint | Average response time of user transactions (unit: $\mu$ s)       |
| resp_total          | bigint | Total response time of user transactions (unit: $\mu$ s)         |
| bg_commit_counter   | bigint | Number of background transactions committed                      |
| bg_rollback_counter | bigint | Number of background transactions rolled back                    |
| bg_resp_min         | bigint | Minimum response time of background transactions (unit: $\mu$ s) |
| bg_resp_max         | bigint | Maximum response time of background transactions (unit: $\mu$ s) |
| bg_resp_avg         | bigint | Average response time of background transactions (unit: $\mu$ s) |
| bg_resp_total       | bigint | Total response time of background transactions (unit: $\mu$ s)   |

### 13.2.6.4 SUMMARY\_WORKLOAD\_TRANSACTION

**SUMMARY\_WORKLOAD\_TRANSACTION** displays information about transactions loaded in the database.

**Table 13-74** SUMMARY\_WORKLOAD\_TRANSACTION columns

| Name     | Type | Description   |
|----------|------|---------------|
| workload | name | Workload name |

| Name                | Type    | Description                                                      |
|---------------------|---------|------------------------------------------------------------------|
| commit_counter      | numeric | Number of user transactions committed                            |
| rollback_counter    | numeric | Number of user transactions rolled back                          |
| resp_min            | bigint  | Minimum response time of user transactions (unit: $\mu$ s)       |
| resp_max            | bigint  | Maximum response time of user transactions (unit: $\mu$ s)       |
| resp_avg            | bigint  | Average response time of user transactions (unit: $\mu$ s)       |
| resp_total          | numeric | Total response time of user transactions (unit: $\mu$ s)         |
| bg_commit_counter   | numeric | Number of background transactions committed                      |
| bg_rollback_counter | numeric | Number of background transactions rolled back                    |
| bg_resp_min         | bigint  | Minimum response time of background transactions (unit: $\mu$ s) |
| bg_resp_max         | bigint  | Maximum response time of background transactions (unit: $\mu$ s) |
| bg_resp_avg         | bigint  | Average response time of background transactions (unit: $\mu$ s) |
| bg_resp_total       | numeric | Total response time of background transactions (unit: $\mu$ s)   |

### 13.2.6.5 GLOBAL\_WORKLOAD\_TRANSACTION

**GLOBAL\_WORKLOAD\_TRANSACTION** displays load information about workloads on each node.

**Table 13-75** GLOBAL\_WORKLOAD\_TRANSACTION columns

| Name             | Type   | Description                                                |
|------------------|--------|------------------------------------------------------------|
| node_name        | name   | Node name                                                  |
| workload         | name   | Workload name                                              |
| commit_counter   | bigint | Number of user transactions committed                      |
| rollback_counter | bigint | Number of user transactions rolled back                    |
| resp_min         | bigint | Minimum response time of user transactions (unit: $\mu$ s) |

| Name                | Type   | Description                                                      |
|---------------------|--------|------------------------------------------------------------------|
| resp_max            | bigint | Maximum response time of user transactions (unit: $\mu$ s)       |
| resp_avg            | bigint | Average response time of user transactions (unit: $\mu$ s)       |
| resp_total          | bigint | Total response time of user transactions (unit: $\mu$ s)         |
| bg_commit_counter   | bigint | Number of background transactions committed                      |
| bg_rollback_counter | bigint | Number of background transactions rolled back                    |
| bg_resp_min         | bigint | Minimum response time of background transactions (unit: $\mu$ s) |
| bg_resp_max         | bigint | Maximum response time of background transactions (unit: $\mu$ s) |
| bg_resp_avg         | bigint | Average response time of background transactions (unit: $\mu$ s) |
| bg_resp_total       | bigint | Total response time of background transactions (unit: $\mu$ s)   |

### 13.2.6.6 WORKLOAD\_SQL\_ELAPSE\_TIME

**WORKLOAD\_SQL\_ELAPSE\_TIME** collects statistics about SUIDs in workloads.

**Table 13-76** WORKLOAD\_SQL\_ELAPSE\_TIME columns

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| workload            | name   | Workload name.                                              |
| total_select_elapse | bigint | Total response time of SELECT statements (unit: $\mu$ s).   |
| max_select_elapse   | bigint | Maximum response time of SELECT statements (unit: $\mu$ s). |
| min_select_elapse   | bigint | Minimum response time of SELECT statements (unit: $\mu$ s). |
| avg_select_elapse   | bigint | Average response time of SELECT statements (unit: $\mu$ s). |
| total_update_elapse | bigint | Total response time of UPDATE statements (unit: $\mu$ s).   |

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| max_update_elapse   | bigint | Maximum response time of UPDATE statements (unit: $\mu$ s). |
| min_update_elapse   | bigint | Minimum response time of UPDATE statements (unit: $\mu$ s). |
| avg_update_elapse   | bigint | Average response time of UPDATE statements (unit: $\mu$ s). |
| total_insert_elapse | bigint | Total response time of INSERT statements (unit: $\mu$ s).   |
| max_insert_elapse   | bigint | Maximum response time of INSERT statements (unit: $\mu$ s). |
| min_insert_elapse   | bigint | Minimum response time of INSERT statements (unit: $\mu$ s). |
| avg_insert_elapse   | bigint | Average response time of INSERT statements (unit: $\mu$ s). |
| total_delete_elapse | bigint | Total response time of DELETE statements (unit: $\mu$ s).   |
| max_delete_elapse   | bigint | Maximum response time of DELETE statements (unit: $\mu$ s). |
| min_delete_elapse   | bigint | Minimum response time of DELETE statements (unit: $\mu$ s). |
| avg_delete_elapse   | bigint | Average response time of DELETE statements (unit: $\mu$ s). |

### 13.2.6.7 SUMMARY\_WORKLOAD\_SQL\_ELAPSE\_TIME

SUMMARY\_WORKLOAD\_SQL\_ELAPSE\_TIME collects statistics about SUIDs in workloads on the primary database node.

**Table 13-77** SUMMARY\_WORKLOAD\_SQL\_ELAPSE\_TIME columns

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| node_name           | name   | Node name.                                                  |
| workload            | name   | Workload name.                                              |
| total_select_elapse | bigint | Total response time of SELECT statements (unit: $\mu$ s).   |
| max_select_elapse   | bigint | Maximum response time of SELECT statements (unit: $\mu$ s). |

| Name                | Type   | Description                                                 |
|---------------------|--------|-------------------------------------------------------------|
| min_select_elapse   | bigint | Minimum response time of SELECT statements (unit: $\mu$ s). |
| avg_select_elapse   | bigint | Average response time of SELECT statements (unit: $\mu$ s). |
| total_update_elapse | bigint | Total response time of UPDATE statements (unit: $\mu$ s).   |
| max_update_elapse   | bigint | Maximum response time of UPDATE statements (unit: $\mu$ s). |
| min_update_elapse   | bigint | Minimum response time of UPDATE statements (unit: $\mu$ s). |
| avg_update_elapse   | bigint | Average response time of UPDATE statements (unit: $\mu$ s). |
| total_insert_elapse | bigint | Total response time of INSERT statements (unit: $\mu$ s).   |
| max_insert_elapse   | bigint | Maximum response time of INSERT statements (unit: $\mu$ s). |
| min_insert_elapse   | bigint | Minimum response time of INSERT statements (unit: $\mu$ s). |
| avg_insert_elapse   | bigint | Average response time of INSERT statements (unit: $\mu$ s). |
| total_delete_elapse | bigint | Total response time of DELETE statements (unit: $\mu$ s).   |
| max_delete_elapse   | bigint | Maximum response time of DELETE statements (unit: $\mu$ s). |
| min_delete_elapse   | bigint | Minimum response time of DELETE statements (unit: $\mu$ s). |
| avg_delete_elapse   | bigint | Average response time of DELETE statements (unit: $\mu$ s). |

### 13.2.6.8 USER\_TRANSACTION

**USER\_TRANSACTION** collects statistics about transactions executed by users. Common users can view only transactions executed by themselves, whereas users with the MONADMIN permission can view transactions executed by all users.

**Table 13-78** USER\_TRANSACTION columns

| Name     | Type | Description |
|----------|------|-------------|
| username | name | Username    |

| Name                | Type   | Description                                                      |
|---------------------|--------|------------------------------------------------------------------|
| commit_counter      | bigint | Number of user transactions committed                            |
| rollback_counter    | bigint | Number of user transactions rolled back                          |
| resp_min            | bigint | Minimum response time of user transactions (unit: $\mu$ s)       |
| resp_max            | bigint | Maximum response time of user transactions (unit: $\mu$ s)       |
| resp_avg            | bigint | Average response time of user transactions (unit: $\mu$ s)       |
| resp_total          | bigint | Total response time of user transactions (unit: $\mu$ s)         |
| bg_commit_counter   | bigint | Number of background transactions committed                      |
| bg_rollback_counter | bigint | Number of background transactions rolled back                    |
| bg_resp_min         | bigint | Minimum response time of background transactions (unit: $\mu$ s) |
| bg_resp_max         | bigint | Maximum response time of background transactions (unit: $\mu$ s) |
| bg_resp_avg         | bigint | Average response time of background transactions (unit: $\mu$ s) |
| bg_resp_total       | bigint | Total response time of background transactions (unit: $\mu$ s)   |

### 13.2.6.9 GLOBAL\_USER\_TRANSACTION

**GLOBAL\_USER\_TRANSACTION** collects statistics about transactions executed by all users.

**Table 13-79** GLOBAL\_USER\_TRANSACTION columns

| Name             | Type   | Description                                                |
|------------------|--------|------------------------------------------------------------|
| node_name        | name   | Node name                                                  |
| username         | name   | Username                                                   |
| commit_counter   | bigint | Number of user transactions committed                      |
| rollback_counter | bigint | Number of user transactions rolled back                    |
| resp_min         | bigint | Minimum response time of user transactions (unit: $\mu$ s) |

| Name                | Type   | Description                                                      |
|---------------------|--------|------------------------------------------------------------------|
| resp_max            | bigint | Maximum response time of user transactions (unit: $\mu$ s)       |
| resp_avg            | bigint | Average response time of user transactions (unit: $\mu$ s)       |
| resp_total          | bigint | Total response time of user transactions (unit: $\mu$ s)         |
| bg_commit_counter   | bigint | Number of background transactions committed                      |
| bg_rollback_counter | bigint | Number of background transactions rolled back                    |
| bg_resp_min         | bigint | Minimum response time of background transactions (unit: $\mu$ s) |
| bg_resp_max         | bigint | Maximum response time of background transactions (unit: $\mu$ s) |
| bg_resp_avg         | bigint | Average response time of background transactions (unit: $\mu$ s) |
| bg_resp_total       | bigint | Total response time of background transactions (unit: $\mu$ s)   |

## 13.2.7 Session/Thread

### 13.2.7.1 SESSION\_STAT

**SESSION\_STAT** collects statistics about session status on the current node based on session threads or the **AutoVacuum** thread.

**Table 13-80** SESSION\_STAT columns

| Name     | Type    | Description                     |
|----------|---------|---------------------------------|
| sessid   | text    | Thread start time and ID        |
| statid   | integer | Statistics ID                   |
| statname | text    | Name of the statistics session  |
| statunit | text    | Unit of the statistics session  |
| value    | bigint  | Value of the statistics session |



### 13.2.7.2 GLOBAL\_SESSION\_STAT

**GLOBAL\_SESSION\_STAT** collects statistics about session status on each node based on session threads or the **AutoVacuum** thread.

**Table 13-81** GLOBAL\_SESSION\_STAT columns

| Name      | Type    | Description                     |
|-----------|---------|---------------------------------|
| node_name | name    | Node name                       |
| sessid    | text    | Thread start time and ID        |
| statid    | integer | Statistics ID                   |
| statname  | text    | Name of the statistics session  |
| statunit  | text    | Unit of the statistics session  |
| value     | bigint  | Value of the statistics session |

### 13.2.7.3 SESSION\_TIME

**SESSION\_TIME** collects statistics about the running time of session threads and time consumed in each execution phase on the current node.

**Table 13-82** SESSION\_TIME columns

| Name      | Type    | Description              |
|-----------|---------|--------------------------|
| sessid    | text    | Thread start time and ID |
| stat_id   | integer | Statistics ID            |
| stat_name | text    | Session type name        |
| value     | bigint  | Session value            |

### 13.2.7.4 GLOBAL\_SESSION\_TIME

**GLOBAL\_SESSION\_TIME** collects statistics about the running time of session threads and time consumed in each execution phase on each node, as shown in [Table 13-83](#).

**Table 13-83** GLOBAL\_SESSION\_TIME columns

| Name      | Type    | Description              |
|-----------|---------|--------------------------|
| node_name | name    | Node name                |
| sessid    | text    | Thread start time and ID |
| stat_id   | integer | Statistics ID            |

| Name      | Type   | Description   |
|-----------|--------|---------------|
| stat_name | text   | Session type  |
| value     | bigint | Session value |

### 13.2.7.5 SESSION\_MEMORY

SESSION\_MEMORY collects statistics about memory usage at the session level in the unit of MB, including all the memory allocated to GaussDB and stream threads on DN for jobs currently executed by users, as shown in [Table 13-84](#).

**Table 13-84** SESSION\_MEMORY columns

| Name     | Type    | Description                                                                       |
|----------|---------|-----------------------------------------------------------------------------------|
| sessid   | text    | Thread start time and ID                                                          |
| init_mem | integer | Memory allocated to the currently executed job before the job enters the executor |
| used_mem | integer | Memory allocated to the currently executed job                                    |
| peak_mem | integer | Peak memory allocated to the currently executed job                               |

### 13.2.7.6 GLOBAL\_SESSION\_MEMORY

GLOBAL\_SESSION\_MEMORY collects statistics about memory usage at the session level on each node in the unit of MB, including all the memory allocated to GaussDB and stream threads on DN for jobs currently executed by users, as shown in [Table 13-85](#).

**Table 13-85** GLOBAL\_SESSION\_MEMORY columns

| Name      | Type    | Description                                                                       |
|-----------|---------|-----------------------------------------------------------------------------------|
| node_name | name    | Node name                                                                         |
| sessid    | text    | Thread start time and ID                                                          |
| init_mem  | integer | Memory allocated to the currently executed job before the job enters the executor |
| used_mem  | integer | Memory allocated to the currently executed job                                    |
| peak_mem  | integer | Peak memory allocated to the currently executed job                               |

### 13.2.7.7 SESSION\_MEMORY\_DETAIL

SESSION\_MEMORY\_DETAIL collects statistics on the memory usage of threads on the current node, in the unit specified by **MemoryContext**, as shown in [Table 13-86](#).

**Table 13-86** SESSION\_MEMORY\_DETAIL columns

| Name        | Type     | Description                               |
|-------------|----------|-------------------------------------------|
| sessid      | text     | Thread start time and ID                  |
| sesstype    | text     | Thread name                               |
| contextname | text     | Name of the memory context                |
| level       | smallint | Level of memory context importance        |
| parent      | text     | Name of the parent memory context         |
| totalsize   | bigint   | Size of the allocated memory (unit: byte) |
| freesize    | bigint   | Size of the idle memory (unit: byte)      |
| usedsize    | bigint   | Size of the used memory (unit: byte)      |

### 13.2.7.8 GLOBAL\_SESSION\_MEMORY\_DETAIL

GLOBAL\_SESSION\_MEMORY\_DETAIL collects statistics about thread memory usage on each node by MemoryContext node.

**Table 13-87** GLOBAL\_SESSION\_MEMORY\_DETAIL columns

| Name        | Type     | Description                               |
|-------------|----------|-------------------------------------------|
| node_name   | name     | Node name                                 |
| sessid      | text     | Thread start time and ID                  |
| sesstype    | text     | Thread name                               |
| contextname | text     | Name of the memory context                |
| level       | smallint | Level of memory context importance        |
| parent      | text     | Name of the parent memory context         |
| totalsize   | bigint   | Size of the allocated memory (unit: byte) |
| freesize    | bigint   | Size of the idle memory (unit: byte)      |
| usedsize    | bigint   | Size of the used memory (unit: byte)      |

### 13.2.7.9 SESSION\_STAT\_ACTIVITY

**SESSION\_STAT\_ACTIVITY** displays information about threads that are running on the current node.

**Table 13-88** SESSION\_STAT\_ACTIVITY columns

| Name             | Type                     | Description                                                                                                                                                                                                                                                              |
|------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| datid            | oid                      | OID of the database that the user session connects to in the backend                                                                                                                                                                                                     |
| datname          | name                     | Name of the database that the user session connects to in the backend                                                                                                                                                                                                    |
| pid              | bigint                   | Backend thread ID                                                                                                                                                                                                                                                        |
| usesysid         | oid                      | OID of the user logged in to the backend                                                                                                                                                                                                                                 |
| username         | name                     | Name of the user logged in to the backend                                                                                                                                                                                                                                |
| application_name | text                     | Name of the application connected to the backend                                                                                                                                                                                                                         |
| client_addr      | inet                     | IP address of the client connected to the backend. If this column is <b>NULL</b> , it indicates either the client is connected through a Unix socket on the server or this is an internal thread, for example, <b>autovacuum</b> is displayed for the thread autovacuum. |
| client_hostname  | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                                                                        |
| client_port      | integer                  | TCP port number that the client uses for communication with this backend ( <b>-1</b> if a Unix socket is used).                                                                                                                                                          |
| backend_start    | timestamp with time zone | Time when this process was started, that is, when the client connected to the server                                                                                                                                                                                     |
| xact_start       | timestamp with time zone | Time when the current transaction was started (null if no transactions are active). If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column.                                                 |

| Name         | Type                     | Description                                                                                                                                                                                                                                                                                |
|--------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| query_start  | timestamp with time zone | Time when the currently active query was started, or time when the last query was started if <b>state</b> is not <b>active</b> . For a stored procedure, function, or package, the first query time is queried and does not change with the running of statements in the stored procedure. |
| state_change | timestamp with time zone | Time when the <b>state</b> was last changed                                                                                                                                                                                                                                                |
| waiting      | boolean                  | Specifies whether the backend is currently waiting on a lock. If so, the value is <b>true</b> .                                                                                                                                                                                            |
| enqueue      | text                     | Unsupported currently.                                                                                                                                                                                                                                                                     |

| Name          | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state         | text   | <p>Overall status of this backend. Its value can be:</p> <ul style="list-style-type: none"> <li>• <b>active</b>: The backend is executing a query.</li> <li>• <b>idle</b>: The backend is waiting for a new client command.</li> <li>• <b>idle in transaction</b>: The backend is in a transaction, but is not currently executing a query.</li> <li>• <b>idle in transaction (aborted)</b>: The backend is in a transaction, but there are statements failed in the transaction.</li> <li>• <b>fastpath function call</b>: The backend is executing a fast-path function.</li> <li>• <b>disabled</b>: This state is reported if <b>track_activities</b> is disabled in this backend.</li> </ul> <p><b>NOTE</b><br/>Common users can view their own session status only. The state information of other accounts is empty. For example, after user <b>judy</b> is connected to the database, the state information of user <b>joe</b> and the initial user <b>omm</b> in <b>pg_stat_activity</b> is empty.</p> <pre>openGauss=# SELECT datname, username, usesysid,state,pid FROM pg_stat_activity; datname    username   usesysid   state    pid -----+-----+-----+-----+----- +-----+-----+-----+-----+----- +-----+-----+-----+-----+----- testdb   omm        10        idle    139968752121616 testdb   omm        10        idle    139968903116560 db_tpcds   judy       16398     active   139968391403280 testdb   omm        10        idle    139968643069712 testdb   omm        10        idle    139968680818448 testdb   joe        16390     idle    139968563377936 (6 rows)</pre> |
| resource_pool | name   | Resource pool used by the user                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| query_id      | bigint | ID of a query                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| query         | text   | Text of this backend's latest query. If the value of <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| unique_sql_id | bigint | Unique SQL statement ID                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

| Name     | Type | Description                                                                              |
|----------|------|------------------------------------------------------------------------------------------|
| trace_id | text | Trace ID specified by the driver, which is used to identify a request of an application. |

### 13.2.7.10 GLOBAL\_SESSION\_STAT\_ACTIVITY

GLOBAL\_SESSION\_STAT\_ACTIVITY displays information about threads that are running on each node in the database.

**Table 13-89** GLOBAL\_SESSION\_STAT\_ACTIVITY columns

| Name             | Type                     | Description                                                                                                                                                                                                                                                              |
|------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| coorname         | text                     | Database process name                                                                                                                                                                                                                                                    |
| datid            | oid                      | OID of the database that the user session connects to in the backend                                                                                                                                                                                                     |
| datname          | text                     | Name of the database that the user session connects to in the backend                                                                                                                                                                                                    |
| pid              | bigint                   | Backend thread ID                                                                                                                                                                                                                                                        |
| usesysid         | oid                      | OID of the user logged in to the backend                                                                                                                                                                                                                                 |
| username         | text                     | Name of the user logged in to the backend                                                                                                                                                                                                                                |
| application_name | text                     | Name of the application connected to the backend                                                                                                                                                                                                                         |
| client_addr      | inet                     | IP address of the client connected to the backend. If this column is <b>NULL</b> , it indicates either the client is connected through a Unix socket on the server or this is an internal thread, for example, <b>autovacuum</b> is displayed for the thread autovacuum. |
| client_hostname  | text                     | Host name of the connected client, as reported by a reverse DNS lookup of <b>client_addr</b> . This column will be non-null only for IP connections and only when <b>log_hostname</b> is enabled.                                                                        |
| client_port      | integer                  | TCP port number that the client uses for communication with this backend (-1 if a Unix socket is used).                                                                                                                                                                  |
| backend_start    | timestamp with time zone | Time when this process was started, that is, when the client connected to the server                                                                                                                                                                                     |

| Name         | Type                     | Description                                                                                                                                                                                                                                                                                             |
|--------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| xact_start   | timestamp with time zone | Time when the current transaction was started (null if no transactions are active). If the current query is the first of its transaction, the value of this column is the same as that of the <b>query_start</b> column.                                                                                |
| query_start  | timestamp with time zone | Time when the currently active query was started, or time when the last query was started if the value of <b>state</b> is not <b>active</b> . For a stored procedure, function, or package, the first query time is queried and does not change with the running of statements in the stored procedure. |
| state_change | timestamp with time zone | Time when the <b>state</b> was last changed                                                                                                                                                                                                                                                             |
| waiting      | Boolean                  | Specifies whether the backend is currently waiting for a lock. If the backend is currently waiting for a lock, the value is <b>true</b> .                                                                                                                                                               |
| enqueue      | text                     | Unsupported currently.                                                                                                                                                                                                                                                                                  |



| Name          | Type   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state         | text   | <p>Overall status of this backend. Its value can be:</p> <ul style="list-style-type: none"><li>• <b>active</b>: The backend is executing a query.</li><li>• <b>idle</b>: The backend is waiting for a new client command.</li><li>• <b>idle in transaction</b>: The backend is in a transaction, but is not currently executing a query.</li><li>• <b>idle in transaction (aborted)</b>: The backend is in a transaction, but there are statements failed in the transaction.</li><li>• <b>fastpath function call</b>: The backend is executing a fast-path function.</li><li>• <b>disabled</b>: This state is reported if <b>track_activities</b> is disabled in this backend.</li></ul> <p><b>NOTE</b><br/>Common users can view their own session status only. The state information of other accounts is empty. For example, after the <b>judy</b> user is connected to the database, the state information of the <b>joe</b> user and the initial user <b>omm</b> in <b>pg_stat_activity</b> is empty.</p> <pre>openGauss=# SELECT datname, username, usesysid,state,pid FROM pg_stat_activity; datname    username   usesysid   state         pid -----+-----+-----+-----+----- +-----+-----+-----+-----+----- 139968752121616 testdb   omm        10                139968903116560 db_tpcds   judy       16398     active   139968391403280 testdb   omm        10                139968643069712 testdb   omm        10                139968680818448 testdb   joe        16390             139968563377936 (6 rows)</pre> |
| resource_pool | name   | Resource pool used by the user                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| query_id      | bigint | ID of a query                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| query         | text   | Text of this backend's latest query. If the value of <b>state</b> is <b>active</b> , this column shows the ongoing query. In all other states, it shows the last query that was executed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| unique_sql_id | bigint | Unique SQL statement ID                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

| Name     | Type | Description                                                                              |
|----------|------|------------------------------------------------------------------------------------------|
| trace_id | text | Trace ID specified by the driver, which is used to identify a request of an application. |

### 13.2.7.11 THREAD\_WAIT\_STATUS

This view allows you to test the blocking status about the backend thread and auxiliary thread of the current node.

**Table 13-90** THREAD\_WAIT\_STATUS columns

| Name            | Type    | Description                                                                                                                                                                                       |
|-----------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name       | text    | Current node name.                                                                                                                                                                                |
| db_name         | text    | Database name.                                                                                                                                                                                    |
| thread_name     | text    | Thread name.                                                                                                                                                                                      |
| query_id        | bigint  | Query ID. The value of this column is the same as that of <b>debug_query_id</b> .                                                                                                                 |
| tid             | bigint  | Thread ID of the current thread.                                                                                                                                                                  |
| sessionid       | bigint  | Session ID.                                                                                                                                                                                       |
| lwtid           | integer | Lightweight thread ID of the current thread.                                                                                                                                                      |
| psessionid      | bigint  | Parent thread of the streaming thread.                                                                                                                                                            |
| tlevel          | integer | Level of the streaming thread.                                                                                                                                                                    |
| smpid           | integer | Concurrent thread ID.                                                                                                                                                                             |
| wait_status     | text    | Waiting status of the current thread.                                                                                                                                                             |
| wait_event      | text    | If <b>wait_status</b> is <b>acquire lock</b> , <b>acquire lwlock</b> , or <b>wait io</b> , this column describes the lock, lightweight lock, or I/O information. Otherwise, this column is empty. |
| locktag         | text    | Information about the lock that the current thread is waiting for.                                                                                                                                |
| lockmode        | text    | Lock mode that the current thread is waiting to obtain. The values include table-level lock, row-level lock, and page-level lock modes.                                                           |
| block_sessionid | bigint  | ID of the session that blocks the current thread from obtaining the lock.                                                                                                                         |

| Name             | Type | Description        |
|------------------|------|--------------------|
| global_sessionid | text | Global session ID. |

### 13.2.7.12 GLOBAL\_THREAD\_WAIT\_STATUS

In this view, you can check the blocking and waiting status of backend threads and auxiliary threads on all nodes.

In GLOBAL\_THREAD\_WAIT\_STATUS, you can see all the call hierarchy relationships between threads of the SQL statements on all nodes in the database, and the blocking status for each thread. With this view, you can easily locate the causes of process hang and similar issues.

The definitions of GLOBAL\_THREAD\_WAIT\_STATUS and THREAD\_WAIT\_STATUS are the same, because the essence of the GLOBAL\_THREAD\_WAIT\_STATUS view is the query summary of the THREAD\_WAIT\_STATUS view on each node in the database.

**Table 13-91** GLOBAL\_THREAD\_WAIT\_STATUS columns

| Name        | Type    | Description                                                                                                                                                                                       |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name   | text    | Node name.                                                                                                                                                                                        |
| db_name     | text    | Database name.                                                                                                                                                                                    |
| thread_name | text    | Thread name.                                                                                                                                                                                      |
| query_id    | bigint  | Query ID. The value of this column is the same as that of <b>debug_query_id</b> .                                                                                                                 |
| tid         | bigint  | Thread ID of the current thread.                                                                                                                                                                  |
| sessionid   | bigint  | Session ID.                                                                                                                                                                                       |
| lwtid       | integer | Lightweight thread ID of the current thread.                                                                                                                                                      |
| psessionid  | bigint  | Parent thread of the streaming thread.                                                                                                                                                            |
| tlevel      | integer | Level of the streaming thread.                                                                                                                                                                    |
| smpid       | integer | Concurrent thread ID.                                                                                                                                                                             |
| wait_status | text    | Waiting status of the current thread.                                                                                                                                                             |
| wait_event  | text    | If <b>wait_status</b> is <b>acquire lock</b> , <b>acquire lwlock</b> , or <b>wait io</b> , this column describes the lock, lightweight lock, or I/O information. Otherwise, this column is empty. |
| locktag     | text    | Information about the lock that the current thread is waiting for.                                                                                                                                |

| Name             | Type   | Description                                                                                                                             |
|------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------|
| lockmode         | text   | Lock mode that the current thread is waiting to obtain. The values include table-level lock, row-level lock, and page-level lock modes. |
| block_sessionid  | bigint | ID of the session that blocks the current thread from obtaining the lock.                                                               |
| global_sessionid | text   | Global session ID.                                                                                                                      |

### 13.2.7.13 LOCAL\_THREADPOOL\_STATUS

LOCAL\_THREADPOOL\_STATUS displays the status of worker threads and sessions in a thread pool. This view is valid only when **enable\_thread\_pool** is set to **on**.

**Table 13-92** LOCAL\_THREADPOOL\_STATUS columns

| Name            | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name       | text    | Node name.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| group_id        | integer | ID of the thread pool group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| bind_numa_id    | integer | NUMA ID to which the thread pool group is bound.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| bind_cpu_number | integer | Information about the CPU to which the thread pool group is bound. If no CPUs are bound, the value is <b>NULL</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| listener        | integer | Number of listener threads in the thread pool group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| worker_info     | text    | Information about threads in the thread pool, including: <ul style="list-style-type: none"><li>• <b>default</b>: Number of initial threads in the thread pool group</li><li>• <b>new</b>: Number of new threads in the thread pool group</li><li>• <b>expect</b>: Expected number of threads in the thread pool group</li><li>• <b>actual</b>: Actual number of threads in the thread pool group</li><li>• <b>idle</b>: Number of idle threads in the thread pool group</li><li>• <b>pending</b>: Number of pending threads in the thread pool group</li></ul> |

| Name         | Type | Description                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| session_info | text | Information about sessions in the thread pool, including: <ul style="list-style-type: none"><li>● <b>total</b>: Total number of sessions in the thread pool group.</li><li>● <b>waiting</b>: Number of sessions pending scheduling in the thread pool group.</li><li>● <b>running</b>: Number of running sessions in the thread pool group.</li><li>● <b>idle</b>: Number of idle sessions in the thread pool group.</li></ul> |
| stream_info  | text | Stream pool information, including: <ul style="list-style-type: none"><li>● <b>total</b>: total number of threads in the stream pool group.</li><li>● <b>running</b>: number of threads that are being executed in the stream pool.</li><li>● <b>idle</b>: number of idle threads in the stream pool.</li></ul>                                                                                                                |

#### 13.2.7.14 GLOBAL\_THREADPOOL\_STATUS

**GLOBAL\_THREADPOOL\_STATUS** displays the status of worker threads and sessions in thread pools on all nodes. Columns in this view are the same as those in [Table 13-92](#).

#### 13.2.7.15 LOCAL\_ACTIVE\_SESSION

**LOCAL\_ACTIVE\_SESSION** displays samples in the **ACTIVE SESSION PROFILE** memory on the current node.

**Table 13-93** LOCAL\_ACTIVE\_SESSION columns

| Name              | Type                     | Description                                                |
|-------------------|--------------------------|------------------------------------------------------------|
| sampleid          | bigint                   | Sample ID.                                                 |
| sample_time       | timestamp with time zone | Sampling time.                                             |
| need_flush_sample | Boolean                  | Specifies whether the sample needs to be flushed to disks. |
| databaseid        | oid                      | Database ID.                                               |
| thread_id         | bigint                   | Thread ID.                                                 |
| sessionid         | bigint                   | Session ID.                                                |

| Name             | Type                     | Description                                                                                                                                                            |
|------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| start_time       | timestamp with time zone | Start time of a session.                                                                                                                                               |
| event            | text                     | Event name.                                                                                                                                                            |
| lwtid            | integer                  | Lightweight thread ID of the current thread.                                                                                                                           |
| psessionid       | bigint                   | Parent thread of the streaming thread.                                                                                                                                 |
| tlevel           | integer                  | Level of the streaming thread, which corresponds to the level ( <b>id</b> ) of the execution plan.                                                                     |
| smpid            | integer                  | Concurrent thread ID in SMP execution mode.                                                                                                                            |
| userid           | oid                      | ID of a session user.                                                                                                                                                  |
| application_name | text                     | Name of an application.                                                                                                                                                |
| client_addr      | inet                     | IP address of a client.                                                                                                                                                |
| client_hostname  | text                     | Name of a client.                                                                                                                                                      |
| client_port      | integer                  | TCP port number used by a client to communicate with the backend.                                                                                                      |
| query_id         | bigint                   | Debug query ID.                                                                                                                                                        |
| unique_query_id  | bigint                   | Unique query ID.                                                                                                                                                       |
| user_id          | oid                      | User ID in the key of the unique query.                                                                                                                                |
| cn_id            | integer                  | CN ID. On a DN, this parameter indicates the ID of the node that delivers the unique SQL statement, that is, the value of <b>cn_id</b> in the key of the unique query. |
| unique_query     | text                     | Standardized text string of the unique SQL statement.                                                                                                                  |
| locktag          | text                     | Information of a lock that the session waits for, which can be parsed using <b>locktag_decode</b>                                                                      |
| lockmode         | text                     | Mode of a lock that the session waits for.                                                                                                                             |
| block_sessionid  | bigint                   | Blocks a session from obtaining the session ID of a lock if the session is waiting for the lock.                                                                       |

| Name                  | Type                     | Description                                                                                                                                                                                             |
|-----------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| final_block_sessionid | bigint                   | ID of the blocked session at the source end.                                                                                                                                                            |
| wait_status           | text                     | Provides more details about an event column.                                                                                                                                                            |
| global_sessionid      | text                     | Global session ID.                                                                                                                                                                                      |
| xact_start_time       | timestamp with time zone | Start time of the transaction.                                                                                                                                                                          |
| query_start_time      | timestamp with time zone | Time when the statement starts to be executed.                                                                                                                                                          |
| state                 | text                     | Current statement state.<br>The value can be <b>active</b> , <b>idle in transaction</b> , <b>fastpath function call</b> , <b>idle in transaction (aborted)</b> , <b>disabled</b> , or <b>retrying</b> . |

## 13.2.8 Transaction

### 13.2.8.1 TRANSACTIONS\_PREPARED\_XACTS

TRANSACTIONS\_PREPARED\_XACTS displays information about transactions that are currently prepared for two-phase commit.

**Table 13-94** TRANSACTIONS\_PREPARED\_XACTS columns

| Name        | Type                     | Description                                                        |
|-------------|--------------------------|--------------------------------------------------------------------|
| transaction | xid                      | Numeric transaction identifier of the prepared transaction         |
| gid         | text                     | Global transaction identifier that was assigned to the transaction |
| prepared    | timestamp with time zone | Time at which the transaction is prepared for commit               |
| owner       | name                     | Name of the user that executes the transaction                     |
| database    | name                     | Name of the database in which the transaction is executed          |

### 13.2.8.2 SUMMARY\_TRANSACTIONS\_PREPARED\_XACTS

SUMMARY\_TRANSACTIONS\_PREPARED\_XACTS displays information about transactions that are currently prepared for two-phase commit on the primary database node in the database.

**Table 13-95** SUMMARY\_TRANSACTIONS\_PREPARED\_XACTS columns

| Name        | Type                     | Description                                                        |
|-------------|--------------------------|--------------------------------------------------------------------|
| transaction | xid                      | Numeric transaction identifier of the prepared transaction         |
| gid         | text                     | Global transaction identifier that was assigned to the transaction |
| prepared    | timestamp with time zone | Time at which the transaction is prepared for commit               |
| owner       | name                     | Name of the user that executes the transaction                     |
| database    | name                     | Name of the database in which the transaction is executed          |

### 13.2.8.3 GLOBAL\_TRANSACTIONS\_PREPARED\_XACTS

GLOBAL\_TRANSACTIONS\_PREPARED\_XACTS displays information about transactions that are currently prepared for two-phase commit on each node.

**Table 13-96** GLOBAL\_TRANSACTIONS\_PREPARED\_XACTS columns

| Name        | Type                     | Description                                                        |
|-------------|--------------------------|--------------------------------------------------------------------|
| transaction | xid                      | Numeric transaction identifier of the prepared transaction         |
| gid         | text                     | Global transaction identifier that was assigned to the transaction |
| prepared    | timestamp with time zone | Time at which the transaction is prepared for commit               |
| owner       | name                     | Name of the user that executes the transaction                     |
| database    | name                     | Name of the database in which the transaction is executed          |

### 13.2.8.4 TRANSACTIONS\_RUNNING\_XACTS

Displays information about running transactions on the current node.



**Table 13-97** TRANSACTIONS\_RUNNING\_XACTS columns

| Name        | Type    | Description                                                                                                                                                                                       |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| handle      | integer | Slot handle in the transaction manager corresponding to a transaction. The value is fixed at <b>-1</b> .                                                                                          |
| gxid        | xid     | Transaction ID.                                                                                                                                                                                   |
| state       | tinyint | Transaction status ( <b>3</b> : prepared; <b>0</b> : starting).                                                                                                                                   |
| node        | text    | Node name.                                                                                                                                                                                        |
| xmin        | xid     | Minimum transaction ID on the node.                                                                                                                                                               |
| vacuum      | Boolean | Specifies whether the current transaction is a lazy vacuum (VACUUM only if necessary) transaction. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no</li></ul> |
| timeline    | bigint  | Number of database restarts.                                                                                                                                                                      |
| prepare_xid | xid     | Transaction ID in the <b>prepared</b> state. If the state is not <b>prepared</b> , the value is <b>0</b> .                                                                                        |
| pid         | bigint  | Thread ID corresponding to the transaction.                                                                                                                                                       |
| next_xid    | xid     | Transaction ID sent by other nodes to the current node. The value is fixed at <b>0</b> .                                                                                                          |

### 13.2.8.5 SUMMARY\_TRANSACTIONS\_RUNNING\_XACTS

SUMMARY\_TRANSACTIONS\_RUNNING\_XACTS displays the summary of running transactions on the primary database node.

**Table 13-98** SUMMARY\_TRANSACTIONS\_RUNNING\_XACTS columns

| Name   | Type    | Description                                                                                              |
|--------|---------|----------------------------------------------------------------------------------------------------------|
| handle | integer | Slot handle in the transaction manager corresponding to a transaction. The value is fixed at <b>-1</b> . |
| gxid   | xid     | Transaction ID.                                                                                          |
| state  | tinyint | Transaction status ( <b>3</b> : prepared; <b>0</b> : starting).                                          |
| node   | text    | Node name.                                                                                               |
| xmin   | xid     | Minimum transaction ID on the node.                                                                      |
| vacuum | Boolean | Specifies whether the current transaction is lazy vacuum.                                                |

| Name        | Type   | Description                                                                                                                                                                                       |
|-------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| timeline    | bigint | Specifies whether the current transaction is a lazy vacuum (VACUUM only if necessary) transaction. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no</li></ul> |
| prepare_xid | xid    | Transaction ID in the <b>prepared</b> state. If the state is not <b>prepared</b> , the value is <b>0</b> .                                                                                        |
| pid         | bigint | Thread ID corresponding to the transaction.                                                                                                                                                       |
| next_xid    | xid    | Transaction ID sent by other nodes to the current node. The value is fixed at <b>0</b> .                                                                                                          |

### 13.2.8.6 GLOBAL\_TRANSACTIONS\_RUNNING\_XACTS

GLOBAL\_TRANSACTIONS\_RUNNING\_XACTS displays the summary of running transactions on the primary database node.

**Table 13-99** GLOBAL\_TRANSACTIONS\_RUNNING\_XACTS columns

| Name        | Type    | Description                                                                                                                                                                                       |
|-------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| handle      | integer | Slot handle in the transaction manager corresponding to a transaction. The value is fixed at <b>-1</b> .                                                                                          |
| gxid        | xid     | Transaction ID.                                                                                                                                                                                   |
| state       | tinyint | Transaction status ( <b>3</b> : prepared; <b>0</b> : starting).                                                                                                                                   |
| node        | text    | Node name.                                                                                                                                                                                        |
| xmin        | xid     | Minimum transaction ID on the node.                                                                                                                                                               |
| vacuum      | Boolean | Specifies whether the current transaction is a lazy vacuum (VACUUM only if necessary) transaction. <ul style="list-style-type: none"><li>• <b>true</b>: yes.</li><li>• <b>false</b>: no</li></ul> |
| timeline    | bigint  | Number of database restarts.                                                                                                                                                                      |
| prepare_xid | xid     | Transaction ID in the <b>prepared</b> state. If the state is not <b>prepared</b> , the value is <b>0</b> .                                                                                        |
| pid         | bigint  | Thread ID corresponding to the transaction.                                                                                                                                                       |
| next_xid    | xid     | Transaction ID sent by other nodes to the current node. The value is fixed at <b>0</b> .                                                                                                          |

## 13.2.9 Query

### 13.2.9.1 STATEMENT

Obtains information about executed statements (normalized SQL statements) on the current node. To query a view, you must have the SYSADMIN or MONADMIN permission. You can view all statistics about normalized SQL statements received by the primary database node and other database nodes, whereas you can view only the statistics about normalized SQL statements executed on other database nodes.

#### NOTE

**unique\_sql\_id** generated by different **savepoint\_name** values is different. When a large number of **savepoint\_name** is used, the number of **unique\_sql\_id** values generated in the system increases rapidly. If the number of **unique\_sql\_id** values is greater than the number of **instr\_unique\_sql\_count** values, the newly generated **unique\_sql\_id** information is not counted.

**Table 13-100** STATEMENT columns

| Name              | Type    | Description                                                                                       |
|-------------------|---------|---------------------------------------------------------------------------------------------------|
| node_name         | name    | Node name.                                                                                        |
| node_id           | integer | ID of a node.                                                                                     |
| user_name         | name    | Username.                                                                                         |
| user_id           | oid     | OID of the user.                                                                                  |
| unique_sql_id     | bigint  | ID of the normalized SQL statement.                                                               |
| query             | text    | Normalized SQL statement.<br>Note: The length is controlled by <b>track_activity_query_size</b> . |
| n_calls           | bigint  | Number of calls.                                                                                  |
| min_elapse_time   | bigint  | Minimum execution time of the SQL statement in the kernel (unit: $\mu$ s).                        |
| max_elapse_time   | bigint  | Maximum execution time of the SQL statement in the kernel (unit: $\mu$ s).                        |
| total_elapse_time | bigint  | Total execution time of the SQL statement in the kernel (unit: $\mu$ s).                          |
| n_returned_rows   | bigint  | Number of rows in the result set returned by the SELECT statement.                                |
| n_tuples_fetched  | bigint  | Number of rows randomly scanned.                                                                  |
| n_tuples_returned | bigint  | Number of rows sequentially scanned.                                                              |
| n_tuples_inserted | bigint  | Number of rows inserted.                                                                          |

| Name                | Type   | Description                                                                                                                                                                                                                           |
|---------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n_tuples_updated    | bigint | Number of rows updated.                                                                                                                                                                                                               |
| n_tuples_deleted    | bigint | Number of rows deleted.                                                                                                                                                                                                               |
| n_blocks_fetched    | bigint | Number of logical reads.                                                                                                                                                                                                              |
| n_blocks_hit        | bigint | Number of memory hits.                                                                                                                                                                                                                |
| n_soft_parse        | bigint | Number of soft parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . |
| n_hard_parse        | bigint | Number of hard parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . |
| db_time             | bigint | Valid internal database time, which is accumulated if multiple threads are involved (unit: $\mu$ s).                                                                                                                                  |
| cpu_time            | bigint | CPU time (unit: $\mu$ s).                                                                                                                                                                                                             |
| execution_time      | bigint | Execution time in the executor (unit: $\mu$ s).                                                                                                                                                                                       |
| parse_time          | bigint | SQL parsing time (unit: $\mu$ s).                                                                                                                                                                                                     |
| plan_time           | bigint | SQL plan generation time (unit: $\mu$ s).                                                                                                                                                                                             |
| rewrite_time        | bigint | SQL rewriting time (unit: $\mu$ s).                                                                                                                                                                                                   |
| pl_execution_time   | bigint | Execution time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                           |
| pl_compilation_time | bigint | Compilation time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                         |
| data_io_time        | bigint | I/O time (unit: $\mu$ s).                                                                                                                                                                                                             |
| net_send_info       | text   | Network status of messages sent through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.                                     |
| net_rcv_info        | text   | Network status of messages received through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.                                 |

| Name                 | Type                     | Description                                                                                                                                                                                          |
|----------------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net_stream_send_info | text                     | Network status of messages sent through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.     |
| net_stream_rcv_info  | text                     | Network status of messages received through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode. |
| sort_count           | bigint                   | Sorting count.                                                                                                                                                                                       |
| sort_time            | bigint                   | Sorting duration (unit: $\mu$ s).                                                                                                                                                                    |
| sort_mem_used        | bigint                   | Size of work memory used during sorting (unit: KB).                                                                                                                                                  |
| sort_spill_count     | bigint                   | Count of file writing when data is flushed to disks during sorting.                                                                                                                                  |
| sort_spill_size      | bigint                   | File size used when data is flushed to disks during sorting (unit: KB).                                                                                                                              |
| hash_count           | bigint                   | Hashing count.                                                                                                                                                                                       |
| hash_time            | bigint                   | Hashing duration (unit: $\mu$ s).                                                                                                                                                                    |
| hash_mem_used        | bigint                   | Size of work memory used during hashing (unit: KB).                                                                                                                                                  |
| hash_spill_count     | bigint                   | Number of file writes when data is flushed to disks during hashing.                                                                                                                                  |
| hash_spill_size      | bigint                   | File size used when data is flushed to disks during hashing (unit: KB).                                                                                                                              |
| last_updated         | timestamp with time zone | Last time when the statement was updated.                                                                                                                                                            |

### 13.2.9.2 SUMMARY\_STATEMENT

SUMMARY\_STATEMENT obtains full information about executed statements (normalized SQL statements) on the primary database node.

**Table 13-101** SUMMARY\_STATEMENT columns

| Name      | Type | Description |
|-----------|------|-------------|
| node_name | name | Node name.  |

| Name              | Type    | Description                                                                                          |
|-------------------|---------|------------------------------------------------------------------------------------------------------|
| node_id           | integer | ID of a node.                                                                                        |
| user_name         | name    | Username.                                                                                            |
| user_id           | oid     | OID of the user.                                                                                     |
| unique_sql_id     | bigint  | ID of the normalized SQL statement.                                                                  |
| query             | text    | Normalized SQL statement.<br>Note: The length is controlled by <b>track_activity_query_size</b> .    |
| n_calls           | bigint  | Number of calls.                                                                                     |
| min_elapse_time   | bigint  | Minimum execution time of the SQL statement in the kernel (unit: $\mu$ s).                           |
| max_elapse_time   | bigint  | Maximum execution time of the SQL statement in the kernel (unit: $\mu$ s).                           |
| total_elapse_time | bigint  | Total execution time of the SQL statement in the kernel (unit: $\mu$ s).                             |
| n_returned_rows   | bigint  | Number of rows in the result set returned by the SELECT statement.                                   |
| n_tuples_fetched  | bigint  | Number of rows randomly scanned.                                                                     |
| n_tuples_returned | bigint  | Number of rows sequentially scanned.                                                                 |
| n_tuples_inserted | bigint  | Number of rows inserted.                                                                             |
| n_tuples_updated  | bigint  | Number of rows updated.                                                                              |
| n_tuples_deleted  | bigint  | Number of rows deleted.                                                                              |
| n_blocks_fetched  | bigint  | Number of logical reads.                                                                             |
| n_blocks_hit      | bigint  | Number of memory hits.                                                                               |
| n_soft_parse      | bigint  | Number of soft parsing times.                                                                        |
| n_hard_parse      | bigint  | Number of hard parsing times.                                                                        |
| db_time           | bigint  | Valid internal database time, which is accumulated if multiple threads are involved (unit: $\mu$ s). |
| cpu_time          | bigint  | CPU time (unit: $\mu$ s).                                                                            |
| execution_time    | bigint  | Execution time in the executor (unit: $\mu$ s).                                                      |
| parse_time        | bigint  | SQL parsing time (unit: $\mu$ s).                                                                    |
| plan_time         | bigint  | SQL plan generation time (unit: $\mu$ s).                                                            |
| rewrite_time      | bigint  | SQL rewriting time (unit: $\mu$ s).                                                                  |

| Name                 | Type                     | Description                                                                                                                                                                                           |
|----------------------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| pl_execution_time    | bigint                   | Execution time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                           |
| pl_compilation_time  | bigint                   | Compilation time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                         |
| data_io_time         | bigint                   | I/O time (unit: $\mu$ s).                                                                                                                                                                             |
| net_send_info        | text                     | Network status of messages sent through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.     |
| net_rcv_info         | text                     | Network status of messages received through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode. |
| net_stream_send_info | text                     | Network status of messages sent through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.      |
| net_stream_rcv_info  | text                     | Network status of messages received through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.  |
| last_updated         | timestamp with time zone | Last time when the statement was updated.                                                                                                                                                             |
| sort_count           | bigint                   | Sorting count.                                                                                                                                                                                        |
| sort_time            | bigint                   | Sorting duration (unit: $\mu$ s).                                                                                                                                                                     |
| sort_mem_used        | bigint                   | Size of work memory used during sorting (unit: KB).                                                                                                                                                   |
| sort_spill_count     | bigint                   | Count of file writing when data is flushed to disks during sorting.                                                                                                                                   |
| sort_spill_size      | bigint                   | File size used when data is flushed to disks during sorting (unit: KB).                                                                                                                               |
| hash_count           | bigint                   | Hashing count.                                                                                                                                                                                        |
| hash_time            | bigint                   | Hashing duration (unit: $\mu$ s).                                                                                                                                                                     |
| hash_mem_used        | bigint                   | Size of work memory used during hashing (unit: KB).                                                                                                                                                   |

| Name             | Type   | Description                                                             |
|------------------|--------|-------------------------------------------------------------------------|
| hash_spill_count | bigint | Number of file writes when data is flushed to disks during hashing.     |
| hash_spill_size  | bigint | File size used when data is flushed to disks during hashing (unit: KB). |

### 13.2.9.3 STATEMENT\_COUNT

Displays statistics about five types of running statements (SELECT, INSERT, UPDATE, DELETE, and MERGE INTO) as well as DDL, DML, and DCL statements on the current node of the database.

#### NOTE

By querying the STATEMENT\_COUNT view, a common user can view statistics only about this user on the current node, whereas an administrator can view statistics about all users on the current node. When the database or the node is restarted, the statistics are cleared and the counting restarts. The system counts when a node receives a query, including a query inside the database. For example, when the primary database node receives a query and the query contains multiple subqueries, the number of subqueries is counted on the database node.

**Table 13-102** STATEMENT\_COUNT columns

| Name                | Type   | Description                                                       |
|---------------------|--------|-------------------------------------------------------------------|
| node_name           | text   | Node name.                                                        |
| user_name           | text   | Username                                                          |
| select_count        | bigint | Statistical result of the <b>SELECT</b> statement                 |
| update_count        | bigint | Statistical result of the <b>UPDATE</b> statement                 |
| insert_count        | bigint | Statistical result of the <b>INSERT</b> statement                 |
| delete_count        | bigint | Statistical result of the <b>DELETE</b> statement                 |
| mergeinto_count     | bigint | Statistical result of the <b>MERGE INTO</b> statement             |
| ddl_count           | bigint | Number of DDL statements                                          |
| dml_count           | bigint | Number of DML statements                                          |
| dcl_count           | bigint | Number of DCL statements                                          |
| total_select_elapse | bigint | Total response time of <b>SELECT</b> statements (unit: $\mu$ s)   |
| avg_select_elapse   | bigint | Average response time of <b>SELECT</b> statements (unit: $\mu$ s) |



| Name                | Type   | Description                                                       |
|---------------------|--------|-------------------------------------------------------------------|
| max_select_elapse   | bigint | Maximum response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| min_select_elapse   | bigint | Minimum response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| total_update_elapse | bigint | Total response time of <b>UPDATE</b> statements (unit: $\mu$ s)   |
| avg_update_elapse   | bigint | Average response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| max_update_elapse   | bigint | Maximum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| min_update_elapse   | bigint | Minimum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| total_insert_elapse | bigint | Total response time of <b>INSERT</b> statements (unit: $\mu$ s)   |
| avg_insert_elapse   | bigint | Average response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| max_insert_elapse   | bigint | Maximum response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| min_insert_elapse   | bigint | Minimum response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| total_delete_elapse | bigint | Total response time of <b>DELETE</b> statements (unit: $\mu$ s)   |
| avg_delete_elapse   | bigint | Average response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| max_delete_elapse   | bigint | Maximum response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| min_delete_elapse   | bigint | Minimum response time of <b>DELETE</b> statements (unit: $\mu$ s) |

#### 13.2.9.4 GLOBAL\_STATEMENT\_COUNT

**GLOBAL\_STATEMENT\_COUNT** displays statistics about five types of running statements (**SELECT**, **INSERT**, **UPDATE**, **DELETE**, and **MERGE INTO**) as well as DDL, DML, and DCL statements on each node of the database.

**Table 13-103** GLOBAL\_STATEMENT\_COUNT columns

| Name      | Type | Description |
|-----------|------|-------------|
| node_name | text | Node name   |

| Name                | Type   | Description                                                       |
|---------------------|--------|-------------------------------------------------------------------|
| user_name           | text   | Username                                                          |
| select_count        | bigint | Statistical result of the <b>SELECT</b> statement                 |
| update_count        | bigint | Statistical result of the <b>UPDATE</b> statement                 |
| insert_count        | bigint | Statistical result of the <b>INSERT</b> statement                 |
| delete_count        | bigint | Statistical result of the <b>DELETE</b> statement                 |
| mergeinto_count     | bigint | Statistical result of the <b>MERGE INTO</b> statement             |
| ddl_count           | bigint | Number of DDL statements                                          |
| dml_count           | bigint | Number of DML statements                                          |
| dcl_count           | bigint | Number of DCL statements                                          |
| total_select_elapse | bigint | Total response time of <b>SELECT</b> statements (unit: $\mu$ s)   |
| avg_select_elapse   | bigint | Average response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| max_select_elapse   | bigint | Maximum response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| min_select_elapse   | bigint | Minimum response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| total_update_elapse | bigint | Total response time of <b>UPDATE</b> statements (unit: $\mu$ s)   |
| avg_update_elapse   | bigint | Average response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| max_update_elapse   | bigint | Maximum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| min_update_elapse   | bigint | Minimum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| total_insert_elapse | bigint | Total response time of <b>INSERT</b> statements (unit: $\mu$ s)   |
| avg_insert_elapse   | bigint | Average response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| max_insert_elapse   | bigint | Maximum response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| min_insert_elapse   | bigint | Minimum response time of <b>INSERT</b> statements (unit: $\mu$ s) |

| Name                | Type   | Description                                                       |
|---------------------|--------|-------------------------------------------------------------------|
| total_delete_elapse | bigint | Total response time of <b>DELETE</b> statements (unit: $\mu$ s)   |
| avg_delete_elapse   | bigint | Average response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| max_delete_elapse   | bigint | Maximum response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| min_delete_elapse   | bigint | Minimum response time of <b>DELETE</b> statements (unit: $\mu$ s) |

### 13.2.9.5 SUMMARY\_STATEMENT\_COUNT

SUMMARY\_STATEMENT\_COUNT displays the summary statistics of five types of statements (SELECT, INSERT, UPDATE, DELETE, and MERGE INTO) and DDL, DML, and DCL statements executed on each database node.

**Table 13-104** SUMMARY\_STATEMENT\_COUNT columns

| Name                | Type    | Description                                                       |
|---------------------|---------|-------------------------------------------------------------------|
| user_name           | text    | Username                                                          |
| select_count        | numeric | Statistical result of the <b>SELECT</b> statement                 |
| update_count        | numeric | Statistical result of the <b>UPDATE</b> statement                 |
| insert_count        | numeric | Statistical result of the <b>INSERT</b> statement                 |
| delete_count        | numeric | Statistical result of the <b>DELETE</b> statement                 |
| mergeinto_count     | numeric | Statistical result of the <b>MERGE INTO</b> statement             |
| ddl_count           | numeric | Number of DDL statements                                          |
| dml_count           | numeric | Number of DML statements                                          |
| dcl_count           | numeric | Number of DCL statements                                          |
| total_select_elapse | numeric | Total response time of <b>SELECT</b> statements (unit: $\mu$ s)   |
| avg_select_elapse   | bigint  | Average response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| max_select_elapse   | bigint  | Maximum response time of <b>SELECT</b> statements (unit: $\mu$ s) |

| Name                | Type    | Description                                                       |
|---------------------|---------|-------------------------------------------------------------------|
| min_select_elapse   | bigint  | Minimum response time of <b>SELECT</b> statements (unit: $\mu$ s) |
| total_update_elapse | numeric | Total response time of <b>UPDATE</b> statements (unit: $\mu$ s)   |
| avg_update_elapse   | bigint  | Average response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| max_update_elapse   | bigint  | Maximum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| min_update_elapse   | bigint  | Minimum response time of <b>UPDATE</b> statements (unit: $\mu$ s) |
| total_insert_elapse | numeric | Total response time of <b>INSERT</b> statements (unit: $\mu$ s)   |
| avg_insert_elapse   | bigint  | Average response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| max_insert_elapse   | bigint  | Maximum response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| min_insert_elapse   | bigint  | Minimum response time of <b>INSERT</b> statements (unit: $\mu$ s) |
| total_delete_elapse | numeric | Total response time of <b>DELETE</b> statements (unit: $\mu$ s)   |
| avg_delete_elapse   | bigint  | Average response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| max_delete_elapse   | bigint  | Maximum response time of <b>DELETE</b> statements (unit: $\mu$ s) |
| min_delete_elapse   | bigint  | Minimum response time of <b>DELETE</b> statements (unit: $\mu$ s) |

### 13.2.9.6 STATEMENT\_RESPONSETIME\_PERCENTILE

STATEMENT\_RESPONSETIME\_PERCENTILE obtains the response times of 80% and 95% SQL statements in the database.

**Table 13-105** STATEMENT\_RESPONSETIME\_PERCENTILE columns

| Name | Type   | Description                                                         |
|------|--------|---------------------------------------------------------------------|
| p80  | bigint | Response time of 80% SQL statements in the database (unit: $\mu$ s) |
| p95  | bigint | Response time of 95% SQL statements in the database (unit: $\mu$ s) |

### 13.2.9.7 STATEMENT\_HISTORY

Displays information about statements executed on the current node. To query a view, you must have the SYSADMIN or MONADMIN permission. The result can be queried only in the system database but cannot be queried in the user database.

**Table 13-106** STATEMENT\_HISTORY columns

| Name               | Type                     | Description                                                                                                                                                                                                                           |
|--------------------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dbname             | name                     | Database name.                                                                                                                                                                                                                        |
| schemaname         | name                     | Schema name.                                                                                                                                                                                                                          |
| origin_node        | integer                  | Node name.                                                                                                                                                                                                                            |
| user_name          | name                     | Username.                                                                                                                                                                                                                             |
| application_name   | text                     | Name of the application that sends a request.                                                                                                                                                                                         |
| client_addr        | text                     | IP address of the client that sends a request.                                                                                                                                                                                        |
| client_port        | integer                  | Port number of the client that sends a request.                                                                                                                                                                                       |
| unique_query_id    | bigint                   | ID of the normalized SQL statement.                                                                                                                                                                                                   |
| debug_query_id     | bigint                   | ID of the unique SQL statement.                                                                                                                                                                                                       |
| query              | text                     | Normalized SQL statement.                                                                                                                                                                                                             |
| start_time         | timestamp with time zone | Time when a statement starts.                                                                                                                                                                                                         |
| finish_time        | timestamp with time zone | Time when a statement ends.                                                                                                                                                                                                           |
| slow_sql_threshold | bigint                   | Standard for slow SQL statement execution.                                                                                                                                                                                            |
| transaction_id     | bigint                   | Transaction ID.                                                                                                                                                                                                                       |
| thread_id          | bigint                   | ID of an execution thread.                                                                                                                                                                                                            |
| session_id         | bigint                   | Session ID of a user.                                                                                                                                                                                                                 |
| n_soft_parse       | bigint                   | Number of soft parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . |

| Name                | Type   | Description                                                                                                                                                                                                                           |
|---------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| n_hard_parse        | bigint | Number of hard parsing times. The value of <b>n_soft_parse</b> plus the value of <b>n_hard_parse</b> may be greater than the value of <b>n_calls</b> because the number of subqueries is not counted in the value of <b>n_calls</b> . |
| query_plan          | text   | Statement execution plan.                                                                                                                                                                                                             |
| n_returned_rows     | bigint | Number of rows in the result set returned by the SELECT statement.                                                                                                                                                                    |
| n_tuples_fetched    | bigint | Number of rows randomly scanned.                                                                                                                                                                                                      |
| n_tuples_returned   | bigint | Number of rows sequentially scanned.                                                                                                                                                                                                  |
| n_tuples_inserted   | bigint | Number of rows inserted.                                                                                                                                                                                                              |
| n_tuples_updated    | bigint | Number of rows updated.                                                                                                                                                                                                               |
| n_tuples_deleted    | bigint | Number of rows deleted.                                                                                                                                                                                                               |
| n_blocks_fetched    | bigint | Number of buffer block access times.                                                                                                                                                                                                  |
| n_blocks_hit        | bigint | Number of buffer block hits.                                                                                                                                                                                                          |
| db_time             | bigint | Valid DB time, which is accumulated if multiple threads are involved (unit: $\mu$ s).                                                                                                                                                 |
| cpu_time            | bigint | CPU time (unit: $\mu$ s).                                                                                                                                                                                                             |
| execution_time      | bigint | Execution time in the executor (unit: $\mu$ s).                                                                                                                                                                                       |
| parse_time          | bigint | SQL parsing time (unit: $\mu$ s).                                                                                                                                                                                                     |
| plan_time           | bigint | SQL plan generation time (unit: $\mu$ s).                                                                                                                                                                                             |
| rewrite_time        | bigint | SQL rewriting time (unit: $\mu$ s).                                                                                                                                                                                                   |
| pl_execution_time   | bigint | Execution time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                           |
| pl_compilation_time | bigint | Compilation time of PL/pgSQL (unit: $\mu$ s).                                                                                                                                                                                         |
| data_io_time        | bigint | I/O time (unit: $\mu$ s).                                                                                                                                                                                                             |

| Name                 | Type   | Description                                                                                                                                                                                           |
|----------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| net_send_info        | text   | Network status of messages sent through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.     |
| net_rcv_info         | text   | Network status of messages received through a physical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode. |
| net_stream_send_info | text   | Network status of messages sent through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.      |
| net_stream_rcv_info  | text   | Network status of messages received through a logical connection, including the time (unit: $\mu$ s), number of calls, and throughput (unit: byte). This column is not supported in standalone mode.  |
| lock_count           | bigint | Number of locks.                                                                                                                                                                                      |
| lock_time            | bigint | Time required for locking.                                                                                                                                                                            |
| lock_wait_count      | bigint | Number of lock waits.                                                                                                                                                                                 |
| lock_wait_time       | bigint | Time required for lock waiting.                                                                                                                                                                       |
| lock_max_count       | bigint | Maximum number of locks.                                                                                                                                                                              |
| lwlock_count         | bigint | Number of lightweight locks (reserved).                                                                                                                                                               |
| lwlock_wait_count    | bigint | Number of lightweight lock waits.                                                                                                                                                                     |
| lwlock_time          | bigint | Time required for lightweight locking (reserved).                                                                                                                                                     |
| lwlock_wait_time     | bigint | Time required for lightweight locking.                                                                                                                                                                |

| Name        | Type    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| details     | bytea   | List of statement lock events, which are recorded in time sequence. The number of records is affected by the <b>track_stmt_details_size</b> parameter.<br>Events include: <ul style="list-style-type: none"><li>• Start locking.</li><li>• Complete locking.</li><li>• Start lock waiting.</li><li>• Complete lock waiting.</li><li>• Start unlocking.</li><li>• Complete unlocking.</li><li>• Start lightweight lock waiting.</li><li>• Complete lightweight lock waiting.</li></ul>                                                                                                                                                                                        |
| is_slow_sql | Boolean | Specifies whether the SQL statement is a slow SQL statement.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| trace_id    | text    | Driver-specific trace ID, which is associated with an application request.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| advise      | text    | Risks which may cause slow SQL statements. (Multiple risks may exist at the same time.) <ul style="list-style-type: none"><li>• Cast Function Cause Index Miss. : Index matching may fail due to implicit conversion.</li><li>• Limit too much rows. : The SQL statement execution may slow down due to a large <b>limit</b> value.</li><li>• Proleakproof of function is false. : The <b>proleakproof</b> of the function is set to <b>false</b>. In this case, the function does not use statistics when generating a plan due to data leakage risks. As a result, the accuracy of the generated plan is affected and the SQL statement execution may slow down.</li></ul> |

### 13.2.9.8 GS\_SLOW\_QUERY\_INFO (Deprecated)

GS\_SLOW\_QUERY\_INFO displays the slow query information that has been dumped on the current node. The data is dumped from the kernel to the system catalog. If the GUC parameter **enable\_resource\_record** is set to **on**, the system imports the query information from the kernel to the GS\_WLM\_SESSION\_QUERY\_INFO\_ALL system catalog every 3 minutes. This operation occupies storage space and affects performance. You can check



GS\_SLOW\_QUERY\_INFO to view the slow query information that has been dumped. This view has been deprecated in this version.

**Table 13-107** GS\_SLOW\_QUERY\_INFO columns

| Name              | Type                     | Description                                                                                 |
|-------------------|--------------------------|---------------------------------------------------------------------------------------------|
| dbname            | text                     | Database name.                                                                              |
| schemaname        | text                     | Schema name.                                                                                |
| nodename          | text                     | Node name.                                                                                  |
| username          | text                     | Username.                                                                                   |
| queryid           | bigint                   | Normalization ID.                                                                           |
| query             | text                     | User statement.                                                                             |
| start_time        | timestamp with time zone | Execution start time.                                                                       |
| finish_time       | timestamp with time zone | Execution end time.                                                                         |
| duration          | bigint                   | Execution duration (unit: ms).                                                              |
| query_plan        | text                     | Plan information.                                                                           |
| n_returned_rows   | bigint                   | Number of rows in the result set returned by the SELECT statement.                          |
| n_tuples_fetched  | bigint                   | Number of rows randomly scanned.                                                            |
| n_tuples_returned | bigint                   | Number of rows sequentially scanned.                                                        |
| n_tuples_inserted | bigint                   | Number of rows inserted.                                                                    |
| n_tuples_updated  | bigint                   | Number of rows updated.                                                                     |
| n_tuples_deleted  | bigint                   | Number of rows deleted.                                                                     |
| n_blocks_fetched  | bigint                   | Number of cache loading times.                                                              |
| n_blocks_hit      | bigint                   | Number of cache hits.                                                                       |
| db_time           | bigint                   | Valid database time, which is accumulated if multiple threads are involved (unit: $\mu$ s). |
| cpu_time          | bigint                   | CPU time (unit: $\mu$ s).                                                                   |

| Name                | Type   | Description                                              |
|---------------------|--------|----------------------------------------------------------|
| execution_time      | bigint | Execution time in the executor (unit: $\mu$ s).          |
| parse_time          | bigint | SQL parsing time (unit: $\mu$ s).                        |
| plan_time           | bigint | SQL plan generation time (unit: $\mu$ s).                |
| rewrite_time        | bigint | SQL rewriting time (unit: $\mu$ s).                      |
| pl_execution_time   | bigint | Execution time of PL/pgSQL statements (unit: $\mu$ s).   |
| pl_compilation_time | bigint | Compilation time of PL/pgSQL statements (unit: $\mu$ s). |
| net_send_time       | bigint | Network time (unit: $\mu$ s).                            |
| data_io_time        | bigint | I/O time (unit: $\mu$ s).                                |

### 13.2.9.9 GS\_SLOW\_QUERY\_HISTORY (Deprecated)

**GS\_SLOW\_QUERY\_HISTORY** displays the slow query information that is not dumped on the current node. For details, see **GS\_SLOW\_QUERY\_INFO**. Only the **system admin** and **monitor admin** users have the permission to query this view. This view is deprecated in this version.

### 13.2.9.10 GLOBAL\_SLOW\_QUERY\_HISTORY (Deprecated)

**GS\_SLOW\_QUERY\_HISTORY** displays the slow query information that is not dumped on all nodes. This view is deprecated in this version. For details, see **GS\_SLOW\_QUERY\_INFO**.

### 13.2.9.11 GLOBAL\_SLOW\_QUERY\_INFO (Deprecated)

**GS\_SLOW\_QUERY\_HISTORY** displays the slow query information that has been dumped on all nodes. This view is deprecated in this version. For details, see **GS\_SLOW\_QUERY\_INFO**.

## 13.2.10 Cache/IO

### 13.2.10.1 STATIO\_USER\_TABLES

**STATIO\_USER\_TABLES** displays I/O status information about all user relationship tables in the namespace.

**Table 13-108** STATIO\_USER\_TABLES columns

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| relid           | oid    | Table OID                                                                   |
| schemaname      | name   | Name of the schema that the table is in                                     |
| relname         | name   | Table name                                                                  |
| heap_blks_read  | bigint | Number of disk blocks read from the table                                   |
| heap_blks_hit   | bigint | Number of cache hits in the table                                           |
| idx_blks_read   | bigint | Number of disk blocks read from indexes in the table                        |
| idx_blks_hit    | bigint | Number of cache hits in indexes in the table                                |
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.2 SUMMARY\_STATIO\_USER\_TABLES

SUMMARY\_STATIO\_USER\_TABLES displays a summary of I/O status information about user relationship tables of each node in the database.

**Table 13-109** SUMMARY\_STATIO\_USER\_TABLES columns

| Name           | Type    | Description                                              |
|----------------|---------|----------------------------------------------------------|
| schemaname     | name    | Name of the schema that contains the table               |
| relname        | name    | Table name.                                              |
| heap_blks_read | numeric | Number of disk blocks read from the table                |
| heap_blks_hit  | numeric | Number of cache hits in the table                        |
| idx_blks_read  | numeric | Number of disk blocks read from all indexes in the table |
| idx_blks_hit   | numeric | Number of cache hits in the table                        |

| Name            | Type    | Description                                                                 |
|-----------------|---------|-----------------------------------------------------------------------------|
| toast_blks_read | numeric | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | numeric | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | numeric | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | numeric | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.3 GLOBAL\_STATIO\_USER\_TABLES

**GLOBAL\_STATIO\_USER\_TABLES** displays I/O status information about all user relationship tables in namespaces on each node.

**Table 13-110** GLOBAL\_STATIO\_USER\_TABLES columns

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| node_name       | name   | Node name                                                                   |
| relid           | oid    | Table OID                                                                   |
| schemaname      | name   | Name of the schema that contains the table                                  |
| relname         | name   | Table name                                                                  |
| heap_blks_read  | bigint | Number of disk blocks read from the table                                   |
| heap_blks_hit   | bigint | Number of cache hits in the table                                           |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in the table                    |
| idx_blks_hit    | bigint | Number of cache hits in the table                                           |
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.4 STATIO\_USER\_INDEXES

**STATIO\_USER\_INDEXES** displays I/O status information about all user relationship table indexes in namespaces on the current node.

**Table 13-111** STATIO\_USER\_INDEXES columns

| Name          | Type   | Description                                     |
|---------------|--------|-------------------------------------------------|
| relid         | oid    | OID of the table that the index is created for  |
| indexrelid    | oid    | OID of the index                                |
| schemaname    | name   | Name of the schema that the index is in         |
| relname       | name   | Name of the table that the index is created for |
| indexrelname  | name   | Index name                                      |
| idx_blks_read | bigint | Number of disk blocks read from the index       |
| idx_blks_hit  | bigint | Number of cache hits in the index               |

### 13.2.10.5 SUMMARY\_STATIO\_USER\_INDEXES

**SUMMARY\_STATIO\_USER\_INDEXES** displays I/O status information about all user relationship table indexes in namespaces in the database.

**Table 13-112** SUMMARY\_STATIO\_USER\_INDEXES columns

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |
| idx_blks_hit  | numeric | Number of cache hits in the index               |

### 13.2.10.6 GLOBAL\_STATIO\_USER\_INDEXES

**GLOBAL\_STATIO\_USER\_INDEXES** displays I/O status information about all user relationship table indexes in namespaces on each node.

**Table 13-113** GLOBAL\_STATIO\_USER\_INDEXES columns

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| node_name     | name    | Node name                                       |
| relid         | oid     | OID of the table that the index is created for  |
| indexrelid    | oid     | OID of the index                                |
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |
| idx_blks_hit  | numeric | Number of cache hits in the index               |

### 13.2.10.7 STATIO\_USER\_SEQUENCES

**STATIO\_USER\_SEQUENCE** displays I/O status information about all user relationship table sequences in namespaces on the current node.

**Table 13-114** STATIO\_USER\_SEQUENCE columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| relid      | oid    | OID of the sequence                          |
| schemaname | name   | Name of the schema that the sequence is in   |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Number of cache hits in the sequence         |

### 13.2.10.8 SUMMARY\_STATIO\_USER\_SEQUENCES

**SUMMARY\_STATIO\_USER\_SEQUENCES** displays I/O status information about all user relationship table sequences in namespaces in the database.

**Table 13-115** SUMMARY\_STATIO\_USER\_SEQUENCES columns

| Name       | Type    | Description                                  |
|------------|---------|----------------------------------------------|
| schemaname | name    | Name of the schema that the sequence is in   |
| relname    | name    | Sequence name                                |
| blks_read  | numeric | Number of disk blocks read from the sequence |
| blks_hit   | numeric | Number of cache hits in the sequence         |

### 13.2.10.9 GLOBAL\_STATIO\_USER\_SEQUENCES

**GLOBAL\_STATIO\_USER\_SEQUENCES** displays I/O status information about all user relationship table sequences in namespaces on each node.

**Table 13-116** GLOBAL\_STATIO\_USER\_SEQUENCES columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| node_name  | name   | Node name                                    |
| relid      | oid    | OID of the sequence                          |
| schemaname | name   | Name of the schema that the sequence is in   |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Number of cache hits in the sequence         |

### 13.2.10.10 STATIO\_SYS\_TABLES

**STATIO\_SYS\_TABLES** displays the I/O status information about all system catalogs in namespaces on the current node.

**Table 13-117** STATIO\_SYS\_TABLES columns

| Name           | Type   | Description                               |
|----------------|--------|-------------------------------------------|
| relid          | oid    | Table OID                                 |
| schemaname     | name   | Name of the schema that the table is in   |
| relname        | name   | Table name                                |
| heap_blks_read | bigint | Number of disk blocks read from the table |

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| heap_blks_hit   | bigint | Number of cache hits in the table                                           |
| idx_blks_read   | bigint | Number of disk blocks read from indexes in the table                        |
| idx_blks_hit    | bigint | Number of cache hits in indexes in the table                                |
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.11 SUMMARY\_STATIO\_SYS\_TABLES

SUMMARY\_STATIO\_SYS\_TABLES displays a summary of I/O status information about system catalogs of each node in the database.

**Table 13-118** SUMMARY\_STATIO\_SYS\_TABLES columns

| Name            | Type    | Description                                                                 |
|-----------------|---------|-----------------------------------------------------------------------------|
| schemaname      | name    | Name of the schema that contains the table                                  |
| relname         | name    | Table name.                                                                 |
| heap_blks_read  | numeric | Number of disk blocks read from the table                                   |
| heap_blks_hit   | numeric | Number of cache hits in the table                                           |
| idx_blks_read   | numeric | Number of disk blocks read from all indexes in the table                    |
| idx_blks_hit    | numeric | Number of cache hits in the table                                           |
| toast_blks_read | numeric | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | numeric | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | numeric | Number of disk blocks read from the TOAST table index (if any) in the table |



| Name          | Type    | Description                                                          |
|---------------|---------|----------------------------------------------------------------------|
| tidx_blks_hit | numeric | Number of buffer-hits in the TOAST table index (if any) in the table |

### 13.2.10.12 GLOBAL\_STATIO\_SYS\_TABLES

**GLOBAL\_STATIO\_SYS\_TABLES** displays I/O status information about all system catalogs in namespaces on each node.

**Table 13-119** GLOBAL\_STATIO\_SYS\_TABLES columns

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| node_name       | name   | Node name                                                                   |
| relid           | oid    | Table OID                                                                   |
| schemaname      | name   | Name of the schema that contains the table                                  |
| relname         | name   | Table name                                                                  |
| heap_blks_read  | bigint | Number of disk blocks read from the table                                   |
| heap_blks_hit   | bigint | Number of cache hits in the table                                           |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in the table                    |
| idx_blks_hit    | bigint | Number of cache hits in the table                                           |
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.13 STATIO\_SYS\_INDEXES

**STATIO\_SYS\_INDEXES** displays the I/O status information about all system catalog indexes in the current namespace.

**Table 13-120** STATIO\_SYS\_INDEXES columns

| Name          | Type   | Description                                     |
|---------------|--------|-------------------------------------------------|
| relid         | oid    | OID of the table that the index is created for  |
| indexrelid    | oid    | OID of the index                                |
| schemaname    | name   | Name of the schema that the index is in         |
| relname       | name   | Name of the table that the index is created for |
| indexrelname  | name   | Index name                                      |
| idx_blks_read | bigint | Number of disk blocks read from the index       |
| idx_blks_hit  | bigint | Number of cache hits in the index               |

#### 13.2.10.14 SUMMARY\_STATIO\_SYS\_INDEXES

SUMMARY\_STATIO\_SYS\_INDEXES displays I/O status information about all system catalog indexes in namespaces in the database.

**Table 13-121** SUMMARY\_STATIO\_SYS\_INDEXES columns

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |
| idx_blks_hit  | numeric | Number of cache hits in the index               |

#### 13.2.10.15 GLOBAL\_STATIO\_SYS\_INDEXES

GLOBAL\_STATIO\_SYS\_INDEXES displays I/O status information about all system catalog indexes in namespaces on each node.

**Table 13-122** GLOBAL\_STATIO\_SYS\_INDEXES columns

| Name      | Type | Description                                    |
|-----------|------|------------------------------------------------|
| node_name | name | Node name                                      |
| relid     | oid  | OID of the table that the index is created for |

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| indexrelid    | oid     | OID of the index                                |
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |
| idx_blks_hit  | numeric | Number of cache hits in the index               |

### 13.2.10.16 STATIO\_SYS\_SEQUENCES

**STATIO\_SYS\_SEQUENCES** shows the I/O status information about all the system sequences in the current namespace.

**Table 13-123** STATIO\_SYS\_SEQUENCES columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| relid      | oid    | OID of the sequence                          |
| schemaname | name   | Name of the schema that the sequence is in   |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Number of cache hits in the sequence         |

### 13.2.10.17 SUMMARY\_STATIO\_SYS\_SEQUENCES

**SUMMARY\_STATIO\_SYS\_SEQUENCES** displays I/O status information about all system sequences in namespaces in the database.

**Table 13-124** SUMMARY\_STATIO\_SYS\_SEQUENCES columns

| Name       | Type    | Description                                  |
|------------|---------|----------------------------------------------|
| schemaname | name    | Name of the schema that the sequence is in   |
| relname    | name    | Sequence name                                |
| blks_read  | numeric | Number of disk blocks read from the sequence |

| Name     | Type    | Description                          |
|----------|---------|--------------------------------------|
| blks_hit | numeric | Number of cache hits in the sequence |

### 13.2.10.18 GLOBAL\_STATIO\_SYS\_SEQUENCES

**GLOBAL\_STATIO\_SYS\_SEQUENCES** displays I/O status information about all system sequences in namespaces on each node.

**Table 13-125** GLOBAL\_STATIO\_SYS\_SEQUENCES columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| node_name  | name   | Node name                                    |
| relid      | oid    | OID of the sequence                          |
| schemaname | name   | Name of the schema that the sequence is in   |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Number of cache hits in the sequence         |

### 13.2.10.19 STATIO\_ALL\_TABLES

**STATIO\_ALL\_TABLES** displays the I/O status information about each table (including TOAST tables) of the current node in the database.

**Table 13-126** STATIO\_ALL\_TABLES columns

| Name           | Type   | Description                                              |
|----------------|--------|----------------------------------------------------------|
| relid          | oid    | Table OID                                                |
| schemaname     | name   | Name of the schema that the table is in                  |
| relname        | name   | Table name                                               |
| heap_blks_read | bigint | Number of disk blocks read from the table                |
| heap_blks_hit  | bigint | Number of cache hits in the table                        |
| idx_blks_read  | bigint | Number of disk blocks read from all indexes in the table |
| idx_blks_hit   | bigint | Number of cache hits in indexes in the table             |

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.20 SUMMARY\_STATIO\_ALL\_TABLES

SUMMARY\_STATIO\_ALL\_TABLES displays a summary of the I/O status information about each table (including TOAST tables) of each node in the database.

**Table 13-127** SUMMARY\_STATIO\_ALL\_TABLES columns

| Name            | Type    | Description                                                                 |
|-----------------|---------|-----------------------------------------------------------------------------|
| schemaname      | name    | Name of the schema that contains the table                                  |
| relname         | name    | Table name.                                                                 |
| heap_blks_read  | numeric | Number of disk blocks read from the table                                   |
| heap_blks_hit   | numeric | Number of cache hits in the table                                           |
| idx_blks_read   | numeric | Number of disk blocks read from all indexes in the table                    |
| idx_blks_hit    | numeric | Number of cache hits in the table                                           |
| toast_blks_read | numeric | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | numeric | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | numeric | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | numeric | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.21 GLOBAL\_STATIO\_ALL\_TABLES

**GLOBAL\_STATIO\_ALL\_TABLES** contains I/O statistics about each table (including TOAST tables) in databases on each node.

**Table 13-128** GLOBAL\_STATIO\_ALL\_TABLES columns

| Name            | Type   | Description                                                                 |
|-----------------|--------|-----------------------------------------------------------------------------|
| node_name       | name   | Node name                                                                   |
| relid           | oid    | Table OID                                                                   |
| schemaname      | name   | Name of the schema that contains the table                                  |
| relname         | name   | Table name                                                                  |
| heap_blks_read  | bigint | Number of disk blocks read from the table                                   |
| heap_blks_hit   | bigint | Number of cache hits in the table                                           |
| idx_blks_read   | bigint | Number of disk blocks read from all indexes in the table                    |
| idx_blks_hit    | bigint | Number of cache hits in the table                                           |
| toast_blks_read | bigint | Number of disk blocks read from the TOAST table (if any) in the table       |
| toast_blks_hit  | bigint | Number of buffer hits in the TOAST table (if any) in the table              |
| tidx_blks_read  | bigint | Number of disk blocks read from the TOAST table index (if any) in the table |
| tidx_blks_hit   | bigint | Number of buffer-hits in the TOAST table index (if any) in the table        |

### 13.2.10.22 STATIO\_ALL\_INDEXES

**STATIO\_ALL\_INDEXES** contains one row for each index in the current database, showing I/O statistics about specific indexes.

**Table 13-129** STATIO\_ALL\_INDEXES columns

| Name         | Type | Description                                     |
|--------------|------|-------------------------------------------------|
| relid        | oid  | OID of the table that the index is created for  |
| indexrelid   | oid  | OID of the index                                |
| schemaname   | name | Name of the schema that the index is in         |
| relname      | name | Name of the table that the index is created for |
| indexrelname | name | Index name                                      |

| Name          | Type   | Description                               |
|---------------|--------|-------------------------------------------|
| idx_blks_read | bigint | Number of disk blocks read from the index |
| idx_blks_hit  | bigint | Number of cache hits in the index         |

### 13.2.10.23 SUMMARY\_STATIO\_ALL\_INDEXES

SUMMARY\_STATIO\_ALL\_INDEXES contains one row for each index in the database, showing a summary of I/O statistics about specific indexes.

**Table 13-130** SUMMARY\_STATIO\_ALL\_INDEXES columns

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |
| idx_blks_hit  | numeric | Number of cache hits in the index               |

### 13.2.10.24 GLOBAL\_STATIO\_ALL\_INDEXES

GLOBAL\_STATIO\_ALL\_INDEXES contains one row for each index in databases on each node, showing I/O statistics about specific indexes.

**Table 13-131** GLOBAL\_STATIO\_ALL\_INDEXES columns

| Name          | Type    | Description                                     |
|---------------|---------|-------------------------------------------------|
| node_name     | name    | Node name                                       |
| relid         | oid     | OID of the table that the index is created for  |
| indexrelid    | oid     | OID of the index                                |
| schemaname    | name    | Name of the schema that the index is in         |
| relname       | name    | Name of the table that the index is created for |
| indexrelname  | name    | Index name                                      |
| idx_blks_read | numeric | Number of disk blocks read from the index       |

| Name         | Type    | Description                       |
|--------------|---------|-----------------------------------|
| idx_blks_hit | numeric | Number of cache hits in the index |

### 13.2.10.25 STATIO\_ALL\_SEQUENCES

**STATIO\_ALL\_SEQUENCES** contains one row for each sequence in the current database, showing I/O statistics about specific sequences.

**Table 13-132** STATIO\_ALL\_SEQUENCES columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| relid      | oid    | OID of this sequence                         |
| schemaname | name   | Name of the schema where the sequence is in  |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Cache hits in the sequence                   |

### 13.2.10.26 SUMMARY\_STATIO\_ALL\_SEQUENCES

**SUMMARY\_STATIO\_ALL\_SEQUENCES** contains one row for each sequence in a database in the entire database system, showing I/O statistics about specific sequences.

**Table 13-133** SUMMARY\_STATIO\_ALL\_SEQUENCES columns

| Name       | Type    | Description                                  |
|------------|---------|----------------------------------------------|
| schemaname | name    | Name of the schema that the sequence is in   |
| relname    | name    | Sequence name                                |
| blks_read  | numeric | Number of disk blocks read from the sequence |
| blks_hit   | numeric | Number of cache hits in the sequence         |

### 13.2.10.27 GLOBAL\_STATIO\_ALL\_SEQUENCES

**GLOBAL\_STATIO\_ALL\_SEQUENCES** contains every row of each sequence in databases on each node, showing I/O statistics about specific sequences.



**Table 13-134** GLOBAL\_STATIO\_ALL\_SEQUENCES columns

| Name       | Type   | Description                                  |
|------------|--------|----------------------------------------------|
| node_name  | name   | Node name                                    |
| relid      | oid    | OID of the sequence                          |
| schemaname | name   | Name of the schema that the sequence is in   |
| relname    | name   | Sequence name                                |
| blks_read  | bigint | Number of disk blocks read from the sequence |
| blks_hit   | bigint | Number of cache hits in the sequence         |

### 13.2.10.28 GLOBAL\_STAT\_DB\_CU

GLOBAL\_STAT\_DB\_CU queries CU hits in each database in GaussDB. You can clear it using pg\_stat\_reset().

**Table 13-135** GLOBAL\_STAT\_DB\_CU columns

| Name           | Type   | Description                             |
|----------------|--------|-----------------------------------------|
| node_name<br>1 | text   | Node name.                              |
| db_name        | text   | Database name.                          |
| mem_hit        | bigint | Number of memory hits.                  |
| hdd_sync_read  | bigint | Number of synchronous hard disk reads.  |
| hdd_async_read | bigint | Number of asynchronous hard disk reads. |

### 13.2.10.29 GLOBAL\_STAT\_SESSION\_CU

GLOBAL\_STAT\_SESSION\_CU is used to query the CU hit rate of running sessions on each node in the database. This data about a session is cleared when you exit this session or restart the database.

**Table 13-136** GLOBAL\_STAT\_SESSION\_CU columns

| Name          | Type    | Description                           |
|---------------|---------|---------------------------------------|
| mem_hit       | integer | Number of memory hits                 |
| hdd_sync_read | integer | Number of synchronous hard disk reads |

| Name          | Type    | Description                            |
|---------------|---------|----------------------------------------|
| hdd_asyn_read | integer | Number of asynchronous hard disk reads |

## 13.2.11 Utility

### 13.2.11.1 REPLICATION\_STAT

REPLICATION\_STAT describes information about log synchronization status, such as the locations where the sender sends logs and where the receiver receives logs.

**Table 13-137** REPLICATION\_STAT columns

| Name                     | Type                     | Description                                                                                                                    |
|--------------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| pid                      | bigint                   | PID of the thread.                                                                                                             |
| usesysid                 | oid                      | User system ID.                                                                                                                |
| username                 | name                     | Username.                                                                                                                      |
| application_name         | text                     | Program name.                                                                                                                  |
| client_addr              | inet                     | Client address.                                                                                                                |
| client_hostname          | text                     | Client name.                                                                                                                   |
| client_port              | integer                  | Port of the client.                                                                                                            |
| backend_start            | timestamp with time zone | Start time of the program.                                                                                                     |
| state                    | text                     | Log replication state: <ul style="list-style-type: none"><li>• Catch-up state.</li><li>• Consistent streaming state.</li></ul> |
| sender_sent_location     | text                     | Location where the sender sends logs.                                                                                          |
| receiver_write_location  | text                     | Location where the receiver writes logs.                                                                                       |
| receiver_flush_location  | text                     | Location where the receiver flushes logs.                                                                                      |
| receiver_replay_location | text                     | Location where the receiver replays logs.                                                                                      |
| sync_priority            | integer                  | Priority of synchronous duplication. (0 indicates asynchronization.)                                                           |

| Name       | Type | Description                                                                                                                                                                |
|------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sync_state | text | Synchronization state: <ul style="list-style-type: none"><li>• Asynchronous replication.</li><li>• Synchronous replication.</li><li>• Potential synchronization.</li></ul> |

### 13.2.11.2 GLOBAL\_REPLICATION\_STAT

GLOBAL\_REPLICATION\_STAT displays information about log synchronization status on each node, such as the locations where the sender sends logs and where the receiver receives logs.

**Table 13-138** GLOBAL\_REPLICATION\_STAT columns

| Name                     | Type                     | Description                                                                                                                    |
|--------------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| node_name                | name                     | Node name.                                                                                                                     |
| pid                      | bigint                   | PID of the thread.                                                                                                             |
| usesysid                 | oid                      | User system ID.                                                                                                                |
| username                 | name                     | Username.                                                                                                                      |
| application_name         | text                     | Program name.                                                                                                                  |
| client_addr              | inet                     | Client address.                                                                                                                |
| client_hostname          | text                     | Client name.                                                                                                                   |
| client_port              | integer                  | Port of the client.                                                                                                            |
| backend_start            | timestamp with time zone | Start time of the program.                                                                                                     |
| state                    | text                     | Log replication state: <ul style="list-style-type: none"><li>• Catch-up state.</li><li>• Consistent streaming state.</li></ul> |
| sender_sent_location     | text                     | Location where the sender sends logs.                                                                                          |
| receiver_write_location  | text                     | Location where the receiver writes logs.                                                                                       |
| receiver_flush_location  | text                     | Location where the receiver flushes logs.                                                                                      |
| receiver_replay_location | text                     | Location where the receiver replays logs.                                                                                      |

| Name          | Type    | Description                                                                                                                                                                |
|---------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sync_priority | integer | Priority of synchronous duplication. (0 indicates asynchronization.)                                                                                                       |
| sync_state    | text    | Synchronization state: <ul style="list-style-type: none"><li>• Asynchronous replication.</li><li>• Synchronous replication.</li><li>• Potential synchronization.</li></ul> |

### 13.2.11.3 REPLICATION\_SLOTS

**REPLICATION\_SLOTS** displays replication slot information.

**Table 13-139** REPLICATION\_SLOTS columns

| Name         | Type    | Description                                                                                                                                                                                                         |
|--------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| slot_name    | text    | Replication slot name.                                                                                                                                                                                              |
| plugin       | text    | Name of the output plug-in corresponding to the logical replication slot.                                                                                                                                           |
| slot_type    | text    | Replication slot type. <ul style="list-style-type: none"><li>• <b>physical</b>: physical replication slot.</li><li>• <b>logical</b>: logical replication slot.</li></ul>                                            |
| datoid       | oid     | OID of the database where the replication slot resides.                                                                                                                                                             |
| database     | name    | Name of the database where the replication slot resides.                                                                                                                                                            |
| active       | boolean | Determines whether the replication slot is activated. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                                                         |
| xmin         | xid     | The earliest transaction that the database must reserve for the replication slot. VACUUM cannot remove a tuple that is deleted by its subsequent transactions.                                                      |
| catalog_xmin | xid     | The earliest transaction that affects the system catalog and must be reserved by the database for the replication slot. VACUUM cannot remove a system catalog tuple that is deleted by its subsequent transactions. |
| restart_lsn  | text    | Physical location of the earliest Xlog required by the replication slot.                                                                                                                                            |

| Name          | Type    | Description                           |
|---------------|---------|---------------------------------------|
| dummy_standby | boolean | Not supported in the current version. |

#### 13.2.11.4 GLOBAL\_REPLICATION\_SLOTS

GLOBAL\_REPLICATION\_SLOTS displays replication slot information of each node in the database.

**Table 13-140** GLOBAL\_REPLICATION\_SLOTS columns

| Name          | Type    | Description                                                                                                                                                                                                         |
|---------------|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name     | name    | Node name.                                                                                                                                                                                                          |
| slot_name     | text    | Replication slot name.                                                                                                                                                                                              |
| plugin        | text    | Name of the output plug-in corresponding to the logical replication slot.                                                                                                                                           |
| slot_type     | text    | Replication slot type. <ul style="list-style-type: none"><li>• <b>physical</b>: physical replication slot.</li><li>• <b>logical</b>: logical replication slot.</li></ul>                                            |
| datoid        | oid     | OID of the database where the replication slot resides.                                                                                                                                                             |
| database      | name    | Name of the database where the replication slot resides.                                                                                                                                                            |
| active        | boolean | Determines whether the replication slot is activated. <ul style="list-style-type: none"><li>• <b>t</b> (true): yes</li><li>• <b>f</b> (false): no</li></ul>                                                         |
| x_min         | xid     | The earliest transaction that the database must reserve for the replication slot. VACUUM cannot remove a tuple that is deleted by its subsequent transactions.                                                      |
| catalog_xmin  | xid     | The earliest transaction that affects the system catalog and must be reserved by the database for the replication slot. VACUUM cannot remove a system catalog tuple that is deleted by its subsequent transactions. |
| restart_lsn   | text    | Physical location of the earliest Xlog required by the replication slot.                                                                                                                                            |
| dummy_standby | boolean | Not supported in the current version.                                                                                                                                                                               |

### 13.2.11.5 BGWRITER\_STAT

**BGWRITER\_STAT** displays statistics about the background writer process's activities.

**Table 13-141** BGWRITER\_STAT columns

| Name                  | Type                     | Description                                                                                                                          |
|-----------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| checkpoints_timed     | bigint                   | Number of scheduled checkpoints that have been performed                                                                             |
| checkpoints_req       | bigint                   | Number of requested checkpoints that have been performed                                                                             |
| checkpoint_write_time | double precision         | Total time spent on performing write operations during checkpointing (unit: ms).                                                     |
| checkpoint_sync_time  | double precision         | Total time spent on performing sync operations during checkpointing (unit: ms).                                                      |
| buffers_checkpoint    | bigint                   | Number of buffers written during checkpointing.                                                                                      |
| buffers_clean         | bigint                   | Number of buffers that have been written to a disk by the BGWRITER thread.                                                           |
| maxwritten_clean      | bigint                   | Number of times that the BGWRITER thread stops cleanup scanning because the number of pages refreshed in a single scan is too large. |
| buffers_backend       | bigint                   | Number of buffers written directly by the backend.                                                                                   |
| buffers_backend_fsync | bigint                   | Number of times that the backend thread automatically executes fsync. Generally, fsync is called by the BGWRITER thread.             |
| buffers_alloc         | bigint                   | Number of buffers allocated                                                                                                          |
| stats_reset           | timestamp with time zone | Time when the statistics of the current row were reset last time.                                                                    |

### 13.2.11.6 GLOBAL\_BGWRITER\_STAT

**GLOBAL\_BGWRITER\_STAT** displays statistics about the bgwriter thread's activities on all instances.

**Table 13-142** GLOBAL\_BGWRITER\_STAT columns

| Name                  | Type                     | Description                                                                                                                          |
|-----------------------|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| node_name             | name                     | Instance name.                                                                                                                       |
| checkpoints_timed     | bigint                   | Number of scheduled checkpoints that have been performed                                                                             |
| checkpoints_req       | bigint                   | Number of requested checkpoints that have been performed                                                                             |
| checkpoint_write_time | double precision         | Total time spent on performing write operations during checkpointing (unit: ms).                                                     |
| checkpoint_sync_time  | double precision         | Total time spent on performing sync operations during checkpointing (unit: ms).                                                      |
| buffers_checkpoint    | bigint                   | Number of buffers written during checkpointing.                                                                                      |
| buffers_clean         | bigint                   | Number of buffers that have been written to a disk by the BGWRITER thread.                                                           |
| maxwritten_clean      | bigint                   | Number of times that the BGWRITER thread stops cleanup scanning because the number of pages refreshed in a single scan is too large. |
| buffers_backend       | bigint                   | Number of buffers written directly by the backend.                                                                                   |
| buffers_backend_fsync | bigint                   | Number of times that the backend thread automatically executes fsync. Generally, fsync is called by the BGWRITER thread.             |
| buffers_alloc         | bigint                   | Number of buffers allocated                                                                                                          |
| stats_reset           | timestamp with time zone | Time when the statistics of the current row were reset last time.                                                                    |

### 13.2.11.7 GLOBAL\_CKPT\_STATUS

GLOBAL\_CKPT\_STATUS displays the information about checkpoints and flushing pages of all instances in the database.

**Table 13-143** GLOBAL\_CKPT\_STATUS columns

| Name                     | Type   | Description                                                                   |
|--------------------------|--------|-------------------------------------------------------------------------------|
| node_name                | text   | Instance name.                                                                |
| ckpt_redo_point          | text   | Checkpoint of the current instance.                                           |
| ckpt_clog_flush_num      | bigint | Number of Clog flushing pages from the startup time to the current time.      |
| ckpt_csnlog_flush_num    | bigint | Number of CSN log flushing pages from the startup time to the current time.   |
| ckpt_multixact_flush_num | bigint | Number of MultiXact flushing pages from the startup time to the current time. |
| ckpt_predicate_flush_num | bigint | Number of predicate flushing pages from the startup time to the current time. |
| ckpt_twophase_flush_num  | bigint | Number of two-phase flushing pages from the startup time to the current time. |

### 13.2.11.8 GLOBAL\_DOUBLE\_WRITE\_STATUS

GLOBAL\_DOUBLE\_WRITE\_STATUS displays doublewrite file status of all instances in the database.

**Table 13-144** GLOBAL\_DOUBLE\_WRITE\_STATUS columns

| Name                 | Type   | Description                                              |
|----------------------|--------|----------------------------------------------------------|
| node_name            | text   | Instance name.                                           |
| curr_dwn             | bigint | Sequence number of the doublewrite file                  |
| curr_start_page      | bigint | Start page for restoring the doublewrite file            |
| file_trunc_num       | bigint | Number of times that the doublewrite file is reused      |
| file_reset_num       | bigint | Number of reset times after the doublewrite file is full |
| total_writes         | bigint | Total number of I/Os of the doublewrite file             |
| low_threshold_writes | bigint | Number of I/Os for inefficient doublewrite files.        |



| Name                  | Type   | Description                                                         |
|-----------------------|--------|---------------------------------------------------------------------|
| high_threshold_writes | bigint | Number of I/Os for efficient doublewrite files.                     |
| total_pages           | bigint | Total number of pages that are flushed to the doublewrite file area |
| low_threshold_pages   | bigint | Number of pages that are flushed with low efficiency                |
| high_threshold_pages  | bigint | Number of pages that are flushed with high efficiency               |
| file_id               | bigint | ID of the current doublewrite file                                  |

### 13.2.11.9 GLOBAL\_PAGEWRITER\_STATUS

GLOBAL\_PAGEWRITER\_STATUS displays the page flushing information and checkpoint information about all instances in the database.

**Table 13-145** GLOBAL\_PAGEWRITER\_STATUS columns

| Name                        | Type    | Description                                                                                 |
|-----------------------------|---------|---------------------------------------------------------------------------------------------|
| node_name                   | text    | Instance name.                                                                              |
| pgwr_actual_flush_total_num | bigint  | Total number of dirty pages flushed from the startup time to the current time               |
| pgwr_last_flush_num         | integer | Number of dirty pages flushed in the previous batch                                         |
| remain_dirty_page_num       | bigint  | Estimated number of dirty pages that are not flushed                                        |
| queue_head_page_rec_lsn     | text    | <b>recovery_lsn</b> of the first dirty page in the dirty page queue of the current instance |
| queue_rec_lsn               | text    | <b>recovery_lsn</b> of the dirty page queue of the current instance                         |
| current_xlog_ckpt_lsn       | text    | The write position of Xlogs in the current instance                                         |
| ckpt_redo_point             | text    | Checkpoint of the current instance                                                          |

### 13.2.11.10 GLOBAL\_RECORD\_RESET\_TIME

GLOBAL\_RECORD\_RESET\_TIME is used to obtain statistics about the reset (restart, primary/standby switchover, and database deletion) time of each node in the database.

**Table 13-146** GLOBAL\_RECORD\_RESET\_TIME columns

| Name       | Type                     | Description       |
|------------|--------------------------|-------------------|
| node_name  | text                     | Node name.        |
| reset_time | timestamp with time zone | Time to be reset. |

### 13.2.11.11 GLOBAL\_REDO\_STATUS

GLOBAL\_REDO\_STATUS displays the log replay status of all instances in the database.

**Table 13-147** GLOBAL\_REDO\_STATUS columns

| Name                   | Type   | Description                                                                                    |
|------------------------|--------|------------------------------------------------------------------------------------------------|
| node_name              | text   | Instance name.                                                                                 |
| redo_start_ptr         | bigint | Start point for replaying the instance logs.                                                   |
| redo_start_time        | bigint | Start time (UTC) when the instance logs are replayed.                                          |
| redo_done_time         | bigint | End time (UTC) when the instance logs are replayed.                                            |
| curr_time              | bigint | Current time (UTC) of the instance.                                                            |
| min_recovery_point     | bigint | Minimum barrier for the current instance log to provide services after the replay is complete. |
| read_ptr               | bigint | Position for reading the instance logs.                                                        |
| last_replayed_read_ptr | bigint | Position for replaying the instance logs.                                                      |
| recovery_done_ptr      | bigint | Replay position after the instance is started.                                                 |
| read_xlog_io_counter   | bigint | Number of I/Os when the current instance reads and replays logs.                               |
| read_xlog_io_total_dur | bigint | Total I/O duration when the current instance reads and replays logs.                           |
| read_data_io_counter   | bigint | Number of data page I/O reads during log replay in the current instance.                       |
| read_data_io_total_dur | bigint | Total I/O duration reads during log replay in the current instance.                            |
| write_data_io_counter  | bigint | Number of data page I/O writes during log replay in the current instance.                      |

| Name                      | Type   | Description                                                                                                                                                                                                                                           |
|---------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| write_data_io_total_dur   | bigint | Total I/O duration writes during log replay in the current instance.                                                                                                                                                                                  |
| process_pending_counter   | bigint | Number of synchronization times of log distribution threads during log replay in the current instance.                                                                                                                                                |
| process_pending_total_dur | bigint | Synchronization duration of log distribution threads during log replay in the current instance.                                                                                                                                                       |
| apply_counter             | bigint | Number of synchronization times of replay threads during log replay in the current instance.                                                                                                                                                          |
| apply_total_dur           | bigint | Synchronization duration of replay threads during log replay in the current instance.                                                                                                                                                                 |
| speed                     | bigint | Log replay rate of the current instance. The value is updated every time when a 256-MB log is replayed. The unit is byte/s.<br>You are advised to run the <b>cm_ctl query -rv</b> command to obtain a more accurate replay speed of the standby node. |
| local_max_ptr             | bigint | Maximum number of replay logs received by the localhost after the instance is started.                                                                                                                                                                |
| primary_flush_ptr         | bigint | Log point where the host flushes logs to a disk.                                                                                                                                                                                                      |
| worker_info               | text   | Replay thread information of the instance. If parallel replay is not enabled, the value is <b>NULL</b> .                                                                                                                                              |

### 13.2.11.12 GLOBAL\_RECOVERY\_STATUS

GLOBAL\_RECOVERY\_STATUS displays log flow control information about the primary and standby nodes.

**Table 13-148** GLOBAL\_RECOVERY\_STATUS columns

| Name              | Type | Description                                          |
|-------------------|------|------------------------------------------------------|
| node_name         | text | Node name (including the primary and standby nodes). |
| standby_node_name | text | Name of the standby node.                            |
| source_ip         | text | IP address of the primary node.                      |

| Name               | Type    | Description                                                                                            |
|--------------------|---------|--------------------------------------------------------------------------------------------------------|
| source_port        | integer | Port number of the primary node.                                                                       |
| dest_ip            | text    | IP address of the standby node.                                                                        |
| dest_port          | integer | Port number of the standby node.                                                                       |
| current_rto        | bigint  | Current log flow control time of the standby node (unit: s).                                           |
| target_rto         | bigint  | Expected flow control time of the standby node specified by the corresponding GUC parameter (unit: s). |
| current_sleep_time | bigint  | Sleep time (in microseconds) required by the host to achieve the expected RTO.                         |

### 13.2.11.13 CLASS\_VITAL\_INFO

**CLASS\_VITAL\_INFO** is used to check whether the OIDs of the same table or index are consistent for WDR snapshots.

**Table 13-149** CLASS\_VITAL\_INFO columns

| Name       | Type   | Description                                                                                                                                                            |
|------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| relid      | oid    | Table OID.                                                                                                                                                             |
| schemaname | name   | Schema name.                                                                                                                                                           |
| relname    | name   | Table name.                                                                                                                                                            |
| relkind    | "char" | Object type. Its value can be: <ul style="list-style-type: none"><li>• <b>r</b>: ordinary table.</li><li>• <b>t</b>: TOAST table.</li><li>• <b>i</b>: index.</li></ul> |

### 13.2.11.14 USER\_LOGIN

**USER\_LOGIN** records the number of user logins and logouts.

**Table 13-150** USER\_LOGIN columns

| Name      | Type | Description           |
|-----------|------|-----------------------|
| node_name | text | Database process name |
| user_name | text | Username.             |

| Name           | Type    | Description                                                                   |
|----------------|---------|-------------------------------------------------------------------------------|
| user_id        | integer | User OID. (Its value is the same as that of <b>oid</b> in <b>pg_authid</b> .) |
| login_counter  | bigint  | Number of logins.                                                             |
| logout_counter | bigint  | Number of logouts.                                                            |

### 13.2.11.15 SUMMARY\_USER\_LOGIN

SUMMARY\_USER\_LOGIN records information about user logins and logouts on the primary database node.

**Table 13-151** SUMMARY\_USER\_LOGIN columns

| Name           | Type    | Description                                                                   |
|----------------|---------|-------------------------------------------------------------------------------|
| node_name      | text    | Database process name                                                         |
| user_name      | text    | Username.                                                                     |
| user_id        | integer | User OID. (Its value is the same as that of <b>oid</b> in <b>pg_authid</b> .) |
| login_counter  | bigint  | Number of logins.                                                             |
| logout_counter | bigint  | Number of logouts.                                                            |

### 13.2.11.16 GLOBAL\_SINGLE\_FLUSH\_DW\_STATUS

GLOBAL\_SINGLE\_FLUSH\_DW\_STATUS displays information about doublewrite files eliminated on a single page of all instances in the database. In the displayed information, the information before the slash (/) indicates the page flushing status of the first version, and the information after the slash (/) indicates the page flushing status of the second version.

**Table 13-152** GLOBAL\_SINGLE\_FLUSH\_DW\_STATUS columns

| Name            | Type | Description                                                       |
|-----------------|------|-------------------------------------------------------------------|
| node_name       | text | Instance name.                                                    |
| curr_dwn        | text | Sequence number of the doublewrite file.                          |
| curr_start_page | text | Start position of the current doublewrite file.                   |
| total_writes    | text | Total number of data write pages in the current doublewrite file. |
| file_trunc_num  | text | Number of times that the doublewrite file is reused.              |

| Name           | Type | Description                                               |
|----------------|------|-----------------------------------------------------------|
| file_reset_num | text | Number of reset times after the doublewrite file is full. |

### 13.2.11.17 GLOBAL\_CANDIDATE\_STATUS

GLOBAL\_CANDIDATE\_STATUS displays the number of candidate buffers and buffer eviction information of all instances in the database.

**Table 13-153** GLOBAL\_GET\_BGWRITER\_STATUS columns

| Name                    | Type    | Description                                                                                                                            |
|-------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------|
| node_name               | text    | Instance name.                                                                                                                         |
| candidate_slots         | integer | Number of pages in the candidate buffer chain of the current normal buffer pool.                                                       |
| get_buf_from_list       | bigint  | Number of times that pages are obtained from the candidate buffer chain during buffer eviction in the current normal buffer pool.      |
| get_buf_clock_sweep     | bigint  | Number of times that pages are obtained from the original eviction solution during buffer eviction in the current normal buffer pool.  |
| seg_candidate_slots     | integer | Number of pages in the candidate buffer chain of the current segment buffer pool.                                                      |
| seg_get_buf_from_list   | bigint  | Number of times that pages are obtained from the candidate buffer chain during buffer eviction in the current segment buffer pool.     |
| seg_get_buf_clock_sweep | bigint  | Number of times that pages are obtained from the original eviction solution during buffer eviction in the current segment buffer pool. |

## 13.2.12 Lock

### 13.2.12.1 LOCKS

LOCKS displays information about locks held by each open transaction.

**Table 13-154** LOCKS columns

| Name               | Type     | Description                                                                                                                                                                                                                        |
|--------------------|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| locktype           | text     | Type of the locked object: <b>relation</b> , <b>extend</b> , <b>page</b> , <b>tuple</b> , <b>transactionid</b> , <b>virtualxid</b> , <b>object</b> , <b>userlock</b> , or <b>advisory</b>                                          |
| database           | oid      | OID of the database in which the locked object exists. <ul style="list-style-type: none"><li>• The OID is <b>0</b> if the object is a shared object.</li><li>• The OID is <b>NULL</b> if the object is a transaction ID.</li></ul> |
| relation           | oid      | OID of the relationship targeted by the lock. The value is <b>NULL</b> if the object is not a relationship or part of a relationship.                                                                                              |
| page               | integer  | Page number targeted by the lock within the relationship. The value is <b>NULL</b> if the object is not a relationship page or row page.                                                                                           |
| tuple              | smallint | Row number targeted by the lock within the page. The value is <b>NULL</b> if the object is not a row.                                                                                                                              |
| bucket             | integer  | Hash bucket ID.                                                                                                                                                                                                                    |
| virtualxid         | text     | Virtual ID of the transaction targeted by the lock. The value is <b>NULL</b> if the object is not a virtual transaction ID.                                                                                                        |
| transactionid      | xid      | ID of the transaction targeted by the lock. The value is <b>NULL</b> if the object is not a transaction ID.                                                                                                                        |
| classid            | oid      | OID of the system catalog that contains the object. The value is <b>NULL</b> if the object is not a general database object.                                                                                                       |
| objid              | oid      | OID of the locked object within its system catalog. The value is <b>NULL</b> if the object is not a general database object.                                                                                                       |
| objsubid           | smallint | Column number for a column in the table ( <b>0</b> if the object is of other object type and <b>NULL</b> if the object is not a general database object)                                                                           |
| virtualtransaction | text     | Virtual ID of the transaction holding or awaiting this lock                                                                                                                                                                        |
| pid                | bigint   | Logical ID of the server thread holding or awaiting this lock. The value is <b>NULL</b> if the lock is held by a prepared transaction.                                                                                             |

| Name             | Type    | Description                                                                                                                                                            |
|------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sessionid        | bigint  | ID of the session holding or awaiting this lock<br>The value is <b>NULL</b> if the lock is held by a prepared transaction.                                             |
| mode             | text    | Lock mode held or desired by this thread                                                                                                                               |
| granted          | Boolean | <ul style="list-style-type: none"><li>The value is <b>TRUE</b> if the lock is a held lock.</li><li>The value is <b>FALSE</b> if the lock is an awaited lock.</li></ul> |
| fastpath         | Boolean | The value is <b>TRUE</b> if the lock is obtained through <b>fast-path</b> , and is <b>FALSE</b> if the lock is obtained through the main lock table.                   |
| locktag          | text    | Lock information that the session waits for. It can be parsed using the <b>locktag_decode()</b> function.                                                              |
| global_sessionid | text    | Global session ID                                                                                                                                                      |

### 13.2.12.2 GLOBAL\_LOCKS

**GLOBAL\_LOCKS** displays information about locks held by open transactions on each node.

**Table 13-155** GLOBAL\_LOCKS columns

| Name      | Type    | Description                                                                                                                                                                                                                    |
|-----------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name | name    | Node name.                                                                                                                                                                                                                     |
| locktype  | text    | Type of the locked object: <b>relation</b> , <b>extend</b> , <b>page</b> , <b>tuple</b> , <b>transactionid</b> , <b>virtualxid</b> , <b>object</b> , <b>userlock</b> , or <b>advisory</b>                                      |
| database  | oid     | OID of the database in which the locked object exists. <ul style="list-style-type: none"><li>The OID is <b>0</b> if the object is a shared object.</li><li>The OID is <b>NULL</b> if the object is a transaction ID.</li></ul> |
| relation  | oid     | OID of the relationship targeted by the lock.<br>The value is <b>NULL</b> if the object is not a relationship or part of a relationship.                                                                                       |
| page      | integer | Page number targeted by the lock within the relationship. The value is <b>NULL</b> if the object is not a relationship page or row page.                                                                                       |



| Name               | Type     | Description                                                                                                                                                                |
|--------------------|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tuple              | smallint | Row number targeted by the lock within the page. The value is <b>NULL</b> if the object is not a row.                                                                      |
| virtualxid         | text     | Virtual ID of the transaction targeted by the lock. The value is <b>NULL</b> if the object is not a virtual transaction ID.                                                |
| transactionid      | xid      | ID of the transaction targeted by the lock. The value is <b>NULL</b> if the object is not a transaction ID.                                                                |
| classid            | oid      | OID of the system catalog that contains the object. The value is <b>NULL</b> if the object is not a general database object.                                               |
| objid              | oid      | OID of the locked object within its system catalog. The value is <b>NULL</b> if the object is not a general database object.                                               |
| objsubid           | smallint | Column number for a column in the table ( <b>0</b> if the target is of other object type and <b>NULL</b> if the object is not a general database object)                   |
| virtualtransaction | text     | Virtual ID of the transaction holding or awaiting this lock                                                                                                                |
| pid                | bigint   | Logical ID of the server thread holding or awaiting this lock The value is <b>NULL</b> if the lock is held by a prepared transaction.                                      |
| mode               | text     | Lock mode held or desired by this thread                                                                                                                                   |
| granted            | Boolean  | <ul style="list-style-type: none"><li>• The value is <b>TRUE</b> if the lock is a held lock.</li><li>• The value is <b>FALSE</b> if the lock is an awaited lock.</li></ul> |
| fastpath           | Boolean  | The value is <b>TRUE</b> if the lock is obtained through <b>fast-path</b> , and is <b>FALSE</b> if the lock is obtained through the main lock table.                       |

## 13.2.13 Wait Events

### 13.2.13.1 WAIT\_EVENTS

**WAIT\_EVENTS** displays statistics about wait events on the current node. For details about the impact of each transaction lock on services, see [LOCK](#).

**Table 13-156** WAIT\_EVENTS columns

| Name            | Type                     | Description                           |
|-----------------|--------------------------|---------------------------------------|
| nodename        | text                     | Database process name.                |
| type            | text                     | Event type.                           |
| event           | text                     | Event name.                           |
| wait            | bigint                   | Number of waiting times.              |
| failed_wait     | bigint                   | Number of waiting failures.           |
| total_wait_time | bigint                   | Total waiting time (unit: $\mu$ s).   |
| avg_wait_time   | bigint                   | Average waiting time (unit: $\mu$ s). |
| max_wait_time   | bigint                   | Maximum waiting time (unit: $\mu$ s). |
| min_wait_time   | bigint                   | Minimum waiting time (unit: $\mu$ s). |
| last_updated    | timestamp with time zone | Last time when the event was updated. |

### 13.2.13.2 GLOBAL\_WAIT\_EVENTS

**GLOBAL\_WAIT\_EVENTS** displays statistics about wait events on each node.

**Table 13-157** GLOBAL\_WAIT\_EVENTS columns

| Name            | Type                     | Description                          |
|-----------------|--------------------------|--------------------------------------|
| nodename        | text                     | Database process name                |
| type            | text                     | Event type                           |
| event           | text                     | Event name                           |
| wait            | bigint                   | Number of waiting times              |
| failed_wait     | bigint                   | Number of waiting failures           |
| total_wait_time | bigint                   | Total waiting time (unit: $\mu$ s)   |
| avg_wait_time   | bigint                   | Average waiting time (unit: $\mu$ s) |
| max_wait_time   | bigint                   | Maximum waiting time (unit: $\mu$ s) |
| min_wait_time   | bigint                   | Minimum waiting time (unit: $\mu$ s) |
| last_updated    | timestamp with time zone | Last time when the event was updated |

## 13.2.14 Configuration

### 13.2.14.1 CONFIG\_SETTINGS

**CONFIG\_SETTINGS** displays information about parameters of the running database.

**Table 13-158** CONFIG\_SETTINGS columns

| Name       | Type   | Description                                                                                                                                                      |
|------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name       | text   | Parameter name                                                                                                                                                   |
| setting    | text   | Current parameter value                                                                                                                                          |
| unit       | text   | Implicit unit of the parameter                                                                                                                                   |
| category   | text   | Logical group of the parameter                                                                                                                                   |
| short_desc | text   | Brief description of the parameter                                                                                                                               |
| extra_desc | text   | Detailed description of the parameter                                                                                                                            |
| context    | text   | Context required to set the parameter value, including <b>internal</b> , <b>postmaster</b> , <b>sighup</b> , <b>backend</b> , <b>superuser</b> , and <b>user</b> |
| vartype    | text   | Parameter type, including <b>bool</b> , <b>enum</b> , <b>integer</b> , <b>real</b> , or <b>string</b>                                                            |
| source     | text   | Method of assigning the parameter value                                                                                                                          |
| min_val    | text   | Maximum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                  |
| max_val    | text   | Minimum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                  |
| enumvals   | text[] | Valid values of an enum-type parameter. If the parameter type is not enum, the value of this column is <b>null</b> .                                             |
| boot_val   | text   | Default parameter value used upon the database startup                                                                                                           |
| reset_val  | text   | Default parameter value used upon the database reset                                                                                                             |
| sourcefile | text   | Configuration file used to set parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> .  |

| Name       | Type    | Description                                                                                                                                                                       |
|------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| sourceline | integer | Row number in the configuration file for setting parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> . |

### 13.2.14.2 GLOBAL\_CONFIG\_SETTINGS

**GLOBAL\_CONFIG\_SETTINGS** displays information about parameters of running databases on each node.

**Table 13-159** GLOBAL\_CONFIG\_SETTINGS columns

| Name       | Type   | Description                                                                                                                                                      |
|------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name  | text   | Node name                                                                                                                                                        |
| name       | text   | Parameter name                                                                                                                                                   |
| setting    | text   | Current parameter value                                                                                                                                          |
| unit       | text   | Implicit unit of the parameter                                                                                                                                   |
| category   | text   | Logical group of the parameter                                                                                                                                   |
| short_desc | text   | Brief description of the parameter                                                                                                                               |
| extra_desc | text   | Detailed description of the parameter                                                                                                                            |
| context    | text   | Context required to set the parameter value, including <b>internal</b> , <b>postmaster</b> , <b>sighup</b> , <b>backend</b> , <b>superuser</b> , and <b>user</b> |
| vartype    | text   | Parameter type, including <b>bool</b> , <b>enum</b> , <b>integer</b> , <b>real</b> , or <b>string</b>                                                            |
| source     | text   | Method of assigning the parameter value                                                                                                                          |
| min_val    | text   | Minimum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                  |
| max_val    | text   | Maximum value of the parameter. If the parameter type is not numeric, the value of this column is <b>null</b> .                                                  |
| enumvals   | text[] | Valid values of an enum-type parameter. If the parameter type is not enum, the value of this column is <b>null</b> .                                             |
| boot_val   | text   | Default parameter value used upon the database startup                                                                                                           |

| Name       | Type    | Description                                                                                                                                                                       |
|------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| reset_val  | text    | Default parameter value used upon the database reset                                                                                                                              |
| sourcefile | text    | Configuration file used to set parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> .                   |
| sourceline | integer | Row number in the configuration file for setting parameter values. If parameter values are not configured using the configuration file, the value of this column is <b>null</b> . |

## 13.2.15 Operator

### 13.2.15.1 OPERATOR\_HISTORY\_TABLE

OPERATOR\_HISTORY\_TABLE displays records about operators of completed jobs. Data is dumped from the kernel to this system view.

**Table 13-160** OPERATOR\_HISTORY\_TABLE columns

| Name            | Type                     | Description                                                            |
|-----------------|--------------------------|------------------------------------------------------------------------|
| queryid         | bigint                   | Internal query ID used for statement execution.                        |
| pid             | bigint                   | Backend thread ID.                                                     |
| plan_node_id    | integer                  | Plan node ID of the execution plan of a query.                         |
| plan_node_name  | text                     | Name of the operator corresponding to the plan node ID.                |
| start_time      | timestamp with time zone | Time when the operator starts to process the first data record.        |
| duration        | bigint                   | Total execution time of the operator (unit: ms).                       |
| query_dop       | integer                  | DOP of the operator.                                                   |
| estimated_rows  | bigint                   | Number of rows estimated by the optimizer.                             |
| tuple_processed | bigint                   | Number of elements returned by the operator.                           |
| min_peak_memory | integer                  | Minimum peak memory used by the operator on database nodes (unit: MB). |
| max_peak_memory | integer                  | Maximum peak memory used by the operator on database nodes (unit: MB). |

| Name                | Type    | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| average_peak_memory | integer | Average peak memory used by the operator on database nodes (unit: MB).                                                                                                                                                                                                                                                                                 |
| memory_skew_percent | integer | Memory usage skew of the operator among database nodes.                                                                                                                                                                                                                                                                                                |
| min_spill_size      | integer | Minimum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0).                                                                                                                                                                                                                                                           |
| max_spill_size      | integer | Maximum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0).                                                                                                                                                                                                                                                           |
| average_spill_size  | integer | Average spilled data among database nodes when a spill occurs (unit: MB) (default value: 0).                                                                                                                                                                                                                                                           |
| spill_skew_percent  | integer | Database node spill skew when a spill occurs.                                                                                                                                                                                                                                                                                                          |
| min_cpu_time        | bigint  | Minimum execution time of the operator on database nodes (unit: ms).                                                                                                                                                                                                                                                                                   |
| max_cpu_time        | bigint  | Maximum execution time of the operator on database nodes (unit: ms).                                                                                                                                                                                                                                                                                   |
| total_cpu_time      | bigint  | Total execution time of the operator on database nodes (unit: ms).                                                                                                                                                                                                                                                                                     |
| cpu_skew_percent    | integer | Execution time skew among database nodes.                                                                                                                                                                                                                                                                                                              |
| warning             | text    | Warning. The following warnings are displayed: <ul style="list-style-type: none"><li>• Sort/SetOp/HashAgg/HashJoin spill</li><li>• Spill file size large than 256MB</li><li>• Broadcast size large than 100MB</li><li>• Early spill</li><li>• Spill times is greater than 3</li><li>• Spill on memory adaptive</li><li>• Hash table conflict</li></ul> |

### 13.2.15.2 OPERATOR\_HISTORY

OPERATOR\_HISTORY displays records of operators in jobs that have been executed by the current user on the current primary database node. Columns in this view are the same as those in GS\_WLM\_OPERATOR\_INFO.

### 13.2.15.3 OPERATOR\_RUNTIME

**OPERATOR\_RUNTIME** displays information about operators of the jobs that are being executed by the current user.

**Table 13-161** OPERATOR\_RUNTIME columns

| Name                | Type                     | Description                                                                                 |
|---------------------|--------------------------|---------------------------------------------------------------------------------------------|
| queryid             | bigint                   | Internal query ID used for statement execution                                              |
| pid                 | bigint                   | Thread ID of the backend                                                                    |
| plan_node_id        | integer                  | Plan node ID of the execution plan of a query                                               |
| plan_node_name      | text                     | Name of the operator corresponding to the plan node ID                                      |
| start_time          | timestamp with time zone | Time when the operator starts to process the first data record                              |
| duration            | bigint                   | Total execution time of the operator, in ms                                                 |
| status              | text                     | Execution status of the current operator, which can be <b>finished</b> or <b>running</b> .  |
| query_dop           | integer                  | DOP of the operator                                                                         |
| estimated_rows      | bigint                   | Number of rows estimated by the optimizer                                                   |
| tuple_processed     | bigint                   | Number of elements returned by the operator                                                 |
| min_peak_memory     | integer                  | Minimum peak memory used by the operator on database nodes (unit: MB)                       |
| max_peak_memory     | integer                  | Maximum peak memory used by the operator on database nodes (unit: MB)                       |
| average_peak_memory | integer                  | Average peak memory used by the operator on database nodes (unit: MB)                       |
| memory_skew_percent | integer                  | Memory usage skew of the operator among database nodes                                      |
| min_spill_size      | integer                  | Minimum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |
| max_spill_size      | integer                  | Maximum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |
| average_spill_size  | integer                  | Average spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |

| Name               | Type    | Description                                                                                                                                                                                                                                                                                                                                            |
|--------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| spill_skew_percent | integer | Database node spill skew when a spill occurs                                                                                                                                                                                                                                                                                                           |
| min_cpu_time       | bigint  | Minimum execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                    |
| max_cpu_time       | bigint  | Maximum execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                    |
| total_cpu_time     | bigint  | Total execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                      |
| cpu_skew_percent   | integer | Execution time skew among database nodes                                                                                                                                                                                                                                                                                                               |
| warning            | text    | Warning. The following warnings are displayed: <ul style="list-style-type: none"><li>• Sort/SetOp/HashAgg/HashJoin spill</li><li>• Spill file size large than 256MB</li><li>• Broadcast size large than 100MB</li><li>• Early spill</li><li>• Spill times is greater than 3</li><li>• Spill on memory adaptive</li><li>• Hash table conflict</li></ul> |

#### 13.2.15.4 GLOBAL\_OPERATOR\_HISTORY

**GLOBAL\_OPERATOR\_HISTORY** displays records of operators in jobs that have been executed by the current user on the primary database node.

**Table 13-162** GLOBAL\_OPERATOR\_HISTORY columns

| Name           | Type                     | Description                                                    |
|----------------|--------------------------|----------------------------------------------------------------|
| queryid        | bigint                   | Internal query ID used for statement execution                 |
| pid            | bigint                   | Thread ID of the backend                                       |
| plan_node_id   | integer                  | Plan node ID of the execution plan                             |
| plan_node_name | text                     | Name of the operator corresponding to the plan node ID         |
| start_time     | timestamp with time zone | Time when the operator starts to process the first data record |
| duration       | bigint                   | Total execution time of the operator, in ms                    |
| query_dop      | integer                  | DOP of the operator                                            |



| Name                | Type    | Description                                                                                 |
|---------------------|---------|---------------------------------------------------------------------------------------------|
| estimated_rows      | bigint  | Number of rows estimated by the optimizer                                                   |
| tuple_processed     | bigint  | Number of elements returned by the operator                                                 |
| min_peak_memory     | integer | Minimum peak memory used by the operator on database nodes (unit: MB)                       |
| max_peak_memory     | integer | Maximum peak memory used by the operator on database nodes (unit: MB)                       |
| average_peak_memory | integer | Average peak memory used by the operator on database nodes (unit: MB)                       |
| memory_skew_percent | integer | Memory usage skew of the operator among database nodes                                      |
| min_spill_size      | integer | Minimum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |
| max_spill_size      | integer | Maximum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |
| average_spill_size  | integer | Average spilled data among database nodes when a spill occurs (unit: MB) (default value: 0) |
| spill_skew_percent  | integer | Database node spill skew when a spill occurs                                                |
| min_cpu_time        | bigint  | Minimum execution time of the operator on database nodes (unit: ms)                         |
| max_cpu_time        | bigint  | Maximum execution time of the operator on database nodes (unit: ms)                         |
| total_cpu_time      | bigint  | Total execution time of the operator on database nodes (unit: ms)                           |
| cpu_skew_percent    | integer | Execution time skew among database nodes                                                    |

| Name    | Type | Description                                                                                                                                                                                                                                                                        |
|---------|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| warning | text | Warning. The following warnings are displayed:<br>1. Sort/SetOp/HashAgg/HashJoin spill<br>2. Spill file size large than 256MB<br>3. Broadcast size large than 100MB<br>4. Early spill<br>5. Spill times is greater than 3<br>6. Spill on memory adaptive<br>7. Hash table conflict |

### 13.2.15.5 GLOBAL\_OPERATOR\_HISTORY\_TABLE

**GLOBAL\_OPERATOR\_HISTORY\_TABLE** displays the records about operators of completed jobs on the primary database node. Data is dumped from the kernel to the system catalog **GS\_WLM\_OPERATOR\_INFO**.

**GLOBAL\_OPERATOR\_HISTORY\_TABLE** is a collection view for querying the system catalog **GS\_WLM\_OPERATOR\_INFO** on the primary database node. Columns in this view are the same as those in [Table 13-162](#).

### 13.2.15.6 GLOBAL\_OPERATOR\_RUNTIME

**GLOBAL\_OPERATOR\_RUNTIME** displays information about operators of the jobs that are being executed by the current user on the primary database node.

**Table 13-163** GLOBAL\_OPERATOR\_RUNTIME columns

| Name           | Type                     | Description                                                                                |
|----------------|--------------------------|--------------------------------------------------------------------------------------------|
| queryid        | bigint                   | Internal query ID used for statement execution                                             |
| pid            | bigint                   | Thread ID of the backend                                                                   |
| plan_node_id   | integer                  | Plan node ID of the execution plan                                                         |
| plan_node_name | text                     | Name of the operator corresponding to the plan node ID                                     |
| start_time     | timestamp with time zone | Time when the operator starts to process the first data record                             |
| duration       | bigint                   | Total execution time of the operator, in ms                                                |
| status         | text                     | Execution status of the current operator, which can be <b>finished</b> or <b>running</b> . |
| query_dop      | integer                  | DOP of the operator                                                                        |
| estimated_rows | bigint                   | Number of rows estimated by the optimizer                                                  |

| Name                | Type    | Description                                                                                                                                                                                                                                                                                                                                            |
|---------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| tuple_processed     | bigint  | Number of elements returned by the operator                                                                                                                                                                                                                                                                                                            |
| min_peak_memory     | integer | Minimum peak memory used by the operator on database nodes (unit: MB)                                                                                                                                                                                                                                                                                  |
| max_peak_memory     | integer | Maximum peak memory used by the operator on database nodes (unit: MB)                                                                                                                                                                                                                                                                                  |
| average_peak_memory | integer | Average peak memory used by the operator on database nodes (unit: MB)                                                                                                                                                                                                                                                                                  |
| memory_skew_percent | integer | Memory usage skew of the operator among database nodes                                                                                                                                                                                                                                                                                                 |
| min_spill_size      | integer | Minimum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0)                                                                                                                                                                                                                                                            |
| max_spill_size      | integer | Maximum spilled data among database nodes when a spill occurs (unit: MB) (default value: 0)                                                                                                                                                                                                                                                            |
| average_spill_size  | integer | Average spilled data among database nodes when a spill occurs (unit: MB) (default value: 0)                                                                                                                                                                                                                                                            |
| spill_skew_percent  | integer | Database node spill skew when a spill occurs                                                                                                                                                                                                                                                                                                           |
| min_cpu_time        | bigint  | Minimum execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                    |
| max_cpu_time        | bigint  | Maximum execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                    |
| total_cpu_time      | bigint  | Total execution time of the operator on database nodes (unit: ms)                                                                                                                                                                                                                                                                                      |
| cpu_skew_percent    | integer | Execution time skew among database nodes                                                                                                                                                                                                                                                                                                               |
| warning             | text    | Warning. The following warnings are displayed: <ul style="list-style-type: none"><li>• Sort/SetOp/HashAgg/HashJoin spill</li><li>• Spill file size large than 256MB</li><li>• Broadcast size large than 100MB</li><li>• Early spill</li><li>• Spill times is greater than 3</li><li>• Spill on memory adaptive</li><li>• Hash table conflict</li></ul> |

## 13.2.16 Workload Manager

### 13.2.16.1 WLM\_USER\_RESOURCE\_CONFIG

**WLM\_USER\_RESOURCE\_CONFIG** displays the resource configuration information of a user.

**Table 13-164** WLM\_USER\_RESOURCE\_CONFIG columns

| Name       | Type    | Description                                             |
|------------|---------|---------------------------------------------------------|
| userid     | oid     | OID of the user                                         |
| username   | name    | Username                                                |
| sysadmin   | boolean | Specifies whether the user has the SYSADMIN permission. |
| rpoid      | oid     | OID of the resource pool                                |
| respool    | name    | Name of a resource pool                                 |
| parentid   | oid     | OID of the parent user                                  |
| totalspace | bigint  | Size of the occupied space                              |
| spacelimit | bigint  | Upper limit of the space size                           |
| childcount | integer | Number of child users                                   |
| childlist  | text    | Child user list                                         |

### 13.2.16.2 WLM\_USER\_RESOURCE\_RUNTIME

**WLM\_USER\_RESOURCE\_RUNTIME** displays resource usage of all users. Only administrators can query this view. This view is valid only when the GUC parameter **use\_workload\_manager** is set to **on**.

**Table 13-165** WLM\_USER\_RESOURCE\_RUNTIME columns

| Name         | Type    | Description                                                                                                                                                    |
|--------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| username     | name    | Username                                                                                                                                                       |
| used_memory  | integer | Used memory (unit: MB)                                                                                                                                         |
| total_memory | integer | Available memory (unit: MB) The value <b>0</b> indicates that the available memory is not limited and depends on the maximum memory available in the database. |
| used_cpu     | integer | Number of CPU cores in use                                                                                                                                     |

| Name              | Type    | Description                                                                                                                                         |
|-------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| total_cpu         | integer | Total number of CPU cores of the Cgroup associated with the user on the node                                                                        |
| used_space        | bigint  | Used storage space (unit: KB)                                                                                                                       |
| total_space       | bigint  | Available storage space (unit: KB) The value -1 indicates that the maximum storage space is not limited.                                            |
| used_temp_space   | bigint  | Used temporary space (reserved column; unit: KB)                                                                                                    |
| total_temp_space  | bigint  | Available temporary storage space (reserved column; unit: KB) The value -1 indicates that the maximum temporary storage space is not limited.       |
| used_spill_space  | bigint  | Used space for storing spilled data (reserved column; unit: KB)                                                                                     |
| total_spill_space | bigint  | Available storage space for spilled data (reserved column; unit: KB) The value -1 indicates that the maximum space for spilled data is not limited. |

## 13.2.17 Global Plancache

Global plan cache (GPC) views are valid only when **enable\_global\_plancache** and the thread pool are enabled.

### 13.2.17.1 GLOBAL\_PLANCACHE\_STATUS

**GLOBAL\_PLANCACHE\_STATUS** displays the GPC status information.

**Table 13-166** GLOBAL\_PLANCACHE\_STATUS columns

| Name        | Type    | Description                                       |
|-------------|---------|---------------------------------------------------|
| nodename    | text    | Name of the node that the plan cache belongs to   |
| query       | text    | Text of query statements                          |
| refcount    | integer | Number of times that the plan cache is referenced |
| valid       | bool    | Specifies whether it is valid.                    |
| databaseid  | oid     | ID of the database that the plan cache belongs to |
| schema_name | text    | Schema that the plan cache belongs to             |

| Name       | Type    | Description                                                                                                                                    |
|------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------|
| params_num | integer | Number of parameters                                                                                                                           |
| func_id    | oid     | OID of the stored procedure where the plan cache is located. If the plancache does not belong to the stored procedure, the value is <b>0</b> . |

### 13.2.17.2 GLOBAL\_PLANCACHE\_CLEAN

**GLOBAL\_PLANCACHE\_CLEAN** clears the global plan cache that is not used on all nodes. The return value is of the Boolean type.

## 13.2.18 RTO & RPO

### 13.2.18.1 global\_rto\_status

Displays log flow control information about the primary and standby nodes (except the current node and standby DN).

**Table 13-167** global\_rto\_status columns

| Parameter | Type | Description                                                                                                                                                                                                                                                               |
|-----------|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| node_name | text | Node name (including the primary and standby nodes)                                                                                                                                                                                                                       |
| rto_info  | text | Flow control information, including the current log flow control time (unit: second) of the standby node, the expected flow control time (unit: second) specified by the GUC parameter, and the primary node sleep time (unit: $\mu$ s) required to reach the expectation |

### 13.2.18.2 global\_streaming\_hadr\_rto\_and\_rpo\_stat

**global\_streaming\_hadr\_rto\_and\_rpo\_stat** displays the log flow control information about the primary and standby database instances in the streaming DR scenario. (This schema can be used only on the primary DN of the primary database instance. Statistics cannot be obtained from the standby DN or the standby database instance.)

**Table 13-168** Parameters

| Parameter               | Type | Description                                                                                                                                   |
|-------------------------|------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| hadr_sender_node_name   | text | Node name, including the primary database instance and the main standby node of the standby database instance.                                |
| hadr_receiver_node_name | text | Name of the main standby node of the standby database instance.                                                                               |
| current_rto             | int  | Flow control information, that is, log RTO time of the current primary and standby database instances (unit: second).                         |
| target_rto              | int  | Flow control information, that is, the RTO time between the target primary and standby database instances (unit: second).                     |
| current_rpo             | int  | Flow control information, that is, log RPO time of the current primary and standby database instances (unit: second).                         |
| target_rpo              | int  | Flow control information, that is, the RPO time between the target primary and standby database instances (unit: second).                     |
| rto_sleep_time          | int  | RTO flow control information, that is, the expected sleep time (unit: $\mu$ s) required by walsender on the host to reach the specified RTO.  |
| rpo_sleep_time          | int  | RPO flow control information, that is, the expected sleep time (unit: $\mu$ s) required by xlogInsert on the host to reach the specified RPO. |

## 13.3 DBE\_PLDEBUGGER Schema

DBE\_PLDEBUGGER Schema system functions are used to debug stored procedures. This chapter describes the APIs supported by DBE\_PLDEBUGGER Schema. Only the administrator has the permission to execute these debugging APIs, but does not have the permission to modify or create functions in this schema. You can debug only non-system functions in the public schema or user-created schema. You are not allowed to debug system functions.

### NOTICE

When a user is created in the function body, the plaintext password is returned when attach, next, continue, info\_code, step, info\_breakpoint, backtrace or finish is called. You are advised not to create a user in the function body.

The administrator can run the following command to grant the **gs\_role\_pldebugger** role and debugger permissions to a user:

```
GRANT gs_role_pldebugger to user;
```

Two clients are required to connect to the database. One client is responsible for executing the debugging API as the debug end, and the other client is responsible for executing the debugging function to control the execution of stored procedures on the server. Example:

- Prepare for debugging.

Use PG\_PROC to find the OID of the stored procedure to be debugged and execute DBE\_PLDEBUGGER.turn\_on(oid). In this case, the client functions as the server.

```
openGauss=# CREATE OR REPLACE PROCEDURE test_debug (IN x INT)
AS
BEGIN
 INSERT INTO t1 (a) VALUES (x);
 DELETE FROM t1 WHERE a = x;
END;
/
CREATE PROCEDURE
openGauss=# SELECT OID FROM PG_PROC WHERE PRONAME='test_debug';
oid

16389
(1 row)
openGauss=# SELECT * FROM DBE_PLDEBUGGER.turn_on(16389);
nodename | port
-----+-----
datanode | 0
(1 row)
```

- Start debugging.

When the server executes the stored procedure, the server hangs before the first SQL statement in the stored procedure and waits for the debugging message sent by the debug end. Debugging is supported only by directly executing a stored procedure and cannot be achieved by calling an executed stored procedure through a trigger.

```
openGauss=# call test_debug(1);
```

Start another client as the debug end and call DBE\_PLDEBUGGER.attach to attach with the stored procedure for debugging based on the data returned by **turn\_on**.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.attach('datanode',0);
funcoid | funcname | lineno | query
-----+-----+-----+-----
16389 | test_debug | 3 | INSERT INTO t1 (a) VALUES (x);
(1 row)
```

Execute the next statement on the client where the attach operation is performed.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.next();
funcoid | funcname | lineno | query
-----+-----+-----+-----
16389 | test_debug | 0 | [EXECUTION FINISHED]
(1 row)
```

Execute the following variable operations on the client where the attach command is performed.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.info_locals(); -- Print all variables.
varname | vartype | value | package_name | isconst
-----+-----+-----+-----+-----
x | int4 | 1 | | f
(1 row)
openGauss=# SELECT * FROM DBE_PLDEBUGGER.set_var('x', 2); -- Assign a value to a variable.
```



```
set_var

t
(1 row)
openGauss=# SELECT * FROM DBE_PLDEBUGGER.print_var('x'); -- Print a single variable.
varname | vartype | value | package_name | isconst
-----+-----+-----+-----+-----
x | int4 | 2 | | f
(1 row)
```

Directly execute the stored procedure that is being debugged.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.continue();
funcoid | funcname | lineno | query
-----+-----+-----+-----
16389 | test_debug | 0 | [EXECUTION FINISHED]
(1 row)
```

Exit the stored procedure that is being debugged and do not execute statements that have not been executed before.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.abort();
abort

t
(1 row)
```

View the code information on the client and identify the line number of the breakpoint that can be set.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.info_code(16389);
lineno | query | canbreak
-----+-----+-----
1 | CREATE OR REPLACE PROCEDURE public.test_debug(IN x INT) | f
2 | AS DECLARE | f
3 | BEGIN | f
4 | INSERT INTO t1 (a) VALUES (x); | t
5 | DELETE FROM t1 WHERE a = x; | t
6 | END; | f
7 | / | f
(7 rows)
```

Sets a breakpoint.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.add_breakpoint(16389,4);
lineno | query | canbreak
-----+-----+-----
1 | CREATE OR REPLACE PROCEDURE public.test_debug(IN x INT) | f
2 | AS DECLARE | f
3 | BEGIN | f
4 | INSERT INTO t1 (a) VALUES (x); | t
5 | DELETE FROM t1 WHERE a = x; | t
6 | END; | f
7 | / | f
(7 rows)
```

View the breakpoint information.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.info_breakpoints();
breakpointno | funcoid | lineno | query | enable
-----+-----+-----+-----+-----
0 | 16389 | 4 | DELETE FROM t1 WHERE a = x; | t
(1 row)
```

Execute to the breakpoint.

```
openGauss=# SELECT * FROM DBE_PLDEBUGGER.continue();
funcoid | funcname | lineno | query
-----+-----+-----+-----
16389 | test_debug | 4 | DELETE FROM t1 WHERE a = x;
(1 row)
```

After the stored procedure is executed, the debugging automatically exits. To debug the stored procedure again, you need to attach again. If the server

does not need to be debugged, run the **turn\_off** command to disable the debugging or exit the session. For details about the debugging APIs, see the following table.

**Table 13-169** DBE\_PLDEBUGGER

| API                                                    | Description                                                                                                                                                          |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">DBE_PLDEBUGGER.turn_on</a>                 | Called by the server, indicating that the stored procedure can be debugged. After the API is called, the stored procedure is hung to wait for debugging information. |
| <a href="#">DBE_PLDEBUGGER.turn_off</a>                | Called by the server, indicating that debugging the stored procedure is disabled.                                                                                    |
| <a href="#">DBE_PLDEBUGGER.local_debug_server_info</a> | Called by the server to print all stored procedures that have been turned on in the current session.                                                                 |
| <a href="#">DBE_PLDEBUGGER.attach</a>                  | Called by the debug end to attach with the stored procedure that is being debugged.                                                                                  |
| <a href="#">DBE_PLDEBUGGER.info_locals</a>             | Called by the debug end to print the current values of variables in the stored procedure that is being debugged.                                                     |
| <a href="#">DBE_PLDEBUGGER.next</a>                    | Called by the debug end to execute the next step.                                                                                                                    |
| <a href="#">DBE_PLDEBUGGER.continue</a>                | Called by the debug end to continue the execution until the breakpoint or stored procedure ends.                                                                     |
| <a href="#">DBE_PLDEBUGGER.abort</a>                   | Called by the debug end to stop debugging. The server reports a long jump error.                                                                                     |
| <a href="#">DBE_PLDEBUGGER.print_var</a>               | Called by the debug end to print the current values of specified variables in the stored procedure that is being debugged.                                           |
| <a href="#">DBE_PLDEBUGGER.info_code</a>               | Called by the debug end or server to print the source statement of a specified stored procedure and the line number corresponding to each line.                      |
| <a href="#">DBE_PLDEBUGGER.step</a>                    | Called by the debug end to execute step by step.                                                                                                                     |
| <a href="#">DBE_PLDEBUGGER.add_breakpoint</a>          | Called by the debug end to add a breakpoint.                                                                                                                         |
| <a href="#">DBE_PLDEBUGGER.delete_breakpoint</a>       | Called by the debug end to delete a breakpoint.                                                                                                                      |
| <a href="#">DBE_PLDEBUGGER.info_breakpoints</a>        | Called by the debug end to view all breakpoints.                                                                                                                     |

| API                                                    | Description                                                                                                                  |
|--------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">DBE_PLDEBUGGER.bac<br/>ktrace</a>          | Called by the debug end to check the current call stack.                                                                     |
| <a href="#">DBE_PLDEBUGGER.ena<br/>ble_breakpoint</a>  | Called by the debug end to enable breakpoints.                                                                               |
| <a href="#">DBE_PLDEBUGGER.disa<br/>ble_breakpoint</a> | Called by the debug end to disable breakpoints.                                                                              |
| <a href="#">DBE_PLDEBUGGER.finis<br/>h</a>             | Called by the debug end to continue the debugging until the breakpoint is reached or the upper-layer call stack is returned. |
| <a href="#">DBE_PLDEBUGGER.set_<br/>var</a>            | Called by the debug end to assign a value to a variable.                                                                     |

### 13.3.1 DBE\_PLDEBUGGER.turn\_on

This function is used to mark a stored procedure as debuggable. After **turn\_on** is executed, the server can execute the stored procedure for debugging. You need to manually obtain the OID of the stored procedure based on the PG\_PROC system catalog and transfer it to the function. After **turn\_on** is executed, the execution of the stored procedure in the current session is hung before the first SQL statement to wait for the debugging instruction from the debug end. This setting is cleared by default after the session is disconnected. Currently, stored procedures and functions with autonomous transactions enabled cannot be debugged.

The function prototype is as follows:

```
DBE_PLDEBUGGER.turn_on(Oid)
RETURN Record;
```

**Table 13-170** turn\_on input parameters and return values

| Name     | Type        | Description                  |
|----------|-------------|------------------------------|
| func_oid | IN oid      | Function OID                 |
| nodename | OUT text    | Node name                    |
| port     | OUT integer | Number of the connected port |

### 13.3.2 DBE\_PLDEBUGGER.turn\_off

This function is used to remove the debugging flag added by **turn\_on**. The return value indicates success or failure. You can run the **DBE\_PLDEBUGGER.local\_debug\_server\_info** command to query the OID of the stored procedure that has been turned on.

The function prototype is as follows:

```
DBE_PLDEBUGGER.turn_off(Oid)
RETURN boolean;
```

**Table 13-171** turn\_off input parameters and return values

| Name     | Type        | Description                               |
|----------|-------------|-------------------------------------------|
| func_oid | IN oid      | Function OID                              |
| turn_off | OUT boolean | Specifies whether turn-off is successful. |

### 13.3.3 DBE\_PLDEBUGGER.local\_debug\_server\_info

This function is used to query the OID of the stored procedure that has been turned on in the current connection. You can use **funcoid** and **pg\_proc** together to determine which stored procedures are to be debugged.

**Table 13-172** local\_debug\_server\_info returned input parameters and return values

| Name     | Type       | Description          |
|----------|------------|----------------------|
| nodename | OUT text   | Node name            |
| port     | OUT bigint | Port number          |
| funcoid  | OUT oid    | Stored procedure OID |

### 13.3.4 DBE\_PLDEBUGGER.attach

When the server executes a stored procedure, the server hangs the execution before the first statement and waits for attaching with the debug end. The debug end calls the attach function and transfers node name and port number to attach with the specified stored procedure.

If an error is reported during debugging, the attach operation automatically becomes invalid. If the debug end is attached to another stored procedure during debugging, the debugging of the attached stored procedure becomes invalid.

**Table 13-173** attach input parameters and return values

| Name     | Type       | Description                  |
|----------|------------|------------------------------|
| nodename | IN text    | Node name                    |
| port     | IN integer | Number of the connected port |
| funcoid  | OUT oid    | Function ID                  |

| Name     | Type        | Description                                                         |
|----------|-------------|---------------------------------------------------------------------|
| funcname | OUT text    | Function name                                                       |
| lineno   | OUT integer | Number of the next line in the current debugging process            |
| query    | OUT text    | Source code of the next line of the function that is being debugged |

### 13.3.5 DBE\_PLDEBUGGER.info\_locals

During debugging on the debug end, **info\_locals** is called to print the variables in the current stored procedure. The input parameter **frameno** of this function indicates the stack layer to be traversed. This function can be called without input parameters. By default, the top-layer stack variable is queried.

**Table 13-174** info\_locals input parameters and return values

| Name         | Type                  | Description                                                                                             |
|--------------|-----------------------|---------------------------------------------------------------------------------------------------------|
| frameno      | IN integer (optional) | Specified stack layer. The default value is the top layer.                                              |
| varname      | OUT text              | Variable name                                                                                           |
| vartype      | OUT text              | Variable type                                                                                           |
| value        | OUT text              | Variable value                                                                                          |
| package_name | OUT text              | Name of the package corresponding to the variable. If the variable is not a package, the value is null. |
| isconst      | OUT boolean           | Specifies whether it is a constant.                                                                     |

### 13.3.6 DBE\_PLDEBUGGER.next

This function is used to execute the current SQL statement in a stored procedure and return the number of the next SQL statement and the corresponding query.

**Table 13-175** next input parameters and return values

| Name     | Type        | Description                                                         |
|----------|-------------|---------------------------------------------------------------------|
| funcoid  | OUT oid     | Function ID                                                         |
| funcname | OUT text    | Function name                                                       |
| lineno   | OUT integer | Number of the next line in the current debugging process            |
| query    | OUT text    | Source code of the next line of the function that is being debugged |

### 13.3.7 DBE\_PLDEBUGGER.continue

Executes the current stored procedure until reaching the next breakpoint or end, and returns the line number of the next execution and the corresponding query.

The function prototype is as follows:

```
DBE_PLDEBUGGER.continue()
RETURN Record;
```

**Table 13-176** continue input parameters and return values

| Name     | Type        | Description                                                         |
|----------|-------------|---------------------------------------------------------------------|
| funcoid  | OUT oid     | Function ID                                                         |
| funcname | OUT text    | Function name                                                       |
| lineno   | OUT integer | Number of the next line in the current debugging process            |
| query    | OUT text    | Source code of the next line of the function that is being debugged |

### 13.3.8 DBE\_PLDEBUGGER.abort

This function is used to abort the stored procedure executed on the server and report an error. The return value specifies whether the abort message is successfully sent.

The function prototype is as follows:

```
DBE_PLDEBUGGER.abort()
RETURN boolean;
```

**Table 13-177** abort input parameters and return values

| Name  | Type        | Description        |
|-------|-------------|--------------------|
| abort | OUT boolean | Success or failure |

### 13.3.9 DBE\_PLDEBUGGER.print\_var

During debugging on the debug end, **print\_var** is called to print the name and value of the specified variable in the current stored procedure. The input parameter **frameno** of this function indicates the stack layer to be traversed. This function can be called without **frameno**. By default, the top-layer stack variable is queried.

**Table 13-178** print\_var input parameters and return values

| Name         | Type                  | Description                                                                                         |
|--------------|-----------------------|-----------------------------------------------------------------------------------------------------|
| var_name     | IN text               | Variable                                                                                            |
| frameno      | IN integer (optional) | Specified stack layer. The default value is the top layer.                                          |
| varname      | OUT text              | Variable name                                                                                       |
| vartype      | OUT text              | Variable type                                                                                       |
| value        | OUT text              | Variable value                                                                                      |
| package_name | OUT text              | Package name corresponding to the variable. This parameter is reserved and is left empty currently. |
| isconst      | OUT boolean           | Specifies whether it is a constant.                                                                 |

### 13.3.10 DBE\_PLDEBUGGER.info\_code

During debugging on the debug end, **info\_code** is called to view the source statement of the specified stored procedure and the line number corresponding to each line. The line number starts from the function body, and the line number in the function header is empty.

**Table 13-179** info\_code input parameters and return values

| Name    | Type   | Description |
|---------|--------|-------------|
| funcoid | IN oid | Function ID |

| Name     | Type        | Description                                              |
|----------|-------------|----------------------------------------------------------|
| lineno   | OUT integer | Line number                                              |
| query    | OUT text    | Source statement                                         |
| canbreak | OUT bool    | Specifies whether the current line supports breakpoints. |

### 13.3.11 DBE\_PLDEBUGGER.step

During debugging on the debug end, if a stored procedure is being executed, the stored procedure continues to be executed and information such as the line number in the first line of the stored procedure is returned. If the executed object is not a stored procedure, the return is the same as that for **next**. After the SQL statement is executed, information such as the line number in the next line is returned.

**Table 13-180** step input parameters and return values

| Name     | Type        | Description                                                         |
|----------|-------------|---------------------------------------------------------------------|
| funcoid  | OUT oid     | Function ID                                                         |
| funcname | OUT text    | Function name                                                       |
| lineno   | OUT integer | Number of the next line in the current debugging process            |
| query    | OUT text    | Source code of the next line of the function that is being debugged |

### 13.3.12 DBE\_PLDEBUGGER.add\_breakpoint

During debugging on the debug end, call **add\_breakpoint** to add a breakpoint. If -1 is returned, the specified breakpoint is invalid. Determine the proper position of the breakpoint based on the **canbreak** column in [DBE\\_PLDEBUGGER.info\\_code](#).

**Table 13-181** add\_breakpoint input parameters and return values

| Name         | Type        | Description       |
|--------------|-------------|-------------------|
| funcoid      | IN text     | Function ID       |
| lineno       | IN integer  | Line number       |
| breakpointno | OUT integer | Breakpoint number |



### 13.3.13 DBE\_PLDEBUGGER.delete\_breakpoint

During debugging on the debug end, call **delete\_breakpoint** to delete the existing breakpoint.

**Table 13-182** delete\_breakpoint input parameters and return values

| Name         | Type       | Description                                     |
|--------------|------------|-------------------------------------------------|
| breakpointno | IN integer | Breakpoint number                               |
| result       | OUT bool   | Specifies whether this operation is successful. |

### 13.3.14 DBE\_PLDEBUGGER.info\_breakpoints

During debugging on the debug end, call **info\_breakpoints** to view the current function breakpoint.

**Table 13-183** info\_breakpoints input parameters and return values

| Name         | Type        | Description        |
|--------------|-------------|--------------------|
| breakpointno | OUT integer | Breakpoint number  |
| funcoid      | OUT oid     | Function ID        |
| lineno       | OUT integer | Line number        |
| query        | OUT text    | Breakpoint content |
| enable       | OUT boolean | Valid or not       |

### 13.3.15 DBE\_PLDEBUGGER.backtrace

During debugging on the debug end, call **backtrace** to view the current call stack.

**Table 13-184** backtrace input parameters and return values

| Name     | Type        | Description        |
|----------|-------------|--------------------|
| frameno  | OUT integer | Call stack ID      |
| funcname | OUT text    | Function name      |
| lineno   | OUT integer | Line number        |
| query    | OUT text    | Breakpoint content |
| funcoid  | OUT oid     | Function OID       |

### 13.3.16 DBE\_PLDEBUGGER.enable\_breakpoint

During debugging on the debug end, calls enable\_breakpoint to enable breakpoints.

**Table 13-185** enable\_breakpoint input parameters and return values

| Name         | Type       | Description                                     |
|--------------|------------|-------------------------------------------------|
| breakpointno | IN integer | Breakpoint number                               |
| result       | OUT bool   | Specifies whether this operation is successful. |

### 13.3.17 DBE\_PLDEBUGGER.disable\_breakpoint

During debugging on the debug end, call disable\_breakpoint to disable breakpoints.

**Table 13-186** disable\_breakpoint input parameters and return values

| Name         | Type       | Description                                     |
|--------------|------------|-------------------------------------------------|
| breakpointno | IN integer | Breakpoint number                               |
| result       | OUT bool   | Specifies whether this operation is successful. |

### 13.3.18 DBE\_PLDEBUGGER.finish

Executes the current SQL statement in the stored procedure until the next breakpoint is triggered or the next line of the upper-layer stack is executed.

**Table 13-187** finish input parameters and return values

| Name     | Type        | Description                                                         |
|----------|-------------|---------------------------------------------------------------------|
| funcoid  | OUT oid     | Function ID                                                         |
| funcname | OUT text    | Function name                                                       |
| lineno   | OUT integer | Number of the next line in the current debugging process            |
| query    | OUT text    | Source code of the next line of the function that is being debugged |

### 13.3.19 DBE\_PLDEBUGGER.set\_var

Changes the variable on the top-layer stack in the specified debugging stored procedure to the value of the input parameter. If a stored procedure contains variables with the same name, set\_var supports only the setting of the first variable value.

**Table 13-188** set\_var input parameters and return values

| Name     | Type        | Description                 |
|----------|-------------|-----------------------------|
| var_name | IN text     | Variable name               |
| value    | IN text     | New value                   |
| result   | OUT boolean | Result (success or failure) |

## 13.4 DBE\_PLDEVELOPER

The DBE\_PLDEVELOPER system catalog records information required for compiling PL/SQL packages, functions, and stored procedures.

### 13.4.1 DBE\_PLDEVELOPER.gs\_source

Records PL/SQL object (stored procedure, function, package, and package body) compilation information. For details, see the following column description.

When the **plsql\_show\_all\_error** parameter is enabled, information about PL/SQL object compilation success or failure is recorded in this table. When the **plsql\_show\_all\_error** parameter is disabled, only information about correct compilation is inserted into this table.

---

**CAUTION**

1. The gs\_source table records only user-defined original object statements. Even if a user uses ALTER to change the created schema or name, the information in the gs\_source table does not change. If the user changes the schema or name of an object, the deleted object still exists in the gs\_source table.
2. The owner in the gs\_source table is the user who creates the table, not the user specified when the user creates the stored procedure or package.
3. By default, row-level security is not configured for the gs\_source table in the database. If you want to use the database isolation feature, run the following statement to add row-level security:  

```
ALTER TABLE dbe_pldeveloper.gs_source ENABLE ROW LEVEL SECURITY;
CREATE ROW LEVEL SECURITY POLICY all_data_rls ON dbe_pldeveloper.gs_source USING(owner =
(select oid from pg_roles where rolname=current_user));
```

**Table 13-189** DBE\_PLDEVELOPER.gs\_source columns

| Name   | Type    | Description                                                     |
|--------|---------|-----------------------------------------------------------------|
| id     | oid     | Object ID.                                                      |
| owner  | bigint  | ID of the user who creates the object.                          |
| nspid  | oid     | Schema ID of an object.                                         |
| name   | name    | Object name.                                                    |
| type   | text    | Object type ( <b>procedure/function/package/package body</b> ). |
| status | boolean | Determines whether the creation is successful.                  |
| src    | text    | Original statement for creating an object.                      |

### 13.4.2 DBE\_PLDEVELOPER.gs\_errors

Records errors that occur during PL/SQL object (stored procedure, function, package, and package body) compilation. For details, see the following column description.

After the **plsql\_show\_all\_error** parameter is enabled, if an error occurs during compilation, the error is skipped and the compilation continues, and the error information is recorded in **gs\_errors**. If the **plsql\_show\_all\_error** parameter is disabled, related information is not inserted into this table.

The owner of the table is the user who creates the table. Modifying the owner of the stored procedure or package does not modify the table information.

**Table 13-190** DBE\_PLDEVELOPER.gs\_errors columns

| Name  | Type    | Description                                                     |
|-------|---------|-----------------------------------------------------------------|
| id    | oid     | Object ID.                                                      |
| owner | bigint  | ID of the user who creates the object.                          |
| nspid | oid     | Schema ID of an object.                                         |
| name  | name    | Object name.                                                    |
| type  | text    | Object type ( <b>procedure/function/package/package body</b> ). |
| line  | integer | Line number.                                                    |
| src   | text    | Error message.                                                  |

## 13.5 DBE\_SQL\_UTIL Schema

The DBE\_SQL\_UTIL schema stores tools for managing SQL patches, including creating, deleting, enabling, and disabling SQL patches. Common users have only the USAGE permission and do not have the CREATE, ALTER, DROP, and COMMENT permissions.

### 13.5.1 DBE\_SQL\_UTIL.create\_hint\_sql\_patch

create\_hint\_sql\_patch is a function used to create SQL patches for hints. It returns whether the execution is successful.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-191** DBE\_SQL\_UTIL.create\_hint\_sql\_patch input parameters and return values

| Parameter     | Type      | Description                                                                  |
|---------------|-----------|------------------------------------------------------------------------------|
| patch_name    | IN name   | Patch name.                                                                  |
| unique_sql_id | IN bigint | Global unique query ID.                                                      |
| hint_string   | IN text   | Hint text.                                                                   |
| description   | IN text   | Patch description. The default value is <b>NULL</b> .                        |
| enabled       | IN bool   | Specifies whether the patch takes effect. The default value is <b>true</b> . |
| result        | OUT bool  | Specifies whether this operation is successful.                              |

### 13.5.2 DBE\_SQL\_UTIL.create\_abort\_sql\_patch

create\_abort\_sql\_patch is a function used to create abort SQL patches. It returns whether the execution is successful.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-192** DBE\_SQL\_UTIL.create\_abort\_sql\_patch input parameters and return values

| Parameter  | Type    | Description |
|------------|---------|-------------|
| patch_name | IN name | Patch name. |

| Parameter     | Type      | Description                                                                  |
|---------------|-----------|------------------------------------------------------------------------------|
| unique_sql_id | IN bigint | Global unique ID.                                                            |
| description   | IN text   | Patch description. The default value is <b>NULL</b> .                        |
| enabled       | IN bool   | Specifies whether the patch takes effect. The default value is <b>true</b> . |
| result        | OUT bool  | Specifies whether this operation is successful.                              |

### 13.5.3 DBE\_SQL\_UTIL.drop\_sql\_patch

drop\_sql\_patch is a function used to delete SQL patches. It returns whether the execution is successful.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-193** DBE\_SQL\_UTIL.drop\_sql\_patch input parameters and return values

| Parameter  | Type     | Description                                     |
|------------|----------|-------------------------------------------------|
| patch_name | IN name  | Patch name.                                     |
| result     | OUT bool | Specifies whether this operation is successful. |

### 13.5.4 DBE\_SQL\_UTIL.enable\_sql\_patch

enable\_sql\_patch is a function used to enable SQL patches. It returns whether the execution is successful.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-194** DBE\_SQL\_UTIL.enable\_sql\_patch input parameters and return values

| Parameter  | Type     | Description                                     |
|------------|----------|-------------------------------------------------|
| patch_name | IN name  | Patch name.                                     |
| result     | OUT bool | Specifies whether this operation is successful. |

### 13.5.5 DBE\_SQL\_UTIL.disable\_sql\_patch

disable\_sql\_patch is a function used to disable SQL patches. It returns whether the execution is successful.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-195** DBE\_SQL\_UTIL.disable\_sql\_patch input parameters and return values

| Parameter  | Type     | Description                                     |
|------------|----------|-------------------------------------------------|
| patch_name | IN name  | Patch name.                                     |
| result     | OUT bool | Specifies whether this operation is successful. |

### 13.5.6 DBE\_SQL\_UTIL.show\_sql\_patch

show\_sql\_patch is a function used to display the SQL patch corresponding to a specified patch name and return the running result.

Only the initial user, SYSADMIN, OPRADMIN, and MONADMIN have the permission to call it.

**Table 13-196** DBE\_SQL\_UTIL.show\_sql\_patch input parameters and return values

| Parameter     | Type       | Description                                   |
|---------------|------------|-----------------------------------------------|
| patch_name    | IN name    | Patch name.                                   |
| unique_sql_id | OUT bigint | Global unique ID.                             |
| enabled       | OUT bool   | Specifies whether the patch takes effect.     |
| abort         | OUT bool   | Specifies whether the patch is an abort hint. |
| hint_str      | OUT text   | Hint text.                                    |

# 14 Logical Replication


## 14.1 Logical Decoding

The data replication capabilities supported by GaussDB are as follows:

Data is periodically synchronized to heterogeneous databases (such as database A) using a data migration tool. Real-time data replication is not supported. Therefore, the requirements for real-time data synchronization to heterogeneous databases are not satisfied.

GaussDB provides the logical decoding function to generate logical logs by decoding Xlogs. A target database parses logical logs to replicate data in real time. For details, see [Figure 14-1](#). Logical replication reduces the restrictions on target databases, allowing for data synchronization between heterogeneous databases and homogeneous databases with different forms. It allows data to be read and written during data synchronization on a target database, reducing the data synchronization latency.

**Figure 14-1** Logical replication



Logical replication consists of logical decoding and data replication. Logical decoding outputs logical logs by transaction. The database service or middleware parses the logical logs to implement data replication. Currently, GaussDB supports only logical decoding. Therefore, this section involves only logical decoding.



## 14.1.1 Overview

### Function

Logical decoding provides basic transaction decoding capabilities for logical replication. GaussDB uses SQL functions for logical decoding. This method features easy function calling, requires no tools to obtain logical logs, and provides specific APIs for interconnecting with external replay tools, saving the need of additional adaptation.

Logical logs are output only after transactions are committed because they use transactions as the unit and logical decoding is driven by users. Therefore, to prevent Xlogs from being recycled by the system when transactions start and prevent required transaction information from being recycled by VACUUM, GaussDB introduces logical replication slots to block Xlog recycling.

A logical replication slot represents a stream of changes that can be re-executed in other databases in the order they were generated in the original database. Each logical replication slot is maintained by the person who obtains the corresponding logical logs.

### Prerequisites

- Currently, logical logs are extracted from DNs. If logical replication is performed, SSL connections must be used. Therefore, ensure that the **ssl** parameter on the corresponding DN is set to **on**.

#### NOTE

For security purposes, ensure that SSL connections are enabled.

- The GUC parameter **wal\_level** is set to **logical**.
- The GUC parameter **max\_replication\_slots** is set to a value greater than or equal to the sum of physical replication slots, backup slots, and logical replication slots required by each DN.

#### NOTE

Plan the number of logical replication slots as follows:

- A logical replication slot can carry changes of only one database for decoding. If multiple databases are involved, create multiple logical replication slots.
- If logical replication is needed by multiple target databases, create multiple logical replication slots in the source database. Each logical replication slot corresponds to one logical replication link.
- A user needs to connect to a database through a DN port before using SQL functions to perform logical decoding. If a CN port is used to connect to the database, **EXECUTE DIRECT ON (datanode\_name) 'statement'** is needed to execute SQL functions.
- Only the initial user and users with the REPLICATION permission can perform this operation. When separation of duties is disabled, database administrators can perform logical replication operations. When separation of duties is enabled, database administrators are not allowed to perform logical replication operations.
- Sufficient memory resources must be configured for parallel decoding. The memory usage consists of the following parts:

- Memory for concatenating TOAST tuples and hash tables = Number of concurrent transactions x Average tuple size. The current version is not controlled by the flushing logic. A large amount of memory may be occupied when TOAST tuples are processed.
- Queue memory = Queue length x Decoding parallelism degree x Average tuple size x 2 (including the read log queue and decoding result queue). The queue memory is controlled by the queue length and decoding parallelism degree.
- Memory for sending and re-ordering parallel decoding logs = Number of concurrent transactions x Number of tuples modified by transactions x Average tuple size. The upper limit can be controlled by the decoding configuration options **max-txn-in-memory** and **max-reorderbuffer-in-memory**.
- Memory for storing transaction metadata = Number of concurrent transactions x Metadata size of each transaction (about 300 bytes).

## Precautions

- Logical decoding does not support DDL statements.
- Decoded data may be lost when a specific DDL statement (for example, to truncate an ordinary table or exchange data between partitioned tables) is executed.
- Decoding of DML operations on data page replication is not supported.
- After a DDL statement (for example, ALTER TABLE) is executed, the physical logs that are not decoded before the DDL statement execution may be lost.
- The size of a single tuple cannot exceed 1 GB, and decoded data may be larger than inserted data. Therefore, it is recommended that the size of a single tuple be less than or equal to 500 MB.
- GaussDB supports the following data types for decoding: INTEGER, BIGINT, SMALLINT, TINYINT, SERIAL, SMALLSERIAL, BIGSERIAL, FLOAT, DOUBLE PRECISION, DATE, TIME[WITHOUT TIME ZONE], TIMESTAMP[WITHOUT TIME ZONE], CHAR(n), VARCHAR(n), and TEXT.
- The name of a logical replication slot can contain no more than 64 characters, including lowercase letters, digits, underscores (\_), question marks (?), hyphens (-), and periods (.). In addition, periods (. or ..) cannot be used as the name of a logical replication slot independently.
- After the database where a logical replication slot resides is deleted, the replication slot becomes unavailable and needs to be manually deleted.
- To decode multiple databases, you need to create a streaming replication slot in each database and start decoding. Logs need to be scanned for decoding of each database.
- Forcible switchover is not supported. After forcible switchover, you need to export all data again.
- To perform decoding on the standby node, set the GUC parameter **enable\_slot\_log** to **on** on the corresponding host.
- During decoding on the standby node, the decoded data may increase during switchover and failover, which needs to be manually filtered out. When the quorum protocol is used, switchover and failover should be performed on the standby node that is to be promoted to primary, and logs must be synchronized from the primary node to the standby node.

- The same replication slot for decoding cannot be used between the primary node and standby node or between different standby nodes at the same time. Otherwise, data inconsistency occurs.
- Replication slots can only be created or deleted on hosts.
- After the database is restarted due to a fault or the logical replication process is restarted, duplicate decoded data may exist. You need to filter out the duplicate data.
- If the computer kernel is faulty, garbled characters may be displayed during decoding, which need to be manually or automatically filtered out.
- Currently, the logical decoding on the standby node does not support enabling the ultimate RTO.
- Ensure that the long transaction is not started during the creation of the logical replication slot. If the long transaction is started, the creation of the logical replication slot will be blocked.
- Interval partitioned tables cannot be replicated.
- Decoding of DML operations on global temporary tables is not supported.
- After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not decoded.
- Do not perform operations on the replication slot on other nodes when the logical replication slot is in use. To delete a replication slot, stop decoding in the replication slot first.
- To parse the UPDATE and DELETE statements of an Astore table, you need to configure the **REPLICA IDENTITY** attribute for the table. If the table does not have a primary key, set the **REPLICA IDENTITY** attribute to **FULL**.
- Do not perform operations on the replication slot on other nodes when the logical replication slot is in use. To delete a replication slot, stop decoding in the replication slot first.
- Considering that the target database may require the system status information of the source database, logical decoding automatically filters only logical logs of system catalogs whose OIDs are less than 16384 in pg\_catalog and pg\_toast schemas. If the target database does not need to copy the content of other related system catalogs, the related system catalogs need to be filtered during logical log replay.
- When logical replication is enabled, if you need to create a primary key index that contains system columns, you must set the **REPLICA IDENTITY** attribute of the table to **FULL** or use USING INDEX to specify a unique, non-local, non-deferrable index that does not contain system columns and contains only columns marked **NOT NULL**.
- After a logical replication slot is used up, delete it in time. Otherwise, Xlog recycling will be blocked.
- If a transaction has too many sub-transactions, too many files are flushed to disks. To exit decoding, you need to run the SQL function pg\_terminate\_backend (walsender thread ID for logical decoding) to manually stop decoding. In addition, the exit delay increases by about 1 minute per 300,000 sub-transactions. Therefore, when logical decoding is enabled, if the number of sub-transactions of a transaction reaches 50,000, a WARNING log is generated.
- When a transaction generates a large number of sub-transactions that need to be flushed to disks, the number of opened file handles may exceed the

upper limit. In this case, set **max\_files\_per\_process** to a value greater than twice the upper limit of sub-transactions.

## SQL Function Decoding Performance

In the Benchmarksql-5.0 with 100 warehouses, when `pg_logical_slot_get_changes` is used:

- If 4000 lines of data (about 5 MB to 10 MB logs) are decoded at a time, the decoding performance ranges from 0.3 MB/s to 0.5 MB/s.
- If 32000 lines of data (about 40 MB to 80 MB logs) are decoded at a time, the decoding performance ranges from 3 MB/s to 5 MB/s.
- If 256000 lines of data (about 320 MB to 640 MB logs) are decoded at a time, the decoding performance ranges from 3 MB/s to 5 MB/s.
- If the amount of data to be decoded at a time still increases, the decoding performance is not significantly improved.

If `pg_logical_slot_peek_changes` and `pg_replication_slot_advance` are used, the decoding performance is 30% to 50% lower than that when `pg_logical_slot_get_changes` is used.

### 14.1.2 Logical Decoding Options

Logical decoding options can provide a restriction on or additional functions for the current logical decoding, for example, specifying whether the decoding result includes a transaction number or whether empty transactions are ignored during decoding. For details about the configuration method and SQL function decoding, see the optional input parameters **options\_name** and **options\_value** of the `pg_logical_slot_peek_changes` function in [Logical Replication Functions](#). For details about JDBC streaming decoding, see the usage of the `withSlotOption` function in the sample code in [Logical Replication](#).

**General Options (Both serial decoding and parallel decoding can be configured, but the settings may be invalid. For details, see the description of related options.)**

- **include-xids:**  
Specifies whether the decoded **data** column contains XID information.  
Valid value: **0** and **1**. The default value is **1**.
  - **0**: The decoded **data** column does not contain XID information.
  - **1**: The decoded **data** column contains XID information.
- **skip-empty-xacts:**  
Specifies whether to ignore empty transaction information during decoding.  
Valid value: **0** and **1**. The default value is **0**.
  - **0**: The empty transaction information is not ignored during decoding.
  - **1**: The empty transaction information is ignored during decoding.
- **include-timestamp:**  
Specifies whether decoded information contains the **commit** timestamp.  
Valid value: **0** and **1**. The default value is **0**.

- **0**: The decoded information does not contain the **commit** timestamp.
  - **1**: The decoded information contains the **commit** timestamp.
- **only-local**:  
Specifies whether to decode only local logs.  
Valid value: **0** and **1**. The default value is **1**.
  - **0**: Non-local logs and local logs are decoded.
  - **1**: Only local logs are decoded.
- **white-table-list**:  
Whitelist parameter, including the schema and table name to be decoded.  
Value range: a string that contains table names in the whitelist. Different tables are separated by commas (,). An asterisk (\*) is used to fuzzily match all tables. Schema names and table names are separated by periods (.). No space character is allowed. For example:

```
select * from pg_logical_slot_peek_changes('slot1', NULL, 4096, 'white-table-list',
'public.t1,public.t2,*,t3,my_schema.*');
```
- **max-txn-in-memory**:  
Memory control parameter. The unit is MB. If the memory occupied by a single transaction is greater than the value of this parameter, data is flushed to disks.  
Value range: an integer ranging from 0 to 100. The default value is **0**, indicating that memory control is disabled.
- **max-reorderbuffer-in-memory**:  
Memory control parameter. The unit is GB. If the total memory (including the cache) of transactions being concatenated in the sender thread is greater than the value of this parameter, the current decoding transaction is flushed to disks.  
Value range: an integer ranging from 0 to 100. The default value is **0**, indicating that memory control is disabled.
- **include-user**:  
Specifies whether the BEGIN logical log of a transaction records the username of the transaction. The username of a transaction refers to the authorized user, that is, the login user who executes the session corresponding to the transaction. The username does not change during the execution of the transaction.  
Valid value: **0** and **1**. The default value is **0**.
  - **0**: The BEGIN logical log of a transaction does not contain the username of the transaction.
  - **1**: The BEGIN logical log of a transaction records the username of the transaction.
- **exclude-userids**:  
Specifies the OID of a blacklisted user.  
Value range: a string, which specifies the OIDs of blacklisted users. Multiple OIDs are separated by commas (,). The system does not check whether the OIDs exist.
- **exclude-users**:  
Name list of blacklisted users.

Value range: a string, which specifies the names of blacklisted users. Multiple names are separated by commas (,). The system does not check whether the names exist.

- **dynamic-resolution:**

Specifies whether to dynamically parse the names of blacklisted users.

Valid value: **0** and **1**. The default value is **1**.

- **0**: If the parameter is set to **0**, an error is reported and the logical decoding exits when the decoding detects that the user does not exist in blacklist **exclude-users**.
- **1**: If the parameter is set to **1**, decoding continues when it detects that the user does not exist in blacklist **exclude-users**.

- **standby-connection:**

Specifies whether to restrict decoding only on the standby node. This option is valid only for streaming decoding.

Value range: Boolean. The default value is **false**.

- **true**: Only the standby node can be connected for decoding. When the primary node is connected for decoding, an error is reported and the system exits.
- **false**: The primary or standby node can be connected for decoding.

- **sender-timeout:**

Specifies the heartbeat timeout threshold between the kernel and the client. This option is valid only for streaming decoding. If no message is received from the client within the period, the logical decoding stops and disconnects from the client. The unit is ms.

Value range: an integer ranging from 0 to 2147483647. The default value depends on the value of the GUC parameter **logical\_sender\_timeout**.

- **parallel-decode-num:**

Specifies the number of decoder threads for parallel decoding. This option is valid only for streaming decoding. When the system function is called, this option is invalid and only the value range is verified.

Value range: The value **1** indicates that decoding is performed based on the original serial logic. Other values indicate that parallel decoding is enabled. The default value is **1**.

---

**NOTICE**

If **parallel-decode-num** is not set (the default value is **1**) or is explicitly set to **1**, the options in the following "Parallel decoding" cannot be configured.

---

## Serial Decoding

- **force-binary:**

Specifies whether to output the decoding result in binary format and display different behaviors in different scenarios.

- For system functions `pg_logical_slot_get_binary_changes` and `pg_logical_slot_peek_binary_changes`:

- Value range: Boolean. The default value is **false**. The value is meaningless. The decoding result is always output in binary format.
- For system functions `pg_logical_slot_get_changes`, `pg_logical_slot_peek_changes`, and `pg_logical_get_area_changes`:  
Value range: Boolean. The value is fixed to **false**. The decoding result is always output in text format.
  - For streaming decoding:  
Value range: Boolean. The default value is **false**. The value is meaningless. The decoding result is always output in text format.

## Parallel Decoding

The following configuration options are set only for streaming decoding.

- **decode-style:**

Specifies the decoding format.

Valid value: **'j'**, **'t'**, or **'b'** of the char type, indicating the JSON, TEXT, or binary format, respectively. The default value is **'b'**, indicating binary decoding.

For the JSON and TEXT formats, in the decoding result sent in batches, the uint32 consisting of the first four bytes of each decoding statement indicates the total number of bytes of the statement (the four bytes occupied by the uint32 are excluded, and **0** indicates that the decoding of this batch ends). The 8-byte uint64 indicates the corresponding LSN (**begin** corresponds to **first\_lsn**, **commit** corresponds to **end\_lsn**, and other values correspond to the LSN of the statement).

 NOTE

The binary encoding rules are as follows:

1. The first four bytes represent the total number of bytes of the decoding result of statements following the statement-level delimiter letter P (excluded) or the batch end character F (excluded). If the value is 0, the decoding of this batch ends.
2. The next eight bytes (uint64) indicate the corresponding LSN (**begin** corresponds to **first\_lsn**, **commit** corresponds to **end\_lsn**, and other values correspond to the LSN of the statement).
3. The next one-byte letter can be **B**, **C**, **I**, **U**, or **D**, representing BEGIN, COMMIT, INSERT, UPDATE, or DELETE.
4. If the letter described in 3 is **B**:
  1. The next eight bytes (uint64) indicate the CSN.
  2. The next eight bytes (uint64) indicate **first\_lsn**.
  3. (Optional) If the next one-byte letter is **T**, the following four bytes (uint32) indicate the timestamp length for committing the transaction. The following characters with the same length are the timestamp character string.
  4. (Optional) If the next one-byte letter is **N**, the following four bytes (uint32) indicate the length of the transaction username. The following characters with the same length are the transaction username.
  5. Because there may still be a decoding statement subsequently, a one-byte letter **P** or **F** is used as a separator between statements. **P** indicates that there are still decoded statements in this batch, and **F** indicates that this batch is completed.
5. If the letter described in 3 is **C**:
  1. (Optional) If the next one-byte letter is **X**, the following eight bytes (uint64) indicate the XID.
  2. (Optional) If the next one-byte letter is **T**, the following four bytes (uint32) indicate the timestamp length. The following characters with the same length are the timestamp character string.
  3. When logs are sent in batches, decoding results of other transactions may still exist after a COMMIT log is decoded. If the next one-byte letter is **P**, the batch still needs to be decoded. If the letter is **F**, the batch decoding ends.
6. If the letter described in 3 is **I**, **U**, or **D**:
  1. The next two bytes (uint16) indicate the length of the schema name.
  2. The schema name is read based on the preceding length.
  3. The next two bytes (uint16) indicate the length of the table name.
  4. The table name is read based on the preceding length.
  5. (Optional) If the next 1-byte letter is **N**, it indicates a new tuple. If the letter is **O**, it indicates an old tuple. In this case, the new tuple is sent first.
    1. The following two bytes (uint16) indicate the number of columns to be decoded for the tuple, which is recorded as **attrnum**.
    2. The following procedure is repeated for *attrnum* times.
      1. The next two bytes (uint16) indicate the length of the column name.
      2. The column name is read based on the preceding length.
      3. The next four bytes (uint32) indicate the OID of the current column type.
      4. The next four bytes (uint32) indicate the length of the value (stored in the character string format) in the current column. If the value is **0xFFFFFFFF**, it indicates null. If the value is 0, it indicates a character string whose length is 0.
      5. The column value is read based on the preceding length.



6. Because there may still be a decoding statement after, if the next one-byte letter is **P**, it indicates that the batch still needs to be decoded, and if the next one-byte letter is **F**, it indicates that decoding of the batch ends.

- **sending-batch:**

Specifies whether to send messages in batches.

Valid value: **0** or **1** of the int type. The default value is **0**.

- **0**: The decoding results are sent one by one.
- **1**: When the accumulated size of decoding results reaches 1 MB, decoding results are sent in batches.

In the scenario where batch sending is enabled, if the decoding format is 'j' or 't', before each original decoding statement, a uint32 number is added indicating the length of the decoding result (excluding the current uint32 number), and a uint64 number is added indicating the LSN corresponding to the current decoding result.

- **parallel-queue-size:**

Specifies the length of the queue for interaction between parallel logical decoding threads.

Value range: an integer ranging from 2 to 1024. The value must be an integer power of 2. The default value is **128**.

The queue length is positively correlated with the memory usage during decoding.

### 14.1.3 Logical Decoding by SQL Functions

In GaussDB, you can call SQL functions to create, delete, and push logical replication slots, as well as obtain decoded transaction logs.

#### Procedure

**Step 1** Log in to the primary node of the GaussDB database as a user who has the REPLICATION permission.

**Step 2** Run the following command to connect to the database.

```
gsql -U user1 -d gaussdb -p 16000 -r
```

In the preceding command, **user1** indicates the username, **gaussdb** indicates the name of the database to be connected, and **16000** indicates the database port number. You can replace them as required.

**Step 3** Create a logical replication slot named **slot1**.

```
openGauss=# SELECT * FROM pg_create_logical_replication_slot('slot1', 'mppdb_decoding');
slotname | xlog_position
-----+-----
slot1 | 0/601C150
(1 row)
```

**Step 4** Create a table **t** in the database and insert data into it.

```
openGauss=# CREATE TABLE t(a int PRIMARY KEY, b int);
openGauss=# INSERT INTO t VALUES(3,3);
```

**Step 5** Read the decoding result of **slot1**. The number of decoded records is 4096.

```
openGauss=# SELECT * FROM pg_logical_slot_peek_changes('slot1', NULL, 4096);
location | xid | data
-----+-----
```

```
+-----+
0/601C188	1010023	BEGIN 1010023
0/601ED60	1010023	COMMIT 1010023 CSN 1010022
0/601ED60	1010024	BEGIN 1010024
0/601ED60	1010024	{"table_name":"public.t","op_type":"INSERT","columns_name":
["a","b"],"columns_type":["integer","integer"],"columns_val":["3","3"],"old_keys_name":[""],"old_keys_type":		
[""],"old_keys_val":[""]}		
0/601EED8	1010024	COMMIT 1010024 CSN 1010023
(5 rows)
```

**Step 6** Delete the logical replication slot **slot1**.

```
openGauss=# SELECT * FROM pg_drop_replication_slot('slot1');
pg_drop_replication_slot

(1 row)
```

----End

## 14.1.4 Replicating Data Using the Logical Replication Tool

Currently, SDR and DRS support GaussDB logical replication. The replication tool extracts logical logs from GaussDB and replays them in the peer database. For details about the code of the replication tool that uses JDBC to connect to the database, see [Logical Replication](#).

# 15 Materialized View

A materialized view is a special physical table, which is relative to an ordinary view. An ordinary view is a virtual table and has many application limitations. Any query on a view is actually converted into a query on an SQL statement, and performance is not actually improved. The materialized view actually stores the results of the statements executed by the SQL statement, and is used to cache the results.

Currently, Ustore does not support the creation and use of materialized views.

## 15.1 Complete-refresh Materialized View

### 15.1.1 Overview

Complete-refresh materialized views can be completely refreshed only. The syntax for creating a complete-refresh materialized view is similar to the CREATE TABLE AS syntax.

### 15.1.2 Usage

#### Syntax

- Create a complete-refresh materialized view.  
`CREATE MATERIALIZED VIEW [ view_name ] AS { query_block };`
- Completely refresh a materialized view.  
`REFRESH MATERIALIZED VIEW [ view_name ];`
- Delete a materialized view.  
`DROP MATERIALIZED VIEW [ view_name ];`
- Query a materialized view.  
`SELECT * FROM [ view_name ];`

#### Examples

```
-- Prepare data.
openGauss=# CREATE TABLE t1(c1 int, c2 int);
openGauss=# INSERT INTO t1 VALUES(1, 1);
openGauss=# INSERT INTO t1 VALUES(2, 2);

-- Create a complete-refresh materialized view.
```

```
openGauss=# CREATE MATERIALIZED VIEW mv AS select count(*) from t1;
CREATE MATERIALIZED VIEW

-- Query the materialized view result.
openGauss=# SELECT * FROM mv;
count

 2
(1 row)

-- Insert data into the base table in the materialized view.
openGauss=# INSERT INTO t1 VALUES(3, 3);
INSERT 0 1

-- Completely refresh the complete-refresh materialized view.
openGauss=# REFRESH MATERIALIZED VIEW mv;
REFRESH MATERIALIZED VIEW

-- Query the materialized view result.
openGauss=# SELECT * FROM mv;
count

 3
(1 row)

-- Delete the materialized view.
openGauss=# DROP MATERIALIZED VIEW mv;
DROP MATERIALIZED VIEW
```

## 15.1.3 Support and Constraints

### Supported Scenarios

- Supports the same query scope as the CREATE TABLE AS statement does.
- Supports index creation in complete-refresh materialized views.
- Supports ANALYZE and EXPLAIN.

### Unsupported Scenarios

Materialized views cannot be added, deleted, or modified. They support only query statements.

### Constraints

When a complete-refresh materialized view is refreshed or deleted, a high-level lock is added to the base table. If the definition of a materialized view involves multiple tables, pay attention to the service logic to avoid deadlock.

## 15.2 Fast-refresh Materialized View

### 15.2.1 Overview

Fast-refresh materialized views can be fastly refreshed. You need to manually execute statements to fastly refresh materialized views in a period of time. The difference between the fast-refresh and the complete-refresh materialized views is that the fast-refresh materialized views support only a small number of scenarios. Currently, only base table scanning statements or UNION ALL can be used to create materialized views.

## 15.2.2 Usage

### Syntax

- Create a fast-refresh materialized view.  
`CREATE INCREMENTAL MATERIALIZED VIEW [ view_name ] AS { query_block };`
- Completely refresh a materialized view.  
`REFRESH MATERIALIZED VIEW [ view_name ];`
- Fastly refresh a materialized view.  
`REFRESH INCREMENTAL MATERIALIZED VIEW [ view_name ];`
- Delete a materialized view.  
`DROP MATERIALIZED VIEW [ view_name ];`
- Query a materialized view.  
`SELECT * FROM [ view_name ];`

### Examples

```
-- Prepare data.
openGauss=# CREATE TABLE t1(c1 int, c2 int);
openGauss=# INSERT INTO t1 VALUES(1, 1);
openGauss=# INSERT INTO t1 VALUES(2, 2);

-- Create a fast-refresh materialized view.
openGauss=# CREATE INCREMENTAL MATERIALIZED VIEW mv AS SELECT * FROM t1;
CREATE MATERIALIZED VIEW

-- Insert data.
openGauss=# INSERT INTO t1 VALUES(3, 3);
INSERT 0 1

-- Fastly refresh the materialized view.
openGauss=# REFRESH INCREMENTAL MATERIALIZED VIEW mv;
REFRESH MATERIALIZED VIEW

-- Query the materialized view result.
openGauss=# SELECT * FROM mv;
c1 | c2
----+----
 1 | 1
 2 | 2
 3 | 3
(3 rows)

-- Insert data.
openGauss=# INSERT INTO t1 VALUES(4, 4);
INSERT 0 1

-- Completely refresh the materialized view.
openGauss=# REFRESH MATERIALIZED VIEW mv;
REFRESH MATERIALIZED VIEW

-- Query the materialized view result.
openGauss=# select * from mv;
c1 | c2
----+----
 1 | 1
 2 | 2
 3 | 3
 4 | 4
(4 rows)

-- Delete the materialized view.
openGauss=# DROP MATERIALIZED VIEW mv;
DROP MATERIALIZED VIEW
```

## 15.2.3 Support and Constraints

### Supported Scenarios

- Supports statements for querying a single table.
- Supports UNION ALL for querying multiple single tables.
- Supports index creation in materialized views.
- Supports the Analyze operation in materialized views.

### Unsupported Scenarios

- Multi-table join plans and subquery plans are not supported in materialized views.
- Except for a few ALTER operations, most DDL operations cannot be performed on base tables in materialized views.
- Materialized views cannot be added, deleted, or modified. They support only query statements.
- Temporary, hash bucket, unlogged, or partitioned tables cannot be used to create materialized views.
- Materialized views cannot be created in nested mode (that is, a materialized view cannot be created in another materialized view).
- Only row-store tables are supported.
- Materialized views of the UNLOGGED type are not supported, and the WITH syntax is not supported.

### Constraints

- If the materialized view definition is UNION ALL, each subquery needs to use a different base table.
- When a fast-refresh materialized view is created, completely refreshed, or deleted, a high-level lock is added to the base table. If the materialized view is defined as UNION ALL, pay attention to the service logic to avoid deadlock.

# 16 Error Log Reference

---

## 16.1 Kernel Error Information

ERRMSG: "unsupported syntax: ENCRYPTED WITH in this operation"

SQLSTATE: 42601

CAUSE: "client encryption feature is not supported this operation."

ACTION: "Check client encryption feature whether supported this operation."

ERRMSG: "invalid grant operation"

SQLSTATE: 0LP01

CAUSE: "Grant options cannot be granted to public."

ACTION: "Grant grant options to roles."

ERRMSG: "unrecognized object kind: %d"

SQLSTATE: XX004

CAUSE: "The object type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "unrecognized GrantStmt.targtype: %d"

SQLSTATE: XX004

CAUSE: "The target type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported target types."

ERRMSG: "invalid grant operation"

SQLSTATE: 0LP01

CAUSE: "Grant to public operation is forbidden in security mode."

ACTION: "Don't grant to public in security mode."

ERRMSG: "unrecognized object type"

SQLSTATE: XX004

CAUSE: "The object type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "invalid grant/revoke operation"

SQLSTATE: 0LP01

CAUSE: "Column privileges are only valid for relations in GRANT/REVOKE."

ACTION: "Use the column privileges only for relations."

ERRMSG: "invalid AccessPriv node"

SQLSTATE: 0LP01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "unrecognized GrantStmt.objtype: %d"

SQLSTATE: XX004

CAUSE: "The object type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "undefined client master key"

SQLSTATE: 42705

CAUSE: "The client master key does not exist."

ACTION: "Check whether the client master key exists."

ERRMSG: "undefined column encryption key"

SQLSTATE: 42705

CAUSE: "The column encryption key does not exist."

ACTION: "Check whether the column encryption key exists."

ERRMSG: "large object %u does not exist"

SQLSTATE: 42704



CAUSE: "The large object does not exist."

ACTION: "Check whether the large object exists."

ERRMSG: "redundant options"

SQLSTATE: 42601

CAUSE: "The syntax 'schemas' is redundant in ALTER DEFAULT PRIVILEGES statement."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax."

ERRMSG: "redundant options"

SQLSTATE: 42601

CAUSE: "The syntax 'roles' is redundant in ALTER DEFAULT PRIVILEGES statement."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax."

ERRMSG: "option '%s' not recognized"

SQLSTATE: 42601

CAUSE: "The option in ALTER DEFAULT PRIVILEGES statement is not supported."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax."

ERRMSG: "unrecognized GrantStmt.objtype: %d"

SQLSTATE: XX004

CAUSE: "The object type is not supported for ALTER DEFAULT PRIVILEGES."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax to obtain the supported object types."

ERRMSG: "invalid alter default privileges operation"

SQLSTATE: 0LP01

CAUSE: "Default privileges cannot be set for columns."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax."

ERRMSG: "unrecognized objtype: %d"

SQLSTATE: XX004

CAUSE: "The object type is not supported for default privileges."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax to obtain the supported object types."

ERRMSG: "could not find tuple for default ACL %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "unexpected default ACL type: %d"

SQLSTATE: 0LP01

CAUSE: "The object type is not supported for default privilege."

ACTION: "Check ALTER DEFAULT PRIVILEGES syntax to obtain the supported object types."

ERRMSG: "invalid object id"

SQLSTATE: 0LP01

CAUSE: "The object type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "undefined column"

SQLSTATE: 42703

CAUSE: "The column of the relation does not exist."

ACTION: "Check whether the column exists."

ERRMSG: "column number out of range"

SQLSTATE: 0LP01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for attribute %d of relation %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for relation %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "unsupported object type"

SQLSTATE: 42809

CAUSE: "Index type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "unsupported object type"

SQLSTATE: 42809

CAUSE: "Composite type is not supported for GRANT/REVOKE."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "wrong object type"

SQLSTATE: 42809

CAUSE: "GRANT/REVOKE SEQUENCE only support sequence objects."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "invalid privilege type USAGE for table"

SQLSTATE: 0LP01

CAUSE: "GRANT/REVOKE TABLE do not support USAGE privilege."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported privilege types for tables."

ERRMSG: "invalid privilege type %s for column"

SQLSTATE: 0LP01

CAUSE: "The privilege type is not supported for column object."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported privilege types for column object."

ERRMSG: "cache lookup failed for database %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for foreign-data wrapper %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for foreign server %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for function %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for language %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "Grant/revoke on untrusted languages if forbidden."

SQLSTATE: 0LP01

CAUSE: "Grant/revoke on untrusted languages if forbidden."

ACTION: "Support grant/revoke on trusted C languages"

ERRMSG: "cache lookup failed for large object %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for namespace %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for tablespace %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for type %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cannot set privileges of array types"

SQLSTATE: 0LP01

CAUSE: "Cannot set privileges of array types."

ACTION: "Set the privileges of the element type instead."

ERRMSG: "wrong object type"

SQLSTATE: 42809

CAUSE: "GRANT/REVOKE DOMAIN only support domain objects."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "cache lookup failed for client master key %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for column encryption key %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "cache lookup failed for directory %u"

SQLSTATE: 29P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "unrecognized privilege type '%s'"

SQLSTATE: 42601

CAUSE: "The privilege type is not supported."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported privilege types."

ERRMSG: "unrecognized privilege: %d"

SQLSTATE: XX004

CAUSE: "The privilege type is not supported."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported privilege types."

ERRMSG: "unrecognized AclResult"

SQLSTATE: XX004

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "permission denied for column '%s' of relation '%s'"

SQLSTATE: 42501

CAUSE: "Insufficient privileges for the column."

ACTION: "Select the system tables to get the acl of the column."

ERRMSG: "role with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "unrecognized objkind: %d"

SQLSTATE: XX004

CAUSE: "The object type is not supported for privilege check."

ACTION: "Check GRANT/REVOKE syntax to obtain the supported object types."

ERRMSG: "attribute %d of relation with OID %u does not exist"

SQLSTATE: 42703

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "the column has been dropped"

SQLSTATE: 42703

CAUSE: "The column does not exist."

ACTION: "Check whether the column exists."

ERRMSG: "relation with OID %u does not exist"

SQLSTATE: 42P01

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "invalid group"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "database with OID %u does not exist"

SQLSTATE: 3D000

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "directory with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "function with OID %u does not exist"

SQLSTATE: 42883

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "client master key with OID %u does not exist"

SQLSTATE: 42705

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "language with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "large object %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "schema with OID %u does not exist"

SQLSTATE: 3F001

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "tablespace with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "foreign-data wrapper with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "foreign server with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "type with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "operator with OID %u does not exist"



SQLSTATE: 42883

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "column encryption key with OID %u does not exist"

SQLSTATE: 42705

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "operator class with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "operator family with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "text search dictionary with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "text search configuration with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "collation with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "conversion with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "synonym with OID %u does not exist"

SQLSTATE: 42704

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "package can not create the same name with schema."

SQLSTATE: 22023

CAUSE: "Package name conflict"

ACTION: "Please rename package name"

ERRMSG: "type is not exists %s."

SQLSTATE: 22023

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "This input type is not supported for tdigest\_in()"

SQLSTATE: 0A000

CAUSE: "input type is not supported"

ACTION: "Check tdigest\_in syntax to obtain the supported privilege types"

ERRMSG: "Failed to apply for memory"

SQLSTATE: 53200

CAUSE: "palloc failed"

ACTION: "Check memory"

ERRMSG: "Failed to get tde info from relation '%s'."

SQLSTATE: XX005

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "SPI\_connect failed: %s"

SQLSTATE: SP001

CAUSE: "System error."

ACTION: "Analyze the error message before the error"

ERRMSG: "permission denied for terminate snapshot thread"

SQLSTATE: 42501

CAUSE: "The user does not have system admin privilege"

ACTION: "Grant system admin to user"

ERRMSG: "terminate snapshot thread failed"

SQLSTATE: OP001

CAUSE: "Execution failed due to: %s"

ACTION: "check if snapshot thread exists"

ERRMSG: "terminate snapshot thread failed"

SQLSTATE: OP001

CAUSE: "restart wdr snapshot thread timeoutor The thread did not respond to the kill signal"

ACTION: "Check the wdr snapshot thread is restarted"

ERRMSG: "set lockwait\_timeout failed"

SQLSTATE: XX000

CAUSE: "System error."

ACTION: "Contact engineer to support."

ERRMSG: "permission denied for create WDR Snapshot"

SQLSTATE: 42501

CAUSE: "The user does not have system admin privilege"

ACTION: "Grant system admin to user"

ERRMSG: "WDR snapshot request can not be accepted, please retry later"

SQLSTATE: OP001

CAUSE: "wdr snapshot thread does not exist"

ACTION: "Check if wdr snapshot thread exists"

ERRMSG: "Cannot respond to WDR snapshot request"

SQLSTATE: OP001

CAUSE: "Execution failed due to: %s"

ACTION: "Check if wdr snapshot thread exists"

ERRMSG: "query(%s) can not get datum values"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "create sequence failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check if sequence can be created"

ERRMSG: "update snapshot end time stamp filled"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "query can not get datum values"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "SPI\_connect failed: %s"

SQLSTATE: XX000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "query(%s) execute failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "clean table of snap\_%s is failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "analyze table failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "insert into tables\_snap\_timestamp start time stamp is failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "insert data failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful and check whether the query can be executed"

ERRMSG: "update tables\_snap\_timestamp end time stamp is failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "clean snapshot id %lu is failed in snapshot table"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful and check whether the query can be executed"

ERRMSG: "clean snapshot failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "can not create snapshot stat table"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "create WDR snapshot data table failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "insert into tables\_snap\_timestamp start time stamp failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "insert into snap\_%s is failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "update tables\_snap\_timestamp end time stamp failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "create index failed"

SQLSTATE: 22000

CAUSE: "System error."

ACTION: "Check whether the query can be executed"

ERRMSG: "analyze table, connection failed: %s"

SQLSTATE: XX000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "snapshot thread SPI\_connect failed: %s"

SQLSTATE: XX000

CAUSE: "System error."

ACTION: "Check whether the snapshot retry is successful"

ERRMSG: "Distributed key column can't be transformed"

SQLSTATE: 42P10

CAUSE: "There is a risk of violating uniqueness when transforming distribution columns."

ACTION: "Change transform column."

ERRMSG: "cannot convert %s to %s"

SQLSTATE: 42804

CAUSE: "There is no conversion path in pg\_cast."

ACTION: "Rewrite or cast the expression."

ERRMSG: "create matview on TDE table failed"

SQLSTATE: 0A000

CAUSE: "create materialized views is not supported on TDE table"

ACTION: "check CREATE syntax about create the materialized views"

ERRMSG: "schema name can not same as package"

SQLSTATE: 22023

CAUSE: "schema name conflict"

ACTION: "rename schema name"

ERRMSG: "Unrecognized commandType when checking read-only attribute."

SQLSTATE: XX004

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Fail to generate subquery plan."

SQLSTATE: XX005

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Unrecognized node type when processing qual condition."

SQLSTATE: XX004

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Unrecognized node type when processing const parameters."

SQLSTATE: XX004

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "SELECT FOR UPDATE/SHARE is not allowed with UNION/INTERSECT/  
EXCEPT"

SQLSTATE: 0A000

CAUSE: "SQL uses unsupported feature."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "GROUP BY cannot be implemented."

SQLSTATE: 0A000

CAUSE: "GROUP BY uses unsupported datatypes."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "TSDB functions cannot be used if enable\_tsdb is off."

SQLSTATE: D0011

CAUSE: "Functions are not loaded."

ACTION: "Turn on enable\_tsdb according to manual."

ERRMSG: "Unrecognized node type when extracting index."

SQLSTATE: XX004

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Ordering operator cannot be identified."



SQLSTATE: 42883

CAUSE: "Grouping set columns must be able to sort their inputs."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "DISTINCT cannot be implemented."

SQLSTATE: 0A000

CAUSE: "DISTINCT uses unsupported datatypes."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "Failed to locate grouping columns."

SQLSTATE: 55000

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Resjunk output columns are not implemented."

SQLSTATE: 20000

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "PARTITION BY cannot be implemented."

SQLSTATE: 0A000

CAUSE: "PARTITION BY uses unsupported datatypes."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "ORDER BY cannot be implemented."

SQLSTATE: 0A000

CAUSE: "ORDER BY uses unsupported datatypes."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "Failed to deconstruct sort operators into partitioning/ordering operators."

SQLSTATE: D0011

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Pool size should not be zero"

SQLSTATE: 22012

CAUSE: "Compute pool configuration file contains error."

ACTION: "Please check the value of 'pl' in cp\_client.conf."

ERRMSG: "Failed to get the runtime info from the compute pool."

SQLSTATE: 22004

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "Version is not compatible between local cluster and the compute pool."

SQLSTATE: XX008

CAUSE: "Compute pool is not installed appropriately."

ACTION: "Configure compute pool according to manual."

ERRMSG: "No optional index path is found."

SQLSTATE: 01000

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "MERGE INTO on replicated table does not yet support using distributed tables."

SQLSTATE: 0A000

CAUSE: "SQL uses unsupported feature."

ACTION: "Modify SQL statement according to the manual."

ERRMSG: "Fail to find ForeignScan node!"

SQLSTATE: P0002

CAUSE: "System error."

ACTION: "Contact Huawei Engineer."

ERRMSG: "sql advisor don't support none table, temp table, system table."

SQLSTATE: 42601

CAUSE: "sql advisor don't support none table, temp table, system table."

ACTION: "check query component"

ERRMSG: "Invalid autonomous transaction return datatypes"

SQLSTATE: P0000

CAUSE: "PL/SQL uses unsupported feature."

ACTION: "Contact Huawei Engineer."

ERRMSG: "new row for relation '%s' violates check constraint '%s'"

SQLSTATE: 23514

CAUSE: "some rows copy failed"

ACTION: "check table defination"

ERRMSG: "new row for relation '%s' violates check constraint '%s'"

SQLSTATE: 23514

CAUSE: "some rows copy failed"

ACTION: "set client\_min\_messages = info for more details"

ERRMSG: "get gauss home path is NULL"

SQLSTATE: XX005

CAUSE: "gauss home path not set"

ACTION: "check if \$GAUSSHOME is exist"

ERRMSG: "unable to open kms\_iam\_info.json file"

SQLSTATE: 58P03

CAUSE: "file not exist or broken"

ACTION: "check the kms\_iam\_info.json file"

ERRMSG: "can not get password plaintext"

SQLSTATE: XX005

CAUSE: "file not exist or broken"

ACTION: "check the password cipher rand file"

ERRMSG: "IAM info json key is NULL"

SQLSTATE: XX005

CAUSE: "IAM info value error"

ACTION: "check tde\_config kms\_iam\_info.json file"

ERRMSG: "get internal password is NULL"

SQLSTATE: XX005

CAUSE: "cipher rand file missing"

ACTION: "check password cipher rand file"

ERRMSG: "KMS info json key is NULL"

SQLSTATE: XX005

CAUSE: "KMS info value error"

ACTION: "check tde\_config kms\_iam\_info.json file"

ERRMSG: "unable to get json file"

SQLSTATE: 58P03

CAUSE: "parse json file failed"

ACTION: "check the kms\_iam\_info.json file format"

ERRMSG: "get JSON tree is NULL"

SQLSTATE: XX005

CAUSE: "get KMS JSON tree failed"

ACTION: "check input parameter or config.ini file"

ERRMSG: "failed to get json tree"

SQLSTATE: XX005

CAUSE: "config.ini json tree error"

ACTION: "check input parameter or config.ini file"

ERRMSG: "failed to set the value of json tree"

SQLSTATE: XX005

CAUSE: "config.ini json tree error"

ACTION: "check input parameter or config.ini file"

ERRMSG: "http request failed"

SQLSTATE: XX005

CAUSE: "http request error"

ACTION: "check KMS or IAM connect or config parameter"

ERRMSG: "get iam token or iam agency token is NULL"

SQLSTATE: XX005

CAUSE: "connect IAM failed"

ACTION: "check if your env can connect with IAM server"

ERRMSG: "KMS dek json key is NULL"

SQLSTATE: XX005

CAUSE: "KMS return value error"

ACTION: "check KMS config parameter"

ERRMSG: "get kms dek is NULL"

SQLSTATE: XX005

CAUSE: "connect KMS failed"

ACTION: "check if your env can connect with KMS server"

ERRMSG: "get http header is NULL"

SQLSTATE: XX005

CAUSE: "http request failed"

ACTION: "check IAM config parameter"

ERRMSG: "create KMS dek failed"

SQLSTATE: XX005

CAUSE: "KMS error"

ACTION: "check KMS connect or config parameter"

ERRMSG: "get KMS dek failed"

SQLSTATE: XX005

CAUSE: "KMS error"

ACTION: "check KMS connect or config parameter"

ERRMSG: "get KMS DEK is NULL"

SQLSTATE: XX005

CAUSE: "get KMS dek\_plaintext failed"

ACTION: "check KMS network or cipher is right"

ERRMSG: "create matview with TDE failed"

SQLSTATE: 0A000

CAUSE: "TDE feature is not supported for Create materialized views"

ACTION: "check CREATE syntax about create the materialized views"

ERRMSG: "failed to add item to the index page"

SQLSTATE: XX002

CAUSE: "System error."

ACTION: "Check WARNINGS for the details."

ERRMSG: "index row size %lu exceeds maximum %lu for index '%s'"

SQLSTATE: 54000

CAUSE: "Values larger than 1/3 of a buffer page cannot be indexed."

ACTION: "Consider a function index of an MD5 hash of the value, or use full text indexing."

## 16.2 CM Error Information

ERRMSG: "Fail to access the cluster static config file."

SQLSTATE: c3000

CAUSE: "The cluster static config file is not generated or is manually deleted."

ACTION: "Please check the cluster static config file."

ERRMSG: "Fail to open the cluster static file."

SQLSTATE: c3000

CAUSE: "The cluster static config file is not generated or is manually deleted."

ACTION: "Please check the cluster static config file."

ERRMSG: "Fail to read the cluster static file."

SQLSTATE: c3001

CAUSE: "The cluster static file permission is insufficient."

ACTION: "Please check the cluster static config file."

ERRMSG: "Failed to read the static config file."

SQLSTATE: c1000

CAUSE: "out of memeory."

ACTION: "Please check the system memory and try again."

ERRMSG: "Could not find the current node in the cluster by the node id %u."

SQLSTATE: c3002

CAUSE: "The static config file probably contained content error."

ACTION: "Please check static config file."

ERRMSG: "Failed to open the logic config file."

SQLSTATE: c3000

CAUSE: "The logic config file is not generated or is manually deleted."

ACTION: "Please check the cluster static config file."

ERRMSG: "Fail to read the logic static config file."

SQLSTATE: c3001

CAUSE: "The logic static config file permission is insufficient."

ACTION: "Please check the logic static config file."

ERRMSG: "Failed to open or read the static config file."

SQLSTATE: c1000

CAUSE: "out of memeory."

ACTION: "Please check the system memory and try again."

ERRMSG: "Failed to open the log file '%s'."

SQLSTATE: c3000

CAUSE: "Log file not found."

ACTION: "Please check the log file."

ERRMSG: "Failed to open the log file '%s'."

SQLSTATE: c3000

CAUSE: "The log file permission is insufficient."

ACTION: "please check the log file."

ERRMSG: "Failed to open the dynamic config file '%s'."

SQLSTATE: c3000

CAUSE: "The dynamic config file permission is insufficient."

ACTION: "Please check the dynamic config file."

ERRMSG: "Failed to malloc memory, size = %lu."

SQLSTATE: c1000

CAUSE: "out of memeory."

ACTION: "Please check the system memory and try again."

ERRMSG: "unrecognized AZ name '%s'."

SQLSTATE: c3000

CAUSE: "The parameter(%s) entered by the user is incorrect."

ACTION: "Please check the parameter entered by the user and try again."

ERRMSG: "unrecognized minorityAz name '%s'."

SQLSTATE: c3000

CAUSE: "The parameter(%s) entered by the user is incorrect."

ACTION: "Please check the parameter entered by the user and try again."

ERRMSG: "Get GAUSSHOM failed."

SQLSTATE: c3000

CAUSE: "The environment variable('GAUSSHOM') is incorrectly configured."

ACTION: "Please check the environment variable('GAUSSHOM')."

ERRMSG: "Get current user name failed."

SQLSTATE: c3000

CAUSE: "N/A"

ACTION: "Please check the environment."

ERRMSG: "-B option must be specified."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-T option must be specified.\n"



SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "can't stop one node or instance with -m normal."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "can't stop one node or instance with -m resume."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "can't stop one availability zone with -m resume."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "log level or cm server arbitration mode must be specified."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "log level or cm server arbitration mode need not be specified."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-R is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-D is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n and -R are needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n and -D are needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "no operation specified."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "no cm directory specified."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "Please check the usage of switchover."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n and -z cannot be specified at the same time."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-m cannot be specified at the same time with -n or -z."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n node(%d) is invalid."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n node is needed."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "%s: -C is needed."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-z value must be 'ALL' when query mppdb cluster."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-v is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-C is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-Cv is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-L value must be 'ALL' when query logic cluster."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "unrecognized LC name '%s'."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n is needed."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "There is no '%s' information in cluster."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-D path is too long.\n"

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-D path is invalid."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-n node(%s) is invalid."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-R only support when the cluster is single-inst."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-t time is invalid."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "-votenum is invalid."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "unrecognized build mode."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "unrecognized build mode '%s'."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "too many command-line arguments (first is '%s')."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "unrecognized operation mode '%s'."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "no cm directory specified."

SQLSTATE: c3000

CAUSE: "%s: The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "Failed to malloc memory."

SQLSTATE: c1000

CAUSE: "out of memeory."

ACTION: "Please check the system memory and try again."

ERRMSG: "Failed to open etcd: %s."

SQLSTATE: c4000

CAUSE: "Etcd is abnormal."

ACTION: "Please check the Cluster Status and try again."

ERRMSG: "[PATCH-ERROR] hotpatch command or path set error."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "no standby datanode in single node cluster."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "restart logic cluster failed."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "restart logic cluster failed"

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

ERRMSG: "The option parameter is not specified."

SQLSTATE: c3000

CAUSE: "The cmdline entered by the user is incorrect."

ACTION: "Please check the cmdline entered by the user(%s)."

# 17

## Configuring GUC Parameters

### 17.1 Viewing Parameter Values

GaussDB uses a set of default running parameters after it is installed. You can modify the parameters to better fit the current service scenarios and data volume.

#### Procedure

**Step 1** Connect to the database.

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database to be connected, and **8000** is the port number of the database primary node.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr
6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
openGauss=#
```

**Step 2** View the parameter values in the database.

- Method 1: Run the **SHOW** command.
  - Run the following command to view the value of a certain parameter:  
openGauss=# **SHOW** *server\_version*;  
*server\_version* indicates the database version.
  - Run the following command to view values of all parameters:  
openGauss=# **SHOW ALL**;
- Method 2: Query the **pg\_settings** view.
  - Run the following command to view the value of a certain parameter:  
openGauss=# **SELECT \* FROM pg\_settings WHERE NAME='server\_version'**;
  - Run the following command to view values of all parameters:  
openGauss=# **SELECT \* FROM pg\_settings**;

----End



## Example

Check the character encoding type of the client.

```
openGauss=# SHOW client_encoding;
client_encoding

UTF8
(1 row)
```

## 17.2 Resetting Parameters

You are advised to modify some parameters on the GaussDB console. If the parameters cannot be modified on the console, evaluate the risks and contact customer service.

### Background

GaussDB provides multiple methods to set the GUC parameters of databases, users, or sessions.

- Parameter names are case-insensitive.
- A parameter value can be an integer, floating point value, string, Boolean value, or enumerated value.
  - The Boolean values can be **on/off**, **true/false**, **yes/no**, or **1/0**, and are case-insensitive.
  - The enumerated value range is specified in the **enumvals** column of the **pg\_settings** system catalog.
- For parameters using units, specify their units during the setting. Otherwise, default units are used.
  - The default units are specified in the **unit** column of **pg\_settings**.
  - The unit of memory can be KB, MB, or GB.
  - The unit of time can be ms, s, min, h, or d.

For details about parameters in the hosts configuration template, see [GUC Parameters](#).

### Setting GUC Parameters

GaussDB provides six types of GUC parameters. For details about parameter types and their setting methods, see [Table 17-1](#).

**Table 17-1** GUC parameters

| Category | Description                                                                                                                                                                         | How to Set |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| INTERNAL | Fixed parameter. It is set during database creation and cannot be modified. Users can only view the parameter by running the <b>SHOW</b> command or in the <b>pg_settings</b> view. | None       |

| Category   | Description                                                                                                                                                                                                                                                        | How to Set                                                                                                                           |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| POSTMASTER | Database server parameter. It can be set when the database is started or in the configuration file.                                                                                                                                                                | Method 1 in <a href="#">Table 17-2</a> .                                                                                             |
| SIGHUP     | Global database parameter. It can be set when the database is started or be modified later.                                                                                                                                                                        | Method 1 or 2 in <a href="#">Table 17-2</a> .                                                                                        |
| BACKEND    | Session connection parameter. It is specified during session connection creation and cannot be modified after that. The parameter setting becomes invalid when the session is disconnected. This is an internal parameter and not recommended for users to set it. | Method 1 or 2 in <a href="#">Table 17-2</a> .<br><b>NOTE</b><br>The parameter setting takes effect when the next session is created. |
| SUSET      | Database administrator parameter. It can be set by common users when or after the database is started. It can also be set by database administrators using SQL statements.                                                                                         | Method 1 or 2 by a common user, or method 3 by a database administrator in <a href="#">Table 17-2</a> .                              |
| USERSET    | Common user parameter. It can be set by any user at any time.                                                                                                                                                                                                      | Method 1, 2, or 3 in <a href="#">Table 17-2</a> .                                                                                    |

You can set GUC parameters in GaussDB using the three methods listed in [Table 17-2](#).

**Table 17-2** Methods for setting GUC parameters

| No.      | How to Set                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Method 1 | <ol style="list-style-type: none"><li>1. Log in to the console.</li><li>2. On the <b>Instances</b> page, click the instance name to go to the <b>Basic Information</b> page.</li><li>3. In the navigation pane on the left, click <b>Parameters</b>. On the displayed page, modify parameters.<br/>If the parameters cannot be modified on the console, evaluate the risks in advance and contact customer service for modification.</li><li>4. Restart the database to make the setting take effect.</li></ol> <b>NOTE</b><br>Rebooting instances will interrupt user operations. Plan a proper execution window before the reboot. |

| No.       | How to Set                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Method 2  | <ol style="list-style-type: none"><li>1. Log in to the console.</li><li>2. On the <b>Instances</b> page, click the instance name to go to the <b>Basic Information</b> page.</li><li>3. In the navigation pane on the left, click <b>Parameters</b>. On the displayed page, modify parameters.<br/>If the parameters cannot be modified on the console, evaluate the risks in advance and contact customer service for modification.</li></ol> |
| Method 3: | <ol style="list-style-type: none"><li>1. Log in to the console.</li><li>2. On the <b>Instances</b> page, click the instance name to go to the <b>Basic Information</b> page.</li><li>3. In the navigation pane on the left, click <b>Parameters</b>. On the displayed page, modify parameters.<br/>If the parameters cannot be modified on the console, evaluate the risks in advance and contact customer service for modification.</li></ol> |

---

**CAUTION**

- If you use method 1 or 2 to set a parameter that does not belong to the current environment, the database displays a message indicating that the parameter is not supported.
  - When you use method 3 to set a parameter, if the parameter value is an integer, leading zeros will be filtered out. For example, **SET *paraname* TO 008192** and **SET *paraname* TO 8192** have the same effect.
- 

## Procedure

The following example shows how to set **hot\_standby** on the primary node of the database using method 1.

**Step 1** Log in as the OS user **omm** to the primary node of the database.

**Step 2** View the value of **hot\_standby**.

```
cat /gaussdb/data/dbnode/postgresql.conf | grep "hot_standby"
hot_standby = on
```

**on** indicates logs are archived.

**Step 3** Set **hot\_standby** to **off** to disable log archiving.

```
gs_guc set -Z datanode -D /gaussdb/data/dbnode -c "hot_standby=off"
```

**NOTE**

You can set **hot\_standby** to **off** for the database nodes.

```
gs_guc set -Z datanode -N all -I all -c "hot_standby=off"
```

**Step 4** Restart the database to make the setting take effect.

```
gs_om -t stop && gs_om -t start
```

**Step 5** Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr
6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.

openGauss=#
```

**Step 6** Check whether the parameter is correctly set.

```
openGauss=# SHOW hot_standby;
hot_standby

off
(1 row)
```

----End

The following example shows how to set **authentication\_timeout** on the primary node of the database using method 2.

**Step 1** Log in as the OS user **omm** to the primary node of the database.

**Step 2** View the value of **authentication\_timeout**.

```
cat /gaussdb/data/dbnode/postgresql.conf | grep authentication_timeout
authentication_timeout = 1min
```

**Step 3** Set **authentication\_timeout** to **59s**.

```
gs_guc reload -Z datanode -N all -I all -c "authentication_timeout = 59s"

Total instances: 2. Failed instances: 0.
Success to perform gs_guc!
```

#### NOTE

You can run the following command to set **authentication\_timeout** to **59s** for the database nodes:

```
gs_guc reload -Z datanode -N all -I all -c "authentication_timeout = 59s"
```

**Step 4** Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr
6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.

openGauss=#
```

**Step 5** Check whether the parameter is correctly set.

```
openGauss=# SHOW authentication_timeout;
authentication_timeout
```

```

59s
(1 row)
```

----End

The following example shows how to set **explain\_perf\_mode** using method 3.

**Step 1** Log in as the OS user **omm** to the primary node of the database.

**Step 2** Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr
6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
openGauss=#
```

**Step 3** View the value of **explain\_perf\_mode**.

```
openGauss=# SHOW explain_perf_mode;
explain_perf_mode

normal
(1 row)
```

**Step 4** Set **explain\_perf\_mode**.

Perform one of the following operations:

- Set a database-level parameter.  
openGauss=# **ALTER DATABASE postgres SET explain\_perf\_mode TO pretty;**

If the following information is displayed, the setting succeeds:

```
ALTER DATABASE
```

The setting takes effect in the next session.

- Set a user-level parameter.  
openGauss=# **ALTER USER omm SET explain\_perf\_mode TO pretty;**

If the following information is displayed, the setting succeeds:

```
ALTER ROLE
```

The setting takes effect in the next session.

- Set a session-level parameter.  
openGauss=# **SET explain\_perf\_mode TO pretty;**

If the following information is displayed, the setting succeeds:

```
SET
```

**Step 5** Check whether the parameter is correctly set.

```
openGauss=# SHOW explain_perf_mode;
explain_perf_mode

pretty
(1 row)
```

----End

## Examples

- Example 1: modifying the maximum connections for the primary database node in GaussDB using method 1

a. Log in as the OS user **omm** to the primary node of the database.

b. Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

c. View the allowed maximum connections.

```
openGauss=# SHOW max_connections;
max_connections
```

```

200
(1 row)
```

d. Run the following command to exit the database:

```
openGauss=# \q
```

e. Change the maximum connections for the primary database node in GaussDB.

```
gs_guc set -Z datanode -N all -I all -c "max_connections = 800"
```

f. Restart the database.

```
gs_om -t stop && gs_om -t start
```

g. Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

h. View the allowed maximum connections.

```
openGauss=# SHOW max_connections;
max_connections
```

```

800
(1 row)
```

- Example 2: setting **authentication\_timeout** (the allowed longest client authentication duration) for the primary database node using method 2

a. Log in as the OS user **omm** to the primary node of the database.

b. Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit
2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

- c. View the value of **authentication\_timeout**.

```
openGauss=# SHOW authentication_timeout;
authentication_timeout

1min
(1 row)
```

- d. Run the following command to exit the database:

```
openGauss=# \q
```

- e. Change the allowed longest client authentication duration of the primary database node.

```
gs_guc reload -Z datanode -N all -I all -c "authentication_timeout = 59s"
```

- f. Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit
2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

- g. View the value of **authentication\_timeout**.

```
openGauss=# SHOW authentication_timeout;
authentication_timeout

59s
(1 row)
```

- Example 3: Change the maximum number of connections between GaussDB database nodes.

- a. Log in as the OS user **omm** to the primary node of the database.

- b. Run the following command to connect to the database:

```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit
2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

- c. View the allowed maximum connections.

```
openGauss=# SHOW max_connections;
max_connections

```

- ```
200
(1 row)
```
- d. Run the following command to exit the database:
- ```
openGauss=# \q
```
- e. Change the maximum number of connections between GaussDB database nodes.
- ```
gs_guc set -Z datanode -N all -I all -c "max_connections = 500"
```
- f. Restart the database.
- ```
gs_om -t stop
gs_om -t start
```
- g. Run the following command to connect to the database:
- ```
gsql -d postgres -p 8000
```
- postgres** is the name of the database, and 8000 is the port number of the primary node of the database.
- If information similar to the following is displayed, the connection succeeds:
- ```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit
2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.

openGauss=#
```
- h. View the allowed maximum connections.
- ```
openGauss=# SHOW max_connections;
max_connections
-----
500
(1 row)
```
- Example 4: setting authentication_timeout (the allowed longest client authentication duration) for database nodes
- a. Log in as the OS user **omm** to the primary node of the database.
- b. Run the following command to connect to the database:
- ```
gsql -d postgres -p 8000
```
- postgres** is the name of the database, and 8000 is the port number of the primary node of the database.
- If information similar to the following is displayed, the connection succeeds:
- ```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit
2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.

openGauss=#
```
- c. View the value of **authentication_timeout**.
- ```
openGauss=# SHOW authentication_timeout;
authentication_timeout

1min
(1 row)
```
- d. Run the following command to exit the database:
- ```
openGauss=# \q
```
- e. Change the allowed longest client authentication duration of GaussDB database nodes.
- ```
gs_guc reload -Z datanode -N all -I all -c "authentication_timeout = 30s"
```
- f. Run the following command to connect to the database:



```
gsql -d postgres -p 8000
```

**postgres** is the name of the database, and 8000 is the port number of the primary node of the database.

If information similar to the following is displayed, the connection succeeds:

```
gsql((GaussDB Kernel VxxxRxxxCxx build f521c606) compiled at 2021-09-16 14:55:22 commit 2935 last mr 6385 release)
Non-SSL connection (SSL connection is recommended when requiring high-security)
Type "help" for help.
```

```
openGauss=#
```

g. View the value of **authentication\_timeout**.

```
openGauss=# SHOW authentication_timeout;
authentication_timeout

30s
(1 row)
```

## 17.3 GUC Parameters

### 17.3.1 GUC Parameter Usage

A database provides many operation parameters. Configurations of these parameters affect the behavior of the database system. Before modifying these parameters, learn the impact of these parameters on the database. Otherwise, unexpected results may occur.

#### Precautions

- If the value range of a parameter is a string, the string should comply with the naming conventions of the path and file name in the OS running the target database.
- If the maximum value of a parameter is *INT\_MAX*, the maximum parameter value varies by OS.
- If the maximum value of a parameter is *DBL\_MAX*, the maximum parameter value varies by OS.

### 17.3.2 File Location

After a database has been installed, three configuration files (**postgresql.conf**, **pg\_hba.conf**, and **pg\_ident.conf**) are automatically generated and saved in the data directory. You can use the methods described in this section to change the names and save paths of these configuration files.

When changing the storage directory of a configuration file, set **data\_directory** in **postgresql.conf** to the actual data directory.

---

#### NOTICE

If a configuration file is incorrectly modified, the database will be seriously affected. Do not modify the configuration files mentioned in this section after installation.

---

## data\_directory

**Parameter description:** Specifies the GaussDB **data** directory. Only users with the **sysadmin** permission can access this parameter. You can set this parameter using one of the following methods:

- Set it when you install the GaussDB.
- This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

**Default value:** Specify this parameter during installation. If this parameter is not specified during installation, the database is not initialized by default.

## config\_file

**Parameter description:** Specifies the configuration file (**postgresql.conf**) of the primary server.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

**Default value:** **postgresql.conf** (The absolute directory of this file may be displayed in the actual situation.)

## hba\_file

**Parameter description:** Specifies the configuration file (**pg\_hba.conf**) for host-based authentication (HBA). This parameter can be specified only in the **postgresql.conf** file and can be accessed only by users with the **sysadmin** permission.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** **pg\_hba.conf** (The absolute directory of this file may be displayed in the actual situation.)

## ident\_file

**Parameter description:** Specifies the name of the configuration file (**pg\_ident.conf**) for client authentication. Only users with the **sysadmin** permission can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** **pg\_ident.conf** (The absolute directory of this file may be displayed in the actual situation.)

## external\_pid\_file

**Parameter description:** Specifies the extra PID file that can be used by the server management program. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

This parameter takes effect only after the database restarts.

---

**Value range:** a string

**Default value:** empty

## 17.3.3 Connection and Authentication

### 17.3.3.1 Connection Settings

This section describes parameters related to client-server connection modes.

#### **listen\_addresses**

**Parameter description:** Specifies the TCP/IP address of the client for a server to listen on.

This parameter specifies the IP address used by the GaussDB server for listening, for example, IPv4. Multiple NICs may exist on the host and each NIC can be bound to multiple IP addresses. This parameter specifies the IP addresses to which GaussDB is bound. The client can use the IP address specified by this parameter to connect to or send requests to GaussDB.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:**

- Host name or IP address. Multiple values are separated with commas (,).
- Asterisk (\*) or **0.0.0.0**, indicating that all IP addresses will be listened to, which is not recommended due to potential security risks. This parameter must be used together with valid addresses (for example, the local IP address). Otherwise, the build may fail. In primary/standby mode, if the value is set to \\* or **0.0.0.0**, the value of **localport** in the **postgresql.conf** file of the database on the primary node cannot be the value of **dataPortBase + 1**. Otherwise, the database cannot be started.
- If the parameter is not specified, the server does not listen on any IP address. In this case, only Unix domain sockets can be used for database connections.

**Default value:** After the database instance is installed, the default value is configured according to the IP address of different instances in the **public\_cloud.conf** configuration file. The default value for the DN instance is the IP address of the **data.net** NIC.

 **NOTE**

The **public\_cloud.conf** file contains the following NIC information: **mgr.net** (management NIC), **data.net** (data NIC), and **virtual.net** (virtual NIC).

## local\_bind\_address

**Parameter description:** Specifies the host IP address bound to the current node for connecting to other nodes in the database.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Default value:** After the database instance is installed, the default value is configured according to the IP address of different instances in the **public\_cloud.conf** configuration file. The default value for the DN instance is the IP address of the **data.net** NIC.

### NOTE

The **public\_cloud.conf** file contains the following NIC information: **mgr.net** (management NIC), **data.net** (data NIC), and **virtual.net** (virtual NIC).

## port

**Parameter description:** Specifies the TCP port listened on by the GaussDB.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

### NOTE

This parameter is specified in the configuration file during installation. Do not modify this parameter unless absolutely necessary. Otherwise, database communication will be affected.

**Value range:** an integer ranging from 1 to 65535

### NOTE

- When setting the port number, ensure that the port number is not in use. When setting the port numbers of multiple instances, ensure that the port numbers do not conflict.
- Ports 1 to 1023 are reserved for the operating system. Do not use them.
- When the database instance is installed using the configuration file, pay attention to the ports reserved in the communication matrix in the configuration file. For example, *dataPortBase* + 1 needs to be reserved as the port used by internal tools, and *dataPortBase* + 6 needs to be reserved as the communication port of the streaming engine message queue. (Due to specification changes, the current version no longer supports this feature. Do not use it.) Therefore, during database instance installation, the maximum port number is **65529** for DNs. Ensure that the port number does not conflict with each other.

**Default value:** **5432** (The actual value is specified in the configuration file during installation.)

## max\_connections

**Parameter description:** Specifies the maximum number of concurrent connections to the database. This parameter influences the concurrent processing capability of the database.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer The minimum value is **10** (greater than **max\_wal\_senders**). The theoretical maximum value is **262143**. The actual

maximum value is a dynamic value, which is calculated using the formula:  $262143 - \text{job\_queue\_processes} - \text{autovacuum\_max\_workers} - \text{AUXILIARY\_BACKENDS} - \text{AV\_LAUNCHER\_PROCS} - \text{max\_inner\_tool\_connections} - \min(\max(\text{newValue}/4, 64), 1024)$ . The values of [job\\_queue\\_processes](#), [autovacuum\\_max\\_workers](#), and [max\\_inner\\_tool\\_connections](#) depend on the settings of the corresponding GUC parameters. **AUXILIARY\_BACKENDS** indicates the number of reserved auxiliary threads, which is fixed to **20**. **AV\_LAUNCHER\_PROCS** indicates the number of reserved autovacuum launcher threads, which is fixed to **2**. In  $\min(\max(\text{newValue}/4, 64), 1024)$ , **newValue** indicates the new value.

For different memory specifications, the value range of this parameter in different instances is as follows:

**Table 17-3** Value ranges of memory specifications for different instances

| Memory Specifications | Value Range in DN |
|-----------------------|-------------------|
| < 32GB                | [10, 100]         |
| [32GB, 64GB)          | [10, 200]         |
| [64GB, 128GB)         | [10, 2048]        |
| [128GB, 256GB)        | [10, 5000]        |
| [256GB, 480GB)        | [10, 11000]       |
| [480GB, 512GB)        | [10, 24000]       |
| [512GB, 640GB)        | [10, 25000]       |
| [640GB, 768GB)        | [10, 34000]       |
| [768GB, 1024GB)       | [10, 40000]       |
| [1024GB, 1536GB)      | [10, 55000]       |
| [1536GB, 2048GB)      | [10, 85000]       |
| >= 2048GB             | [10, 110000]      |

**Default value:**

**55000** (128-core CPU/1024 GB memory, 104-core CPU/1024 GB memory, 96-core CPU/1024 GB memory); **40000** (96-core CPU/768 GB memory); **25000** (64-core CPU/512 GB memory); **24000** (60-core CPU/480 GB memory); **11000** (32-core CPU/256 GB memory); **5000** (16-core CPU/128 GB memory); **2048** (8-core CPU/64 GB memory);

**Setting suggestions:**

Retain the default value of this parameter on the primary node of the databases.

**Impact of incorrect configuration:**

- If the value of **max\_connections** exceeds the maximum dynamic value, the node fails to be started and the following error message is displayed: "invalid

value for parameter "max\_connections". Alternatively, the memory fails to be allocated during the node startup and the following error message is displayed: "Cannot allocate memory".

- If only the value of *max\_connections* is increased while the memory parameter is not adjusted in proportion according to the external egress specifications, when the service load is heavy, the memory may be insufficient, and the error message "memory is temporarily unavailable" is displayed.

#### NOTE

- If the number of connections of the administrator exceeds the value of *max\_connections*, the administrator can still connect to the database after the connections are used up by common users. If the number of connections exceeds the value of *sysadmin\_reserved\_connections*, an error is reported. That is, the maximum number of connections of the administrator is equal to the value of *max\_connections* + *sysadmin\_reserved\_connections*.
- For common users, internal jobs use some connections. Therefore, the value of this parameter is slightly less than that of *max\_connections*. The value depends on the number of internal connections.

## max\_inner\_tool\_connections

**Parameter description:** Specifies the maximum number of concurrent connections of a tool which is allowed to connect to the database. This parameter influences the concurrent connection capability of the GaussDB tool.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to *MIN* (which takes the smaller value between **262143** and *max\_connections*). For details about how to calculate the value of *max\_connections*, see the preceding description.

**Default value:** 50 for each database node. If the default value is greater than the maximum value supported by the kernel (determined when the **gs\_initdb** command is executed), an error message is displayed.

#### Setting suggestions:

Retain the default value of this parameter on the primary node of the databases.

If this parameter is set to a large value, GaussDB requires more System V shared memory or semaphores, which may exceed the default maximum configuration of the OS. In this case, modify the value as needed.

## sysadmin\_reserved\_connections

**Parameter description:** Specifies the minimum number of connections reserved for administrators. You are advised not to set this parameter to a large value. This parameter is used together with the *max\_connections* parameter. The maximum number of connections of the administrator is equal to the value of *max\_connections* + *sysadmin\_reserved\_connections*.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *MIN* (which takes the smaller value between **262143** and *max\_connections*). For details about how to calculate the value of *max\_connections*, see the preceding description.

**Default value:** 3

**Note:** When the thread pool function is enabled, if the thread pool is fully occupied, a processing bottleneck occurs. As a result, connections reserved by the administrator cannot be established. In this case, you can use `gsql` to establish connections through the primary port number + 1 to clear useless sessions.

## unix\_socket\_directory

**Parameter description:** Specifies the Unix domain socket directory for the GaussDB server to listen to connections from the client.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

The parameter length limit varies by OS. If the length is exceeded, the error "Unix-domain socket path xxx is too long" will be reported.

**Value range:** a string

**Default value:** empty. The actual value is specified by the configuration file during installation.

## unix\_socket\_group

**Parameter description:** Specifies the group of the Unix domain socket (the user of a socket is the user that starts the server). This parameter can work with [unix\\_socket\\_permissions](#) to control socket access.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. If this parameter is set to an empty string, the default group of the current user is used.

**Default value:** empty

## unix\_socket\_permissions

**Parameter description:** Specifies access permissions for the Unix domain socket.

The Unix domain socket uses the common permission set of the Unix file system. The value of this parameter should be a number (acceptable for the **chmod** and **umask** commands). If a user-defined octal format is used, the number must start with 0.

You are advised to set it to **0770** (only allowing access from users connecting to the database and users in the same group as them) or **0700** (only allowing access from users connecting to the database).

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0000 to 0777

**Default value:** 0700

#### NOTE

In the Linux OS, a document has one document attribute and nine permission attributes, which consists of the read (r), write (w), and execute (x) permissions of the Owner, Group, and Others groups.

The r, w, and x permissions are represented by the following numbers:

r: 4

w: 2

x: 1

-: 0

The three attributes in a group are accumulative.

For example, **-rwxrwx---** indicates the following permissions:

owner = rwx = 4+2+1 = 7

group = rwx = 4+2+1 = 7

others = --- = 0+0+0 = 0

The permission of the file is 0770.

## application\_name

**Parameter description:** Specifies the client name used in the current connection request.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

When a standby node requests to replicate logs on the primary node, if this parameter is not an empty string, it is used as the name of the streaming replication slot of the standby node on the primary node. In this case, if the length of this parameter exceeds 61 bytes, only the first 61 bytes are used as the streaming replication slot name.

**Value range:** a string

**Default value:** empty (The actual value is the name of the application connected to the backend.)

## connection\_info

**Parameter description:** Specifies the database connection information, including the driver type, driver version, driver deployment path, and process owner.

This is a USERSET parameter used for O&M. You are advised not to change the parameter value.

**Value range:** a string

**Default value:** empty



 NOTE

- An empty string indicates that the driver connected to the database does not support automatic setting of the **connection\_info** parameter or the parameter is not set by users in applications.
- The following is an example of the concatenated value of **connection\_info**:  

```
{"driver_name":"ODBC","driver_version": "(GaussDB Kernel VxxxRxxxCxx build 290d125f) compiled at 2020-05-08 02:59:43 commit 2143 last mr 131 release","driver_path":"/usr/local/lib/psqlodbcw.so","os_user":"omm"}
```

**driver\_name** and **driver\_version** are displayed by default. Whether **driver\_path** and **os\_user** are displayed is determined by users.

## check\_disconnect\_query

**Parameter description:** Specifies whether to terminate the execution of statements on the GaussDB server after the client is disconnected abnormally (for example, socketTimeout is triggered by JDBC, rwttimeout is triggered by libpq and the connection is closed, or the client process is terminated during service running).

**Parameter type:** Boolean

**Unit:** none

**Value range:**

- **on:** indicates that the GaussDB server stops running the corresponding statements after the client is disconnected unexpectedly.
- **off:** indicates that the GaussDB server does not stop running the corresponding statements after the client is disconnected unexpectedly.

**Default value:** on

**Setting method:** This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Setting suggestion:** Retain the default value.

---

 CAUTION

If this parameter is set and an upgrade is required, check whether the target version supports this parameter before the upgrade. If this parameter is not supported, delete it from the configuration file first before proceeding with the upgrade.

---

### 17.3.3.2 Security and Authentication (postgresql.conf)

This section describes parameters about client-to-server authentication.

## authentication\_timeout

**Parameter description:** Specifies the longest duration to wait before the client authentication times out. If a client is not authenticated by the server within the period, the server automatically disconnects from the client so that the client does not occupy connection resources.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 600. The smallest unit is s.

**Default value:** 1min

## auth\_iteration\_count

**Parameter description:** Specifies the number of iterations during the generation of encryption information for authentication.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 2048 to 134217728

**Default value:** 10000

---

### NOTICE

If the number of iterations is too small, the password storage security is reduced. If the number of iterations is too large, the performance deteriorates in scenarios involving password encryption, such as authentication and user creation. Set the number of iterations based on actual hardware conditions. You are advised to retain the default value.

---

## session\_authorization

**Parameter description:** Specifies the user ID of the current session.

This is a USERSET parameter and can be set only using the SET SESSION AUTHORIZATION syntax.

**Value range:** a string

**Default value:** NULL

## session\_timeout

**Parameter description:** Specifies the longest duration allowed when no operations are performed on a client after it is connected to the server.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 86400. The smallest unit is s. **0** indicates that the timeout is disabled.

**Default value:** 1800s

---

### NOTICE

The gsql client of GaussDB has an automatic reconnection mechanism. For local connection of initialized users, the client reconnects to the server if the connection breaks after the timeout.

---

## ssl

**Parameter description:** Specifies whether SSL connections are enabled.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that SSL connections are enabled.
- **off** indicates that SSL connections are not enabled.

---

### NOTICE

To enable SSL connections, you also need to configure parameters such as **ssl\_cert\_file**, **ssl\_key\_file**, and **ssl\_ca\_file** and the corresponding files. Incorrect configurations may cause startup failure of the database.

---

**Default value:** on

## require\_ssl

**Parameter description:** Specifies whether the server requires the SSL connection. This parameter is valid only when **ssl** is set to **on**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the server requires SSL connections.
- **off** indicates that the server does not require SSL connections.

---

### NOTICE

GaussDB supports SSL when a client connects to a the primary node of the database. It is recommended that the SSL connection be enabled only on the primary node of the databases.

---

**Default value:** off

## ssl\_ciphers

**Parameter description:** Specifies the list of encryption algorithms supported by SSL. Only the sysadmin user can access the list.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. Separate multiple encryption algorithms by semicolons (;).

**Default value:** ALL

---

### NOTICE

If **ssl\_ciphers** is set incorrectly, the database cannot be started properly.

---

## ssl\_renegotiation\_limit

**Parameter description:** Specifies the allowed traffic volume over an SSL-encrypted channel before the session key is renegotiated. The renegotiation mechanism reduces the probability that attackers use the password analysis method to crack the key based on a huge amount of data but causes big performance losses. The traffic indicates the sum of transmitted and received traffic. The SSL renegotiation mechanism has been disabled because of potential risks. This parameter is reserved for version compatibility and does not take effect.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647. The unit is KB. **0** indicates that the renegotiation mechanism is disabled.

**Default value:** 0

## ssl\_cert\_file

**Parameter description:** Specifies the name of the file that contains the SSL server certificate. The relative path is relative to the data directory.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** server.crt

## ssl\_key\_file

**Parameter description:** Specifies the name of the file that contains the SSL private key. The relative path is relative to the data directory.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** server.key

## ssl\_ca\_file

**Parameter description:** Specifies the name of a file that contains CA information. The relative path is relative to the data directory.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. If it is an empty string, no CA file is loaded and client certificate verification is not performed.

**Default value:** cacert.pem

## ssl\_crl\_file

**Parameter description:** Specifies the certificate revocation list (CRL). If the certificate of a client is in the list, the certificate is invalid. A relative path must be used. The path depends on the data directory.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that there is no CRL.

**Default value:** empty

## krb\_server\_keyfile

**Parameter description:** Specifies the location of the main configuration file of the Kerberos service. Only the sysadmin user can access the file.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** empty

## krb\_srvname

**Parameter description:** Specifies the Kerberos service name.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** postgres

## krb\_caseins\_users

**Parameter description:** Specifies whether the Kerberos username is case-sensitive.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the Kerberos username is case-insensitive.
- **off** indicates that the Kerberos username is case-sensitive.

**Default value:** off

## modify\_initial\_password

**Parameter description:** After GaussDB is installed, there is only one initial user account (whose UID is 10) in the database. When a user logs in to the database using this initial account for the first time, this parameter determines whether the password of the initial account needs to be modified.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

If the initial user password is not specified during the installation, the initial user password is empty by default after the installation. Before performing other operations, you need to set the initial user password using the gsql client. This parameter no longer takes effect and is reserved only for compatibility with upgrade scenarios.

---

**Value range:** Boolean

- **on** indicates that the password of the initial user needs to be modified upon the first login after the database is successfully installed.
- **off** indicates that the password of the initial user does not need to be modified upon the first login after the database is successfully installed.

**Default value:** off

## password\_policy

**Parameter description:** Specifies whether to check the password complexity when you run the **CREATE ROLE/USER** or **ALTER ROLE/USER** command to create or modify a GaussDB account.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

For security purposes, do not disable the password complexity policy.

---

**Value range:** 0 and 1

- **0** indicates that no password complexity policy is enabled.
- **1** indicates that the default password complexity policy is enabled.

**Default value:** 1

## password\_reuse\_time

**Parameter description:** Specifies whether to check the reuse interval of the new password when you run the **ALTER USER** or **ALTER ROLE** command to change a user password.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

When you change the password, the system checks the values of [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#).

- If the values of [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#) are both positive numbers, an old password can be reused when it meets either of the reuse restrictions.
  - If the value of [password\\_reuse\\_time](#) is **0**, password reuse is restricted based on the number of reuse times, and not on the reuse interval.
  - If the value of [password\\_reuse\\_max](#) is **0**, password reuse is restricted based on the reuse interval, and not on the number of reuse times.
  - If the values of both [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#) are **0**, password reuse is not restricted.
- 

**Value range:** a floating-point number ranging from 0 to 3650. The unit is day.

- **0** indicates that the password reuse interval is not checked.
- A positive number indicates that a new password cannot be chosen from passwords in history that are newer than the specified number of days.

**Default value:** 0

## password\_reuse\_max

**Parameter description:** Specifies whether to check the reuse times of the new password when you run the **ALTER USER** or **ALTER ROLE** command to change a user password. Only the **sysadmin** user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

When you change the password, the system checks the values of [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#).

- If the values of [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#) are both positive numbers, an old password can be reused when it meets either of the reuse restrictions.
- If the value of [password\\_reuse\\_time](#) is **0**, password reuse is restricted based on the number of reuse times, and not on the reuse interval.
- If the value of [password\\_reuse\\_max](#) is **0**, password reuse is restricted based on the reuse interval, and not on the number of reuse times.
- If the values of both [password\\_reuse\\_time](#) and [password\\_reuse\\_max](#) are **0**, password reuse is not restricted.

---

**Value range:** an integer ranging from 0 to 1000

- **0** indicates that the password reuse times are not checked.
- A positive number indicates that the new password cannot be the one whose reuse times exceed the specified number.

**Default value:** 0

## password\_lock\_time

**Parameter description:** Specifies the duration before a locked account is automatically unlocked.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

The locking and unlocking functions take effect only when the values of [password\\_lock\\_time](#) and [failed\\_login\\_attempts](#) are positive numbers.

---

**Value range:** a floating-point number ranging from 0 to 365. The unit is day.

- **0** indicates that an account is not automatically locked if the password verification fails.
- A positive number indicates the duration after which a locked account is automatically unlocked.

**Default value:** 1

## failed\_login\_attempts

**Parameter description:** Specifies the maximum number of incorrect password attempts before an account is locked. The account will be automatically unlocked after the time specified by **password\_lock\_time**. Only the **sysadmin** user can access the account. The automatic account locking policy applies in scenarios such as login and password modification using the **ALTER USER** command.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

The locking and unlocking functions take effect only when the values of **failed\_login\_attempts** and **password\_lock\_time** are positive numbers.

---

**Value range:** an integer ranging from 0 to 1000

- **0** indicates that the automatic locking function does not take effect.
- A positive number indicates that an account is locked when the number of incorrect password attempts reaches the specified number.

**Default value:** 10

## password\_encryption\_type

**Parameter description:** Specifies the encryption type of a user password. Changing the value of this parameter does not change the password encryption type of existing users. The new encryption type is applied to passwords of new users or passwords modified after the parameter value is changed.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0, 1, 2, or 3

- **0** indicates that passwords are encrypted with MD5.
- **1** indicates that passwords are encrypted with SHA-256 and MD5.
- **2** indicates that passwords are encrypted with SHA-256.
- **3** indicates that the passwords are encrypted in sm3 mode.

---

### NOTICE

The MD5 encryption algorithm is not recommended because it has lower security and poses security risks.

---

**Default value:** 2



## password\_min\_length

**Parameter description:** Specifies the minimum length of an account password. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. A password can contain 6 to 999 characters.

**Default value:** 8

## password\_max\_length

**Parameter description:** Specifies the maximum length of an account password. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. A password can contain 6 to 999 characters.

**Default value:** 32

## password\_min\_uppercase

**Parameter description:** Specifies the minimum number of uppercase letters that an account password must contain. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 999

- 0 means no limit.
- An integer ranging from 1 to 999 indicates the minimum number of uppercase letters required in a password when you create an account.

**Default value:** 0

## password\_min\_lowercase

**Parameter description:** Specifies the minimum number of lowercase letters that an account password must contain. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 999

- 0 means no limit.
- An integer ranging from 1 to 999 indicates the minimum number of lowercase letters required in a password when you create an account.

**Default value:** 0

## password\_min\_digital

**Parameter description:** Specifies the minimum number of digits that an account password must contain. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 999

- 0 means no limit.
- An integer ranging from 1 to 999 indicates the minimum number of digits required in a password when you create an account.

**Default value:** 0

## password\_min\_special

**Parameter description:** Specifies the minimum number of special characters that an account password must contain. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 999

- 0 means no limit.
- An integer ranging from 1 to 999 indicates the minimum number of special characters required in a password when you create an account.

**Default value:** 0

## password\_effect\_time

**Parameter description:** Specifies the validity period of an account password.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating-point number ranging from 0 to 999. The unit is day.

- 0 indicates that the validity period restriction is disabled.
- A floating-point number from 1 to 999 indicates the number of days for which an account password is valid. When the password is about to expire or has expired, the system prompts the user to change the password.

**Default value:** 0

## password\_notify\_time

**Parameter description:** Specifies how many days in advance a user is notified before a password expires.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 999. The unit is day.

- 0 indicates that the reminder is disabled.
- An integer ranging from 1 to 999 indicates the number of days prior to password expiration that a user will receive a reminder.

**Default value:** 7

## ssl\_cert\_notify\_time

**Parameter description:** Specifies the number of days prior to SSL server certificate expiration that a user will receive a reminder. When the SSL certificate is initialized during connection establishment, if the duration from the current time to the certificate expiration time is shorter than the specified value, an expiration notification is recorded in the log.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 7 to 180. The unit is day.

**Default value:** 90

### 17.3.3.3 Communication Library Parameters

This section describes parameter settings and value ranges for communication libraries.

## tcp\_keepalives\_idle

**Parameter description:** Specifies the interval for transmitting keepalive signals on an OS that supports the TCP\_KEEPIDLE socket option. If no keepalive signal is transmitted, the connection is in idle mode.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

- If the OS does not support TCP\_KEEPIDLE, set this parameter to 0.
  - The parameter is ignored on an OS where connections are established using the Unix domain socket.
  - If this parameter is set to 0, the system value is used.
  - This parameter is not shared among different sessions. That is, different session connections may have different values.
  - The parameter value in the current session connection, not the value of the GUC copy, is displayed.
- 

**Value range:** 0 to 3600. The unit is s.

**Default value:** 60

## tcp\_keepalives\_interval

**Parameter description:** Specifies the response time before retransmission on an OS that supports the TCP\_KEEPINTVL socket option.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 to 180. The unit is s.

**Default value:**

---

**NOTICE**

- If the OS does not support TCP\_KEEPINTVL, set this parameter to **0**.
  - The parameter is ignored on an OS where connections are established using the Unix domain socket.
  - If this parameter is set to **0**, the system value is used.
  - This parameter is not shared among different sessions. That is, different session connections may have different values.
  - The parameter value in the current session connection, not the value of the GUC copy, is displayed.
- 

## tcp\_keepalives\_count

**Parameter description:** Specifies the number of keepalive signals that can be waited before the GaussDB server is disconnected from the client on an OS that supports the TCP\_KEEPCNT socket option.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- If the OS does not support TCP\_KEEPCNT, set this parameter to **0**.
  - The parameter is ignored on an OS where connections are established using the Unix domain socket.
  - If this parameter is set to **0**, the system value is used.
  - This parameter is not shared among different sessions. That is, different session connections may have different values.
  - The parameter value in the current session connection, not the value of the GUC copy, is displayed.
- 

**Value range:** 0 to 100. **0** indicates that the connection is immediately broken if GaussDB does not receive a keepalive signal from the client.

**Default value:** 20

## tcp\_user\_timeout

**Parameter description:** Specifies the maximum duration for which the transmitted data can remain in the unacknowledged state before the TCP connection is forcibly closed when the GaussDB sends data on the OS that supports the TCP\_USER\_TIMEOUT socket option.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**NOTICE**

- If the OS does not support the TCP\_USER\_TIMEOUT option, the value of this parameter does not take effect. The default value is 0.
- The parameter is ignored on an OS where connections are established using the Unix domain socket.

**Value Range:** 0 to 3600000. The unit is ms. The value 0 indicates that the value is set based on the OS.

**Default value:** 0

Note that the effective result of this parameter varies according to the OS kernel.

- For AArch64 EulerOS (Linux kernel version: 4.19), the timeout interval is the value of this parameter.
- For x86 EulerOS 2.5 (Linux kernel version: 3.10), the timeout interval is not the set value of this parameter but the maximum value in different ranges. That is, the maximum value of the total Linux TCP retransmission duration range to which the set value of **tcp\_user\_timeout** belongs. For example, if **tcp\_user\_timeout** is set to 40000, the actual total retransmission duration is 51 seconds.

**Table 17-4** Value of tcp\_user\_timeout for x86 EulerOS 2.5 (Linux kernel version: 3.10)

| Number of Linux TCP Retransmission Times | Total Linux TCP Retransmission Duration Range (s) | Example of tcp_user_timeout (ms) | Actual Linux TCP Retransmission Duration (s) |
|------------------------------------------|---------------------------------------------------|----------------------------------|----------------------------------------------|
| 1                                        | (0.2,0.6]                                         | 400                              | 0.6                                          |
| 2                                        | (0.6,1.4]                                         | 1000                             | 1.4                                          |
| 3                                        | (1.4,3]                                           | 2000                             | 3                                            |
| 4                                        | (3,6.2]                                           | 4000                             | 6.2                                          |
| 5                                        | (6.2,12.6]                                        | 10000                            | 12.6                                         |
| 6                                        | (12.6,25.4]                                       | 20000                            | 25.4                                         |
| 7                                        | (25.4,51]                                         | 40000                            | 51                                           |
| 8                                        | (51,102.2]                                        | 80000                            | 102.2                                        |
| 9                                        | (102.2,204.6]                                     | 150000                           | 204.6                                        |
| 10                                       | (204.6,324.6]                                     | 260000                           | 324.6                                        |
| 11                                       | (324.6,444.6]                                     | 400000                           | 444.6                                        |

Note: The duration of each TCP retransmission increases exponentially with the number of retransmission times until it reaches 120s. When the duration of a TCP retransmission reaches 120s, the duration of each subsequent retransmission remains 120s.

## comm\_proxy\_attr

**Parameter description:** Specifies the parameters related to the communication proxy library.

### NOTE

- This parameter applies only to the centralized Arm-based standalone system running EulerOS 2.9.
- This function takes effect when the thread pool is enabled, that is, **enable\_thread\_pool** is set to **on**.
- When setting this parameter, you need to set the GUC parameter **local\_bind\_address** to the IP address of the NIC of the **libos\_kni**.
- **Parameter template:** `comm_proxy_attr = '{enable_libnet:true, enable_dfx:false, numa_num:4, numa_bind:[[30,31],[62,63],[94,95],[126,127]]}'`
- Parameters that need to be configured include:
  - **enable\_libnet:** whether to enable the user-mode protocol. The options are as follows: **true** and **false**.
  - **enable\_dfx:** whether to enable the communication proxy library view. The options are as follows: **true** and **false**.
  - **numa\_num:** number of NUMA nodes in the system. 2P and 4P servers are supported. The value can be: **4** or **8**.
  - **numa\_bind:** core binding parameter of the agent thread. Each numa has two CPUs. There are a total of **numa\_num** groups. The value range is as follows: [0, Number of CPUs – 1].

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

**Default value:** none

## 17.3.4 Resource Consumption

### 17.3.4.1 Memory

This section describes memory parameters.

#### NOTICE

These parameters, except **local\_syscache\_threshold**, take effect only after the database restarts.

## memorypool\_enable

**Parameter description:** Specifies whether to enable a memory pool.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the memory pool is enabled.
- **off** indicates that the memory pool is disabled.

**Default value:** off

## memorypool\_size

**Parameter description:** Specifies the memory pool size.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 128 x 1024 to *INT\_MAX*/2. The unit is KB.

**Default value:** 512 MB

## enable\_memory\_limit

**Parameter description:** Specifies whether to enable the logical memory management module.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the logical memory management module is enabled.
- **off** indicates that the logical memory management module is disabled.

**Default value:** on

---

### CAUTION

- If the value of **max\_process\_memory** minus **shared\_buffer** minus **cstore\_buffers** minus metadata size is less than 2 GB, GaussDB forcibly sets **enable\_memory\_limit** to **off**. Metadata is the memory used in GaussDB and is related to some concurrent parameters, such as **max\_connections**, **thread\_pool\_attr** and **max\_prepared\_transactions**.
  - If this parameter is set to **off**, the memory used by the database is not limited. When a large number of concurrent or complex queries are performed, too much memory is used, which may cause OS OOM problems.
- 

## max\_process\_memory

**Parameter description:** Specifies the maximum physical memory of a database node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 2 x 1024 x 1024 to *INT\_MAX*. The unit is KB.

**Default value:**

**900 GB** (128-core CPU/1024 GB memory, 104-core CPU/1024 GB memory, 96-core CPU/1024 GB memory); **680 GB** (96-core CPU/768 GB memory); **450 GB** (64-core

CPU/512 GB memory); **420 GB** (60-core CPU/480 GB memory); **200 GB** (32-core CPU/256 GB memory); **90 GB** (16-core CPU/128 GB memory); **40 GB** (8-core CPU/64 GB memory);

**Setting suggestions:**

The value on the database node is determined based on the physical memory of the system and the number of master database nodes deployed on a single node. The recommended calculation formula is as follows:  $(\text{Physical memory} - \text{vm.min\_free\_kbytes}) \times 0.7 / (1 + \text{Number of primary nodes})$ . This parameter is used to prevent node OOM caused by memory usage increase, ensuring system reliability. **vm.min\_free\_kbytes** indicates the OS memory reserved for the kernel to receive and send data. Its value is at least 5% of the total memory. That is,  $\text{max\_process\_memory} = \text{Physical memory} \times 0.665 / (1 + \text{Number of primary nodes})$

---

 **CAUTION**

If this parameter is set to a value greater than the physical memory of the server, the OS OOM problem may occur.

---

## enable\_memory\_context\_control

**Parameter description:** Specifies whether to enable the function of checking whether the number of memory contexts exceeds the specified limit. This parameter applies only to the DEBUG version.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the function of checking the number of memory contexts is enabled.
- **off** indicates that the function of checking the number of memory contexts is disabled.

**Default value:** off

## uncontrolled\_memory\_context

**Parameter description:** Specifies which memory context will not be checked when the function of checking whether the number of memory contexts exceeds the specified limit is enabled. This parameter applies only to the DEBUG version.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

During the query, the title meaning string "MemoryContext white list:" is added to the beginning of the parameter value.

**Value range:** a string

**Default value:** empty



## shared\_buffers

**Parameter description:** Specifies the size of shared memory used by GaussDB. Increasing the value of this parameter causes GaussDB to request more System V shared memory than the default configuration allows.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 16 to 1073741823. The unit is 8 KB.

The value of **shared\_buffers** must be an integer multiple of **BLCKSZ**. Currently, **BLCKSZ** is set to **8 KB**. That is, the value of **shared\_buffers** must be an integer multiple of 8 KB. The minimum value changes according to the value of **BLCKSZ**.

**Default value:**

**360 GB** (128-core CPU/1024 GB memory, 104-core CPU/1024 GB memory, 96-core CPU/1024 GB memory); **270 GB** (96-core CPU/768 GB memory); **180 GB** (64-core CPU/512 GB memory); **160 GB** (60-core CPU/480 GB memory); **80 GB** (32-core CPU/256 GB memory); **36 GB** (16-core CPU/128 GB memory); **16 GB** (8-core CPU/64 GB memory);

**Setting suggestions:**

1. Set **shared\_buffers** to a value less than 40% of the memory.
2. If **shared\_buffers** is set to a larger value, increase the value of **checkpoint\_segments** because a longer period of time is required to write a large amount of new or changed data.
3. If the process fails to be restarted after the value of **shared\_buffers** is changed, perform either of the following operations based on the error information:
  - a. Adjust the **kernel.shmall**, **kernel.shmmax**, and **kernel.shmmin** OS parameters. For details, see "Preparing for Installation" > "Modifying OS Configuration" > "Configuring Other OS Parameters" in *Installation Guide*.
  - b. Run the **free -g** command to check whether the available memory and swap space of the OS are sufficient. If the memory is insufficient, manually stop other user programs that occupy much memory.
  - c. Do not set **shared\_buffers** to an excessively large or small value.

## segment\_buffers

**Parameter description:** Specifies the memory size of a GaussDB segment-paged metadata page.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 16 to 1073741823. The unit is 8 KB.

The value of **segment\_buffers** must be an integer multiple of **BLCKSZ**. Currently, **BLCKSZ** is set to **8 KB**. That is, the value of **segment\_buffers** must be an integer multiple of 8 KB. The minimum value changes according to the value of **BLCKSZ**.

**Default value:** 8 MB

**Setting suggestions:**

**segment\_buffers** is used to cache the content of segment-paged headers, which is key metadata information. To improve performance, it is recommended that the segment headers of common tables be cached in the buffer and not be replaced. You are advised to set this parameter based on the following formula: Number of tables (including indexes and toast tables) x Number of partitions x 3 + 128. This is because each table (partition) has some extra metadata segments. Generally, a table has three segments. At last, + 128 is added because segment- and page-based tablespace management requires a certain number of buffers.

If this parameter is set to a small value, it takes a long time to create a segment-paged table for the first time. Therefore, you are advised to set this parameter to the recommended value.

## bulk\_write\_ring\_size

**Parameter description:** Specifies the size of the ring buffer used by the operation when a large amount of data is written (for example, the copy operation).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 16384 to 2147483647. The unit is KB.

**Default value:** 2 GB

**Setting suggestions:** Increase the value of this parameter on database nodes if a huge amount of data will be imported.

## standby\_shared\_buffers\_fraction

**Parameter description:** Specifies the **shared\_buffers** proportion used on the server where a standby instance is deployed.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a double-precision floating-point number ranging from 0.1 to 1.0

**Default value:** 1

## temp\_buffers

**Parameter description:** Specifies the maximum size of local temporary buffers used by a database session.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**temp\_buffers** can be modified only before the first use of temporary tables within each session. Subsequent attempts to change the value of this parameter will not take effect on that session.

A session allocates temporary buffers based on the value of **temp\_buffers**. If a large value is set in a session that does not require many temporary buffers, only the overhead of one buffer descriptor is added. If a buffer is used, additional 8192 bytes will be consumed for it.

**Value range:** an integer ranging from 100 to 1073741823. The unit is 8 KB.

**Default value:** 1 MB

## max\_prepared\_transactions

**Parameter description:** Specifies the maximum number of transactions that can stay in the **prepared** state simultaneously. Increasing the value of this parameter causes GaussDB to request more System V shared memory than the default configuration allows.

When GaussDB is deployed as an HA system, set this parameter on standby servers to a value equal to or greater than that on primary servers. Otherwise, queries will fail on the standby servers.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 262143

**Default value:** 200

### NOTE

Generally, explicit PREPARE operations are not required for transactions. If explicit PREPARE operations are performed for transactions, increase the value of this parameter to be greater than the number of concurrent services that require PREPARE to prevent preparation failures.

## work\_mem

**Parameter description:** Specifies the amount of memory to be used by internal sort operations and hash tables before they write data into temporary disk files. Sort operations are required for **ORDER BY**, **DISTINCT**, and merge joins. Hash tables are used in hash joins, hash-based aggregation, and hash-based processing of **IN** subqueries.

In a complex query, several sort or hash operations may run in parallel; each operation will be allowed to use as much memory as this parameter specifies. If the memory is insufficient, data will be written into temporary files. In addition, several running sessions could be performing such operations concurrently. Therefore, the total memory used may be many times the value of **work\_mem**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 64 to 2147483647. The unit is KB.

**Default value:**

**256 MB** (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory, and 96-core CPU/768-GB memory); **128 MB** (64-core CPU/512-GB memory, 60-core CPU/480-GB memory, 32-core CPU/256-GB memory, and 16-core CPU/128-GB memory); **64 MB** (8-core CPU/64-GB memory);

**NOTICE****Setting suggestions:**

If the physical memory specified by **work\_mem** is insufficient, additional operator calculation data will be written into temporary tables based on query characteristics and the degree of parallelism. This reduces performance by five to ten times, and prolongs the query response time from seconds to minutes.

- For complex serial queries, each query requires five to ten associated operations. Set **work\_mem** using the following formula: **work\_mem** = 50% of the memory/10.
- For simple serial queries, each query requires two to five associated operations. Set **work\_mem** using the following formula: **work\_mem** = 50% of the memory/5.
- For concurrent queries, set **work\_mem** using the following formula: **work\_mem** = **work\_mem** for serial queries/Number of concurrent SQL statements.
- BitmapScan hash tables are also restricted by **work\_mem**, but will not be forcibly flushed to disks. In the case of complete lossify, every 1 MB memory occupied by the hash table corresponds to a 16 GB page of BitmapHeapScan (32 GB for Ustore). After the upper limit of **work\_mem** is reached, the memory increases linearly with the data access traffic based on this ratio.

## query\_mem

**Parameter description:** Specifies the memory used by a query.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or an integer greater than 32 MB. The default unit is KB.

**Default value:** 0

**NOTICE**

- If the value of **query\_mem** is greater than 0, the optimizer adjusts the memory cost estimate to this value when generating an execution plan.
- If the value is set to a negative value or a positive integer less than 32 MB, the default value 0 is used. In this case, the optimizer does not adjust the estimated query memory.

## query\_max\_mem

**Parameter description:** Specifies the maximum memory that can be used by a query.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or an integer greater than 32 MB. The default unit is KB.

**Default value:** 0

---

**NOTICE**

- If the value of **query\_max\_mem** is greater than 0, an error is reported when the query memory usage exceeds the value.
  - If the value is set to a negative value or a positive integer less than 32 MB, the default value **0** is used. In this case, the optimizer does not limit the query memory.
- 

## **maintenance\_work\_mem**

**Parameter description:** Specifies the maximum amount of memory to be used by maintenance operations, such as **VACUUM**, **CREATE INDEX**, and **ALTER TABLE ADD FOREIGN KEY**. This parameter may affect the execution efficiency of **VACUUM**, **VACUUM FULL**, **CLUSTER**, and **CREATE INDEX**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1024 to *INT\_MAX*. The unit is KB.

**Default value:**

**2 GB** (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory, 96-core CPU/768-GB memory, 64-core CPU/512-GB memory, and 60-core CPU/480-GB memory); **1 GB** (32-core CPU/256-GB memory); **512 MB** (16-core CPU/128-GB memory); **256 MB** (8-core CPU/64-GB memory);

---

**NOTICE**

**Setting suggestions:**

- The value of this parameter must be greater than that of **work\_mem** so that database dumps can be cleared or restored more quickly. In a database session, only one maintenance operation can be performed at a time. Maintenance is usually performed when there are not many running sessions.
  - When **Automatic Vacuuming** threads are running, the memory equal to a multiple of the value of **autovacuum\_max\_workers** is allocated. In this case, set **maintenance\_work\_mem** to a value equal to or greater than that of **work\_mem**.
  - If a large amount of data is to be clustered, increase the value of this parameter in the session.
- 

## **max\_stack\_depth**

**Parameter description:** Specifies the maximum safe depth of the GaussDB execution stack. The safety margin is required because the stack depth is not checked in every routine in the server, but only in key potentially-recursive routines, such as expression evaluation.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 100 to *INT\_MAX*. The unit is KB.

**Default value:**

- If the value of **ulimit -s** minus 640 KB is equal to or greater than 2 MB, the default value of this parameter is **2 MB**.
- If the value of **ulimit -s** minus 640 KB is less than 2 MB, the default value of this parameter is the value of **ulimit -s** minus 640 KB.

---

**NOTICE**

When setting this parameter, comply with the following principles:

- The database needs to reserve 640 KB stack depth. Therefore, the ideal value of this parameter is the actual stack size limit enforced by the OS kernel (as set by **ulimit -s**) minus 640 KB.
  - If the value of this parameter is greater than the value of **ulimit -s** minus 640 KB before the database is started, the database fails to be started. During database running, if the value of this parameter is greater than the value of **ulimit -s** minus 640 KB, this parameter does not take effect.
  - If the value of **ulimit -s** minus 640 KB is less than the minimum value of this parameter, the database fails to be started.
  - Setting this parameter to a value greater than the actual kernel limit means that a running recursive function may crash an individual backend process.
  - Since not all OSs provide this function, you are advised to set a specific value for this parameter.
  - The default value is **2 MB**, which is relatively small and does not easily cause system breakdown.
- 

## bulk\_read\_ring\_size

**Parameter description:** Specifies the size of the ring buffer used by the operation when a large amount of data is queried (for example, during large table scanning).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 256 to 2147483647. The unit is KB.

**Default value:** 16 MB

## enable\_early\_free

**Parameter description:** Specifies whether the operator memory can be released in advance.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the operator memory can be released in advance.
- **off** indicates that the operator memory cannot be released in advance.

**Default value:** on

## memory\_trace\_level

**Parameter description:** Specifies the control level for recording memory allocation information after the dynamic memory usage exceeds 90% of the maximum dynamic memory. This parameter takes effect only when **use\_workload\_manager** and **enable\_memory\_limit** are enabled. This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **none:** indicates that memory allocation information is not recorded.
- **level1:** After the dynamic memory usage exceeds 90% of the maximum dynamic memory, the following information is recorded and the recorded memory information is saved in the *\$GAUSSLOG/mem\_log* directory.
  - Global memory overview.
  - Memory usage of the top 20 memory contexts of the instance, session, and thread types.
  - The **totalsize** and **freesize** columns for each memory context.
- **level2:** After the dynamic memory usage exceeds 90% of the maximum dynamic memory, the following information is recorded and the recorded memory information is saved in the *\$GAUSSLOG/mem\_log* directory.
  - Global memory overview.
  - Memory usage of the top 20 memory contexts of the instance, session, and thread types.
  - The **totalsize** and **freesize** columns for each memory context.
  - Detailed information about all memory allocations in each memory context, including the file where the allocated memory is located, line number, and size.

**Default value:** level1

---

### NOTICE

- If this parameter is set to **level2**, the memory allocation details (file, line, and size) of each memory context are recorded, which greatly affects the performance. Therefore, exercise caution when setting this parameter.
  - The recorded memory snapshot information can be queried by using the system function **gs\_get\_history\_memory\_detail(cstring)**.
  - The recorded memory context is obtained after all memory contexts of the same type with the same name are summarized.
- 

## resilience\_memory\_reject\_percent

**Parameter description:** Specifies the dynamic memory usage percentage for escape from memory overload. This parameter takes effect only when the GUC parameters **use\_workload\_manager** and **enable\_memory\_limit** are enabled. This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

This parameter consists of **recover\_memory\_percent** and **overload\_memory\_percent**. The meanings of the two parts are as follows:

- **recover\_memory\_percent**: Percentage of the dynamic memory usage when the memory recovers from overload to the maximum dynamic memory. When the dynamic memory usage is less than the maximum dynamic memory multiplied by the value of this parameter, the overload escape function is disabled and new connections are allowed. The value ranges from 0 to 100. The value indicates a percentage.
- **overload\_memory\_percent**: Percentage of the dynamic memory usage to the maximum dynamic memory when the memory is overloaded. When the dynamic memory usage is greater than the maximum dynamic memory multiplied by the value of this parameter, the current memory is overloaded. In this case, the overload escape function is triggered to kill sessions and new connections are prohibited. The value ranges from 0 to 100. The value indicates a percentage.

**Default value:** '0,0', indicating that the escape from memory overload function is disabled.

**Example:**

```
resilience_memory_reject_percent = '70,90'
```

When the memory usage exceeds 90% of the upper limit, new connections are forbidden and stacked sessions are killed. When the memory usage is less than 70% of the upper limit, session killing is stopped and new connections are allowed.

---

**NOTICE**

- You can query the maximum dynamic memory and used dynamic memory in the **gs\_total\_memory\_detail** view. **max\_dynamic\_memory** indicates the maximum dynamic memory, and **dynamic\_used\_memory** indicates the used dynamic memory.
  - If this parameter is set to a small value, the escape from memory overload process is frequently triggered. As a result, ongoing sessions are forcibly logged out, and new connections fail to be connected for a short period of time. Therefore, exercise caution when setting this parameter based on the actual memory usage.
  - The values of **recover\_memory\_percent** and **overload\_memory\_percent** can be 0 at the same time. In addition, the value of **recover\_memory\_percent** must be smaller than that of **overload\_memory\_percent**. Otherwise, the setting does not take effect.
- 

### 17.3.4.2 Disk Space

This section describes the disk space parameters, which are used to set limits on the disk space for storing temporary files.

#### sql\_use\_spacelimit

**Parameter description:** Specifies the space size for files to be flushed to disks when a single SQL statement is executed on a single database node. The



managed space includes the space occupied by ordinary tables, temporary tables, and intermediate result sets to be flushed to disks.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to 2147483647. The unit is KB. -1 indicates no limit.

**Default value:** -1

## temp\_file\_limit

**Parameter description:** Specifies the limit on the size of a temporary file spilled to disk in a session. The temporary file can be a sort or hash temporary file, or the storage file for a held cursor.

This is a session-level setting.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

This parameter does not apply to disk space used for temporary tables during the SQL query process.

---

**Value range:** an integer ranging from -1 to 2147483647. The unit is KB. -1 indicates no limit.

**Default value:** -1

## 17.3.4.3 Kernel Resource Usage

This section describes kernel resource parameters. Whether these parameters take effect depends on OS settings.

## max\_files\_per\_process

**Parameter description:** Specifies the maximum number of simultaneously open files allowed by each server process. If the kernel is enforcing a proper limit, setting this parameter is not required.

However, on some platforms, such as most Berkeley Software Distribution (BSD) systems, the kernel allows individual processes to open many more files than the system can support. If the message "Too many open files" is displayed, set this parameter to a smaller value. Generally, the system must meet this requirement:  
Number of file descriptors  $\geq$  Maximum number of concurrent statements  $\times$   
Number of database nodes  $\times$  **max\_files\_per\_process**  $\times$  3

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 25 to 2147483647

**Default value:** 1024

## shared\_preload\_libraries

**Parameter description:** Specifies one or more shared libraries to be preloaded at server start. If multiple libraries are to be loaded, separate their names using commas (,). Only the sysadmin user can access this parameter. For example, **\$libdir/mylib** will cause **mylib.so** (or on some platforms, **mylib.sl**) to be preloaded before the loading of the standard library directory.

You can preinstall the GaussDB's stored procedure library using the **\$libdir/plXXX** syntax as described in the preceding text. **XXX** can only be **pgsql**, **perl**, **tcl**, or **python**.

By preloading a shared library and initializing it as required, the library startup time is avoided when the library is first used. However, the time to start each new server process may increase, even if that process never uses the library. Therefore, set this parameter only for libraries that will be used in most sessions.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

- If a specified library is not found, the GaussDB service will fail to start.
  - Each GaussDB-supported library has a special mark that is checked to guarantee compatibility. Therefore, libraries that do not support GaussDB cannot be loaded in this way.
- 

**Value range:** a string

**Default value:** **security\_plugin**

### 17.3.4.4 Cost-based Vacuum Delay

This feature allows administrators to reduce the I/O impact of the VACUUM and ANALYZE statements on concurrent database activities. It is often more important to prevent maintenance statements, such as VACUUM and ANALYZE, from affecting other database operations than to run them quickly. Cost-based vacuum delay provides a way for administrators to achieve this purpose.

---

### NOTICE

Certain vacuum operations hold critical locks and should be complete as quickly as possible. In GaussDB, cost-based vacuum delays do not take effect during such operations. To avoid uselessly long delays in such cases, the actual delay is the larger of the two calculated values:

- **vacuum\_cost\_delay** x **accumulated\_balance/vacuum\_cost\_limit**
  - **vacuum\_cost\_delay** x 4
- 

## Background

During the execution of the ANALYZE | ANALYSE and VACUUM statements, the system maintains an internal counter that keeps track of the estimated cost of the

various I/O operations that are performed. When the accumulated cost reaches a limit (specified by **vacuum\_cost\_limit**), the process performing the operation will sleep for a short period of time (specified by **vacuum\_cost\_delay**). Then, the counter resets and the operation continues.

By default, this feature is disabled. To enable this feature, set **vacuum\_cost\_delay** to a non-zero value.

## **vacuum\_cost\_delay**

**Parameter description:** Specifies the length of time that a process will sleep when **vacuum\_cost\_limit** has been exceeded.

On many systems, the effective resolution of the sleep length is 10 milliseconds. Therefore, setting **vacuum\_cost\_delay** to a value that is not a multiple of 10 has the same effect as setting it to the next higher multiple of 10.

This parameter is usually set to a small value, such as 10 or 20 milliseconds. To adjust the resource usage of this feature, you are advised to adjust other parameters instead of this parameter.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 100. A positive number enables cost-based vacuum delay and **0** disables cost-based vacuum delay.

**Default value:** 1

## **vacuum\_cost\_page\_hit**

**Parameter description:** Specifies the estimated cost for vacuuming a buffer found in the shared buffer. It represents the cost to lock the buffer pool, look up the shared hash table, and scan the content of the page.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 10000

**Default value:** 1

## **vacuum\_cost\_page\_miss**

**Parameter description:** Specifies the estimated cost for vacuuming a buffer read from the disk. It represents the cost to lock the buffer pool, look up the shared hash table, read the desired block from the disk, and scan the block.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 10000

**Default value:** 10

## **vacuum\_cost\_page\_dirty**

**Parameter description:** Specifies the estimated cost charged when vacuum modifies a block that was previously clean. It represents the extra cost required to update the dirty block out to the disk again.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 10000

**Default value:** 20

## **vacuum\_cost\_limit**

**Parameter description:** Specifies the cost limit. The vacuuming process will sleep if this limit is exceeded.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 10000

**Default value:** 1000

## **17.3.4.5 Background Writer**

This section describes background writer parameters. The background writer process is used to write dirty data (new or modified data) in shared buffers to disks. This mechanism ensures that database processes seldom or never need to wait for a write action to occur when handling user queries.

It also mitigates performance deterioration caused by checkpoints because only a few of dirty pages need to be flushed to the disk when the checkpoints arrive. This mechanism, however, increases the overall net I/O load because while a repeatedly-dirtied page may otherwise be written only once per checkpoint interval, the background writer may write it several times as it is dirtied in the same interval. In most cases, continuous light loads are preferred, instead of periodic load peaks. The parameters discussed in this section can be set based on actual requirements.

## **bgwriter\_delay**

**Parameter description:** Specifies the interval at which the background writer writes dirty shared buffers. Each time, the background writer process initiates write operations for some dirty buffers. In full checkpoint mode, the **bgwriter\_lru\_maxpages** parameter is used to control the amount of data to be written each time, and the process is restarted after *bgwriter\_delay* ms hibernation. In incremental checkpoint mode, the number of target idle buffer pages is calculated based on the value of **candidate\_buf\_percent\_target**. If the number of idle buffer pages is insufficient, a batch of pages is flushed to disks every *bgwriter\_delay* ms. The number of flushed pages is calculated based on the target difference percentage. The maximum number of flushed pages is limited by **max\_io\_capacity**.

In many systems, the effective resolution of sleep delays is 10 milliseconds. Therefore, setting this parameter to a value that is not a multiple of 10 has the same effect as setting it to the next higher multiple of 10.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 10000. The unit is millisecond.

**Default value:** 2s

**Setting suggestion:** Reduce this value in slow data writing scenarios to reduce the checkpoint load.

## candidate\_buf\_percent\_target

**Parameter description:** Specifies the expected percentage of available buffers in the **shared\_buffer** memory buffer in the candidate buffer chain when the incremental checkpoint is enabled. If the number of available buffers in the current candidate chain is less than the target value, the bgwriter thread starts flushing dirty pages that meet the requirements.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a double-precision floating point number ranging from 0.1 to 0.85

**Default value:** 0.3

## bgwriter\_lru\_maxpages

**Parameter description:** Specifies the number of dirty buffers the background writer can write in each round.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000

### NOTE

When this parameter is set to **0**, the background writer is disabled. This setting does not affect checkpoints.

**Default value:** 100

## bgwriter\_lru\_multiplier

**Parameter description:** Specifies the coefficient used to estimate the number of dirty buffers the background writer can write in the next round.

The number of dirty buffers written in each round depends on the number of buffers used by server processes during recent rounds. The estimated number of buffers required in the next round is calculated using the following formula: Average number of recently used buffers x **bgwriter\_lru\_multiplier**. The background writer writes dirty buffers until sufficient, clean and reusable buffers are available. The number of buffers the background writer writes in each round is always less than or equal to the value of **bgwriter\_lru\_maxpages**.

Therefore, the value **1.0** of **bgwriter\_lru\_multiplier** represents a just-in-time policy of writing exactly the number of dirty buffers predicted to be required. Larger values provide some cushion against spikes in demand, whereas smaller values intentionally leave more writes to be done by server processes.

Smaller values of **bgwriter\_lru\_maxpages** and **bgwriter\_lru\_multiplier** reduce the extra I/O load caused by the background writer, but make it more likely that server processes will have to issue writes for themselves, delaying interactive queries.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to 10

**Default value:** 2

## pagewriter\_thread\_num

**Parameter description:** Specifies the number of threads for background page flushing after the incremental checkpoint is enabled. Dirty pages are flushed in sequence to disks, promoting recovery points.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 16

**Default value:** 4

## dirty\_page\_percent\_max

**Parameter description:** Specifies the percentage of dirty pages to **shared\_buffers** after the incremental checkpoint is enabled. When the value of this parameter is reached, the background page flush thread flushes dirty pages based on the maximum value of **max\_io\_capacity**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0.1 to 1

**Default value:** 0.9

## pagewriter\_sleep

**Parameter description:** Specifies the interval for the pagewriter thread to flush dirty pages to disks after the incremental checkpoint is enabled. When the ratio of dirty pages to **shared\_buffers** reaches **dirty\_page\_percent\_max**, the number of pages in each batch is calculated based on the value of **max\_io\_capacity**. In other cases, the number of pages in each batch decreases proportionally.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600000. The unit is ms.

**Default value:** 2000ms(2s)

## max\_io\_capacity

**Parameter description:** Specifies the maximum I/O per second for the background writer process to flush pages in batches. Set this parameter based on the service scenario and disk I/O capability of the host. If the RTO is short or the data volume is much larger than the shared memory, and the service access data volume is random, the value of this parameter cannot be too small. A small parameter value reduces the number of pages flushed by the background writer process. If a large number of pages are eliminated due to service triggering, the services are affected.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 30720 to 10485760. The unit is KB.

**Default value:** 512000 KB (500 MB)

## **enable\_consider\_usecount**

**Parameter description:** Specifies whether the backend thread considers the page popularity during page replacement. You are advised to enable this parameter in large-capacity scenarios.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on/true:** The page popularity is considered.
- **off/false:** The page popularity is not considered.

**Default value:** off

## **dw\_file\_num**

**Parameter description:** Specifies the number of doublewrite files to be written in batches. The value is related to **pagewriter\_thread\_num** and cannot be greater than **pagewriter\_thread\_num**. If the value is too large, it will be corrected to the value of **pagewriter\_thread\_num**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 16

**Default value:** 1

## **dw\_file\_size**

**Parameter description:** Specifies the size of each doublewrite file.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer, in the range [32,256]

**Default value:** 256

### **17.3.4.6 Asynchronous I/O**

## **checkpoint\_flush\_after**

**Parameter description:** Specifies the threshold for the number of pages flushed by the checkpoint thread. If the threshold is exceeded, the operating system is instructed to flush the pages cached in the operating system asynchronously. In GaussDB, the disk page size is 8 KB.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 256. **0** indicates that the asynchronous flush function is disabled. For example, if the value is **32**, the checkpoint thread continuously writes 32 disk pages (that is, 32 x 8 = 256 KB) before asynchronous flush.

**Default value:** 256 KB

## bgwriter\_flush\_after

**Parameter description:** Specifies the threshold for the number of pages flushed by the background writer thread. If the threshold is exceeded, the background writer thread instructs the operating system to asynchronously flush the pages cached in the operating system to disks. In GaussDB, the disk page size is 8 KB.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 256. **0** indicates that the asynchronous flush function is disabled. The size of a single page is 8 KB. For example, if the value is **64**, the background writer thread continuously writes 64 disk pages (that is,  $64 \times 8 = 512$  KB) before asynchronous flush.

**Default value:** 512 KB (64 pages)

## backend\_flush\_after

**Parameter description:** Specifies the threshold for the number of pages flushed by the backend thread. If the number of pages exceeds the threshold, the backend thread instructs the operating system to asynchronously flush the pages cached in the operating system to disks. In GaussDB, the disk page size is 8 KB.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 256. **0** indicates that the asynchronous flush function is disabled. The size of a single page is 8 KB. For example, if the value is **64**, the backend thread continuously writes 64 disk pages (that is,  $64 \times 8 = 512$  KB) before asynchronous flush.

**Default value:** 0

## 17.3.5 Data Import and Export

This section describes the parameters for importing and exporting data.

### raise\_errors\_if\_no\_files

**Parameter description:** Specifies whether to distinguish between the problems "the number of imported file records is empty" and "the imported file does not exist". If this parameter is set to **TRUE** and the problem "the imported file does not exist" occurs, GaussDB will report the error message "file does not exist".

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the messages of "the number of imported file records is empty" and "the imported file does not exist" are distinguished when files are imported.
- **off** indicates that the messages of "the number of imported file records is empty" and "the imported file does not exist" are the same when files are imported.

**Default value:** off



## safe\_data\_path

**Parameter description:** Specifies the path prefix restriction except for the initial user. Currently, the restrictions are posed on COPY and advanced package paths. The path cannot end with a slash (/) or contain periods (..).

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string of up to 4096 characters

**Default value:** NULL

## enable\_copy\_server\_files

**Parameter description:** Specifies whether to enable the permission to copy server files.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the permission to copy server files is enabled.
- **off** indicates that the permission to copy server files is disabled.

**Default value:** off

---

### NOTICE

When the **enable\_copy\_server\_files** parameter is disabled, only the initial user is allowed to run the **COPY FROM FILENAME** or **COPY TO FILENAME** statement. When the **enable\_copy\_server\_files** parameter is enabled, users with the SYSADMIN permission or users who inherit the **gs\_role\_copy\_files** permission of the built-in role are allowed to run the **COPY FROM FILENAME** or **COPY TO FILENAME** statement.

---

## 17.3.6 Write Ahead Log

### 17.3.6.1 Settings

#### wal\_level

**Parameter description:** Specifies the level of information to be written to the WAL. The value cannot be empty or commented out.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**NOTICE**

- To enable WAL archiving and data streaming replication between primary and standby servers, set this parameter to **archive**, **hot\_standby**, or **logical**.
- If this parameter is set to **minimal**, **hot\_standby** must be set to **off** and **max\_wal\_senders** must be set to **0**. In addition, the database must be deployed in a standalone system. Otherwise, the database cannot be started.
- If this parameter is set to **archive**, **hot\_standby** must be set to **off**. Otherwise, the database startup fails. However, **hot\_standby** cannot be set to **off** in a primary-standby deployment. For details, see the description of the **hot\_standby** parameter.

**Value range:** enumerated values

- **minimal**  
Advantages: Certain bulk operations (including creating tables and indexes, executing cluster operations, and copying tables) are safely skipped in logging, which can make those operations much faster.  
Disadvantages: WALs contain only basic information required for recovery from a database server crash or an emergency shutdown. Data cannot be restored from archived WALs.
- **archive**  
Adds logging required for WAL archiving, supporting the database restoration from archives.
- **hot\_standby**
  - Further adds information required to run SQL queries on a standby server and takes effect after a server restart.
  - To enable read-only queries on a standby server, the **wal\_level** parameter must be set to **hot\_standby** on the primary server and the same value must be set on the standby server. There are few measurable differences in performance between using **hot\_standby** and **archive** levels. However, feedback is welcome if any differences in their impacts on product performance are noticeable.
- **logical**  
This parameter indicates that WALs support logical replication.

**Default value:** **hot\_standby**

## **fsync**

**Parameter description:** Specifies whether the GaussDB server uses the **fsync()** function (see [wal\\_sync\\_method](#)) to ensure that updates can be written to disks in a timely manner.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- Using the **fsync()** function ensures that the data can be recovered to a known state when an OS or a hardware crashes.
  - Setting this parameter to **off** may result in unrecoverable data corruption in a system crash.
- 

**Value range:** Boolean

- **on** indicates that the **fsync()** function is used.
- **off** indicates that the **fsync()** function is not used.

**Default value:** on

## synchronous\_commit

**Parameter description:** Specifies the synchronization mode of the current transaction.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

Generally, logs generated by a transaction are synchronized in the following sequence:

1. The primary node writes the log content to the local memory.
2. The primary node writes the logs in the local memory to the local file system.
3. The primary node flushes logs in the local file system to disks.
4. The primary node sends the log content to the standby node.
5. The standby node receives the logs and saves them to the local memory.
6. The standby node writes logs in the local memory to the local file system.
7. The standby node flushes logs in the local file system to disks.
8. The standby node replays logs to complete the incremental update of data files.

**Value range:** enumerated values

- **on:** The primary node waits for the standby node to flush logs to disks before committing a transaction.
- **off:** The primary node commits transactions without waiting for the primary node to flush logs to disks. This mode is also called asynchronous commit.
- **local:** The primary node waits for the primary node to flush logs to disks before committing a transaction. This mode is also called local commit.
- **remote\_write:** The primary node waits for the standby node to write logs to the file system before committing a transaction. (The logs do not need to be flushed to disks.)
- **remote\_receive:** The primary node waits for the standby node to receive logs before committing a transaction. (The logs do not need to be written to the file system.)
- **remote\_apply:** The primary node waits for the standby node to complete log replay before committing a transaction.

- **true**: same as **on**.
- **false**: same as **off**.
- **yes**: same as **on**.
- **no**: same as **off**.
- **1**: same as **on**.
- **0**: same as **off**.
- **2**: same as **remote\_apply**.

**Default value:** on

## wal\_sync\_method

**Parameter description:** Specifies the method used for forcing WAL updates out to disk.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

If **fsync** is set to **off**, the setting of this parameter does not take effect because WAL updates will not be forcibly written to disks.

---

**Value range:** enumerated values

- **open\_datasync** indicates that WAL files are opened with the O\_DSYNC option.
- **fdasync** indicates that **fdasync()** is called at each commit. (SUSE 10 and SUSE 11 are supported.)
- **fsync\_writethrough** indicates that **fsync()** is called at each commit to force data in the buffer to be written to the disk.

### NOTE

**wal\_sync\_method** can be set to **fsync\_writethrough** on a Windows platform, but this setting has the same effect as setting the parameter to **fsync**.

- **fsync** indicates that **fsync()** is invoked at each commit (SUSE 10 and SUSE 11 are supported).
- **open\_sync** indicates that the **open()** with the O\_SYNC option is used to write WAL files (SUSE 10 and SUSE 11 are supported).

### NOTE

Not all platforms support the preceding parameters.

**Default value:** fdasync

## full\_page\_writes

**Parameter description:** Specifies whether the GaussDB server writes the entire content of each disk page to WALs during the first modification of that page after a checkpoint. When the incremental checkpoint function and **enable\_double\_write** are enabled at the same time, **full\_page\_writes** is not used.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- This parameter is needed because a page write that is in process during an OS crash might be only partially completed, leading to an on-disk page that contains a mix of old and new data. The row-level change data normally stored in WALs will not be enough to completely restore such a page during post-crash recovery. Storing the full page image guarantees that the page can be correctly restored, but at the price of increasing the amount of data that must be written to WALs.
  - Setting this parameter to **off** might lead to unrecoverable data corruption after a system failure. It might be safe to set this parameter to **off** if you have hardware (such as a battery-backed disk controller) or file-system software (such as ReiserFS 4) that reduces the risk of partial page writes to an acceptably low level.
- 

**Value range:** Boolean

- **on** indicates that this feature is enabled.
- **off** indicates that this feature is disabled.

**Default value:** on

## wal\_log\_hints

**Parameter description:** Specifies whether to write an entire page to WALs during the first modification of that page after a checkpoint, even for non-critical modifications of so-called hint bits. You are advised not to modify the setting.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the entire page is written to WALs.
- **off** indicates that the entire page is not written to WALs.

**Default value:** on

## wal\_buffers

**Parameter description:** Specifies the number of **XLOG\_BLCKSZ** used for storing WAL data. The size of each **XLOG\_BLCKSZ** is 8 KB.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** -1 to  $2^{18}$ . The minimum value is -1 and the maximum value is 262144. The unit is 8 KB.

- If this parameter is set to -1, the value of **wal\_buffers** is automatically adjusted to 1/64 of **shared\_buffers**. The minimum value is 8 **XLOG\_BLCKSZ** and the maximum value is 2048 **XLOG\_BLCKSZ**, if the automatically adjusted value is less than the minimum value, the value is adjusted to the minimum value. If the automatically adjusted value is greater than the maximum value, the value is adjusted to the maximum value.

- If this parameter is set to a value other than **-1** and smaller than **4**, the value **4** is used.

**Default value:**

**1 GB** (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory, 96-core CPU/768-GB memory, 64-core CPU/512-GB memory, 60-core CPU/480-GB memory, and 32-core CPU/256-GB memory); **512 MB** (16-core CPU/128-GB memory); **256 MB** (8-core CPU/64-GB memory);

**Setting suggestions:** The content of the WAL buffers is written to disks at every transaction commit, so extremely large values are unlikely to provide a significant increase in system performance. However, setting this parameter to hundreds of megabytes can improve the disk write performance on a server to which a large number of transactions are committed at the same time. The default value meets user requirements in most cases.

## wal\_writer\_delay

**Parameter description:** Specifies the delay between activity rounds for the WAL writer.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

A longer delay might lead to insufficient WAL buffer and a shorter delay leads to continuously writing of the WALs, thereby increasing the load of disk I/O.

---

**Value range:** an integer ranging from 1 to 10000. The unit is millisecond.

**Default value:** 200ms

## commit\_delay

**Parameter description:** Specifies the duration for committed data to be stored in the WAL buffer.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- When this parameter is set to a non-zero value, the committed transaction is stored in the WAL buffer instead of being written to the WAL immediately. Then the WAL writer process flushes the buffer out to disks periodically.
  - If system load is high, other transactions are probably ready to be committed within the delay. If no other transactions are ready to be committed, the delay is a waste of time.
- 

**Value range:** an integer ranging from 0 to 100000. The unit is microsecond. **0** indicates no delay.

**Default value:** 0

## commit\_siblings

**Parameter description:** Specifies a threshold on the number of concurrent open transactions. If the number of concurrent open transactions is greater than the value of this parameter, a transaction that initiates a commit request will wait for a period of time specified by [commit\\_delay](#). Otherwise, this transaction is written into WALs immediately.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000

**Default value:** 5

## wal\_block\_size

**Parameter description:** Specifies the size of a WAL disk block.

This parameter is a fixed INTERNAL parameter and cannot be modified.

**Value range:** an integer. The unit is byte.

**Default value:** 8192

## wal\_segment\_size

**Parameter description:** Specifies the size of a WAL segment file.

This parameter is a fixed INTERNAL parameter and cannot be modified.

**Value range:** an integer. The unit is 8 KB.

**Default value:** 16 MB (2048 x 8 KB)

## walwriter\_cpu\_bind

**Parameter description:** Specifies the CPU core bound to the WAL write thread. This parameter is used together with the thread pool parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to the number of cores minus 1

**Default value:** -1

## walwriter\_sleep\_threshold

**Parameter description:** Specifies the number of times that idle Xlogs are refreshed before the xlogflusher enters the sleep state. If the number of times reaches the threshold, the xlogflusher enters the sleep state.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 50000

**Default value:** 500 (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory, 96-core CPU/768-GB memory, 64-core CPU/512-GB memory, and 60-core CPU/480-GB memory); 50 (32-core CPU/256-GB memory, 16-core CPU/128-GB memory, 8-core CPU/64-GB memory, )

## wal\_file\_init\_num

**Parameter description:** Specifies the number of Xlog segment files to be created by the WAL writer.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000

**Default value:** **10** (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory, 96-core CPU/768-GB memory, 64-core CPU/512-GB memory, and 60-core CPU/480-GB memory); **0** (32-core CPU/256-GB memory, 16-core CPU/128-GB memory, 8-core CPU/64-GB memory, )

## xlog\_file\_path

**Parameter description:** Specifies the path of the Xlog shared disk in dual-database instance shared storage scenarios. This parameter is configured by the OM during database system initialization. You are not advised to modify the configuration.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** NULL

## xlog\_file\_size

**Parameter description:** Specifies the size of the Xlog shared disk in dual-database instance shared storage scenarios. This parameter is configured by the OM during database system initialization. You are not advised to modify the configuration.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a long integer ranging from 5053733504 to 576460752303423487. The unit is byte.

**Default value:** 549755813888

## xlog\_lock\_file\_path

**Parameter description:** Specifies the path of the lock file preempted by the Xlog shared disk in dual-database instance shared storage scenarios. This parameter is configured by the OM during database system initialization. You are not advised to modify the configuration.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** NULL

## force\_promote

**Parameter description:** Specifies whether to enable the forcible switchover function on the standby node.



When a database instance is faulty, the forcible switchover enables the database instance to recover services as soon as possible at the cost of losing some data. This is an escape method used when the database instance is unavailable. You are not advised to trigger this method frequently. You are not advised not to use this function if you are not clear about the impact of data loss on services.

To use this function, you need to enable it on the DN and CM Server and restart the database instance for the setting to take effect. For details about how to enable the forcible switchover function on the standby node, see "Emergency Handling" > "Performing a Forcible Primary/Standby Switchover" in *Troubleshooting*.

**Value range:** 0 or 1

The value **0** indicates that the function is disabled, and the value **1** indicates that the function is enabled.

**Default value:** 0

## wal\_debug

**Parameter description:** Specifies whether to output WAL-related debugging information. This parameter is available only when **WAL\_DEBUG** is enabled during compilation.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** false

## wal\_flush\_timeout

**Parameter description:** Specifies the timeout interval for traversing **WalInsertStatusEntryTbl**. It is the maximum wait time for the adaptive Xlog disk flushing I/O to traverse **WalInsertStatusEntryTbl**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

If the timeout interval is too long, the Xlog flushing frequency may decrease, reducing the Xlog processing performance.

---

**Value range:** an integer ranging from 0 to 900000000. The unit is microsecond.

**Default value:** 2

## wal\_flush\_delay

**Parameter description:** Specifies the wait interval when an entry in the **WAL\_NOT\_COPIED** state is encountered during **WalInsertStatusEntryTbl** traversal.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 900000000. The unit is microsecond.

**Default value:** 1

## 17.3.6.2 Checkpoints

### checkpoint\_segments

**Parameter description:** Specifies the minimum number of WAL segment files in the period specified by [checkpoint\\_timeout](#). The size of each log file is 16 MB.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. The minimum value is 1.

Increasing the value of this parameter speeds up the import of a large amount of data. Set this parameter based on [checkpoint\\_timeout](#) and [shared\\_buffers](#). This parameter affects the number of WAL segment files that can be reused. Generally, the maximum number of reused files in the **pg\_xlog** folder is twice the number of **checkpoint\_segments**. The reused files are not deleted and are renamed to the WAL segment files which will be later used.

**Default value:** 1024

### checkpoint\_timeout

**Parameter description:** Specifies the maximum time between automatic WAL checkpoints.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 30 to 3600. The unit is s.

If the value of [checkpoint\\_segments](#) is increased, you need to increase the value of this parameter. The increase of these two parameters further requires the increase of [shared\\_buffers](#). Consider all these parameters during setting.

**Default value:** 15min

### checkpoint\_completion\_target

**Parameter description:** Specifies the completion target of each checkpoint, as a fraction of total time between checkpoints.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a double-precision floating point number ranging from 0.0 to 1.0

**Default value:** 0.5

#### NOTE

**0.5** indicates that each checkpoint should be complete within 50% of the interval between checkpoints.

## checkpoint\_warning

**Parameter description:** Specifies a time in seconds. If the checkpoint interval is close to this time due to filling of checkpoint segment files, a message is sent to the server log to suggest an increase in the value of [checkpoint\\_segments](#).

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is second. **0** indicates that the warning is disabled.

**Default value:** 5min

**Recommended value:** 5min

## checkpoint\_wait\_timeout

**Parameter description:** Sets the longest time that the checkpoint waits for the checkpoint thread to start.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 2 to 3600. The unit is s.

**Default value:** 1min

## enable\_incremental\_checkpoint

**Parameter description:** Specifies whether to enable incremental checkpointing.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** on

## enable\_double\_write

**Parameter description:** Specifies whether to enable double writing. When the incremental checkpoint function is enabled and **enable\_double\_write** is enabled, the **enable\_double\_write** dual-write feature is used for protection, and **full\_page\_writes** is not used to prevent half-page write.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** on

## incremental\_checkpoint\_timeout

**Parameter description:** Specifies the maximum interval between automatic WAL checkpoints when the incremental checkpointing is enabled.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 3600. The unit is s.

**Default value:** 1min

## enable\_xlog\_prune

**Parameter description:** Specifies whether the primary node recycles logs if the size of Xlogs exceeds the value of **max\_size\_for\_xlog\_prune** when any standby node is disconnected.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- If this parameter is set to **on**, the primary node recycles logs when any standby node is disconnected.
- If this parameter is set to **off**, the primary node does not recycle logs when any standby node is disconnected.

**Default value:** on

## max\_size\_for\_xlog\_prune

**Parameter description:** This parameter takes effect when **enable\_xlog\_prune** is enabled. The working mechanism is as follows:

1. If all standby nodes specified by the **replconninfo** series GUC parameters are connected to the primary node, this parameter does not take effect.
2. If any standby node specified by the **replconninfo** series GUC parameters is not connected to the primary node, this parameter takes effect. When the number of historical logs on the primary node is greater than the value of this parameter, the logs are forcibly recycled. Exception: In synchronous commit mode (that is, the value of **synchronous\_commit** is not **local** or **off**), if there are connected standby nodes, the primary node retains the logs that meet the minimum log receiving requirements on the majority of standby nodes. In this case, the number of reserved logs may exceed the value of **max\_size\_for\_xlog\_prune**.
3. If any standby node is being built, this parameter does not take effect. All logs of the primary node are retained to prevent build failures due to log recycling.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647. The unit is KB.

**Default value:** 256 GB

### 17.3.6.3 Log Replay

## recovery\_time\_target

**Parameter description:** Specifies the time for a standby server to write and replay logs.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600. The unit is s.

**0** indicates that log flow control is disabled. A value from **1** to **3600** indicates that a standby server can write and replay logs within the period specified by the value, so that the standby server can quickly assume the primary role. If this parameter

is set to a small value, the performance of the primary node is affected. If it is set to a large value, the log flow is not effectively controlled.

**Default value:** 60

## recovery\_max\_workers

**Parameter description:** Specifies the maximum number of concurrent replay threads.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 20

**Default value:** 4 (For better performance, the default value in tool installation is 4.)

## recovery\_parse\_workers

**Parameter description:** Specifies the number of **ParseRedoRecord** threads for the ultimate RTO feature.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 16

**recovery\_parse\_workers** can be set to a value greater than 1 only when the ultimate RTO feature is enabled. In addition, it must be used together with **recovery\_redo\_workers**. If both **recovery\_parse\_workers** and **recovery\_max\_workers** are enabled, the setting of **recovery\_parse\_workers** prevails and the concurrent replay function is disabled. The ultimate RTO feature does not support the hot standby mode. Therefore, **recovery\_parse\_workers** can be set to a value greater than 1 only when **hot\_standby** is set to **off** and **replication\_type** to 1.

**Default value:** 1

## recovery\_redo\_workers

**Parameter description:** Specifies the number of **PageRedoWorker** threads corresponding to each **ParseRedoRecord** thread when the ultimate RTO feature is enabled.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 8

This parameter must be used together with **recovery\_parse\_workers**. The setting of **recovery\_redo\_workers** takes effect only when **recovery\_parse\_workers** is set to a value greater than 0.

**Default value:** 1

## recovery\_parallelism

**Parameter description:** Specifies the actual number of replay threads. This parameter is read-only.

This is a POSTMASTER parameter and is affected by **recovery\_max\_workers** and **recovery\_parse\_workers**. If any value is greater than 0, **recover\_parallelism** will be recalculated.

**Value range:** an integer ranging from 1 to 2147483647

**Default value:** 1

## enable\_page\_lsn\_check

**Parameter description:** Specifies whether to enable the data page LSN check. During replay, the current LSN of the data page is checked to see if it is the expected one.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** on

## recovery\_min\_apply\_delay

**Parameter description:** Specifies the replay delay of the standby node.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

- This parameter does not take effect on the primary node. It must be set on the standby node that requires a delay. You are advised to set this parameter on the asynchronous standby node. If the delay is set on the asynchronous standby node, the RTO will be long after the node is promoted to primary.
  - The delay time is calculated based on the transaction commit timestamp on the primary server and the current time on the standby server. Therefore, ensure that the clocks of the primary and standby servers are synchronized.
  - If the delay time is too long, the disk where the Xlog file is located on the standby node may be full. Therefore, you need to set the delay time based on the disk size.
  - Operations without transactions are not delayed.
  - After the primary/standby switchover, if the original primary node needs to be delayed, you need to manually set this parameter.
  - When **synchronous\_commit** is set to **remote\_apply**, synchronous replication is affected by the delay. Each commit message is returned only after the replay on the standby server is complete.
  - Using this feature also delays **hot\_standby\_feedback**, which may cause the primary server to bloat, so be careful when using both.
  - If a DDL operation (such as DROP or TRUNCATE) that holds an AccessExclusive lock is performed on the primary node, the query operation on the operation object on the standby node will be returned only after the lock is released during the delayed replay of the record on the standby node.
  - The MOT table is not supported.
-

**Value range:** an integer ranging from 0 to INT\_MAX. The unit is ms.

**Default value:** 0 (no delay added)

## redo\_bind\_cpu\_attr

**Parameter description:** Specifies the core binding operation of the replayer thread. Only the **sysadmin** user can access this parameter. This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string of more than 0 characters. The value is case-insensitive.

- **'nobind'**: The thread is not bound to a core.
- **'nodebind: 1, 2'**: Use the CPU cores in NUMA groups 1 and 2 to bind threads.
- **'cpubind: 0-30'**: Use the CPU cores 0 to 30 to bind threads.

**Default value:** 'nobind'

## 17.3.6.4 Archiving

### archive\_timeout

**Parameter description:** Specifies the archiving period.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

- The server is forced to switch to a new WAL segment file when the period specified by this parameter has elapsed since the last file switch.
- Archived files that are closed early due to a forced switch are still of the same length as full files. Therefore, a very short **archive\_timeout** will bloat the archive storage. You are advised to set **archive\_timeout** to 60s.

---

**Value range:** an integer ranging from 0 to 1073741823. The unit is s. 0 indicates that archiving timeout is disabled.

**Default value:** 0

### archive\_interval

**Parameter description:** Specifies the archiving interval.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

- Log files are forcibly archived when the period specified by this parameter has elapsed.
  - Archiving involves I/O operations. Therefore, frequent archiving is not allowed. In addition, the RPO cannot be set to a large value; otherwise, the PITR will be affected. You are advised to use the default value.
-

**Value range:** an integer ranging from 1 to 1000. The unit is s.

**Default value:** 1

## time\_to\_target\_rpo

### Parameter description:

In the remote DR mode of dual database instances, set the allowed *time\_to\_target\_rpo* seconds from the time when an exception occurs on the primary database instance to the time when data is archived to the OBS recovery point.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600. The unit is s.

In remote DR mode of dual database instances, logs of the primary database instance are archived to OBS. **0** indicates that log flow control is disabled. 1 to 3600 indicates the maximum *time\_to\_target\_rpo* seconds from the time when an exception occurs on the primary database instance to the time when data is archived to the recovery point of OBS. This ensures that the maximum duration of data loss is within the allowed range when the primary database instance breaks down due to a disaster. If this parameter is set to a small value, the performance of the primary node is affected. If it is set to a large value, the log flow is not effectively controlled.

**Default value:** 0

## 17.3.7 HA Replication

### 17.3.7.1 Sending Server

#### max\_wal\_senders

**Parameter description:** Specifies the maximum number of simultaneously running WAL sender processes. The value cannot be equal to or greater than that of [max\\_connections](#).

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

[wal\\_level](#) must be set to **archive**, **hot\_standby**, or **logical** to allow the connection from standby servers.

---

**Value range:** an integer ranging from 0 to 1024. The recommended value range is 8 to 100.

#### NOTE

This parameter can be set to 0 only when a single DN is used and there is no primary/standby instance.



**Default value:** 20

## wal\_keep\_segments

**Parameter description:** Specifies the number of Xlog file segments. Specifies the minimum number of transaction log files stored in the **pg\_xlog** directory. The standby node obtains log files from the primary node for streaming replication.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 2 to *INT\_MAX*

**Default value:** 1024

**Setting suggestions:**

- During WAL archiving or recovery from a checkpoint on the server, the system may retain more log files than the number specified by **wal\_keep\_segments**.
- If this parameter is set to an excessively small value, a transaction log may have been overwritten by a new transaction before requested by the standby server. As a result, the request fails and the connection between the primary and standby servers is terminated.
- If the HA system uses asynchronous transmission, increase the value of **wal\_keep\_segments** when data greater than 4 GB is continuously imported in COPY mode. Take T6000 board as an example. If the data to be imported reaches 50 GB, you are advised to set this parameter to **1000**. You can dynamically restore the setting of this parameter after data import is complete and the WAL synchronization is proper.
- If the synchronous commit level is lower than **LOCAL\_FLUSH**, you are advised to set this parameter to **1000** when rebuilding the standby node to prevent rebuilding failures caused by primary node log recycling during the rebuilding.

## wal\_sender\_timeout

**Parameter description:** Specifies the maximum duration that the sending server waits for the WAL reception in the receiver.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

- If the data volume on the primary node is huge, the value of this parameter must be increased for the database rebuilding on a standby node. For example, if the data volume on the primary node reaches 500 GB, you are advised to set this parameter to 600 seconds.
  - This parameter cannot be set to a value larger than the value of **wal\_receiver\_timeout** or the timeout parameter for database rebuilding.
- 

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is ms.

**Default value:** 6s

## max\_replication\_slots

**Parameter description:** Specifies the number of log replication slots on the primary node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1024. The recommended value range is 8 to 100.

**Default value:** 20

### Setting suggestions:

When HA replication, backup and restoration, and logical decoding are used, you are advised to set this parameter to: Number of current physical replication slots + Number of backup slots + Number of required logical replication slots. If the actual value is smaller than the recommended value, these functions may be unavailable or abnormal.

- Physical streaming replication slots provide an automatic method to ensure that Xlogs are not removed from a primary node before they are received by all the standby nodes. That is, physical replication slots are used to support primary/standby HA. The number of physical streaming replication slots required by a database is equal to the ratio of standby nodes to the primary node. For example, if an HA database has 1 primary node and 2 standby nodes, the number of required physical replication slots will be 2. If an HA database has 1 primary node and 3 standby nodes, the number of required physical replication slots will be 3.
- Backup slot records replication information during backup execution. Full backup and incremental backup correspond to two independent backup slots.
- Plan the number of logical replication slots as follows:
  - A logical replication slot can carry changes of only one database for decoding. If multiple databases are involved, create multiple logical replication slots.
  - If logical replication is needed by multiple target databases, create multiple logical replication slots in the source database. Each logical replication slot corresponds to one logical replication link.

## enable\_slot\_log

**Parameter description:** Specifies whether to enable primary/standby synchronization for logical replication slots.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that primary/standby synchronization is enabled for logical replication slots.
- **off** indicates that primary/standby synchronization is disabled for logical replication slots.

**Default value:** on

## max\_changes\_in\_memory

**Parameter description:** Specifies the maximum number of DML statements cached in memory for a single transaction during logical decoding.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647

**Default value:** 4096

## max\_cached\_tuplebufs

**Parameter description:** Specifies the upper limit of the total tuple information cached in the memory during logical decoding. You are advised to set this parameter to a value equal to or greater than twice of [max\\_changes\\_in\\_memory](#).

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647

**Default value:** 8192

## logical\_decode\_options\_default

**Parameter description:** Specifies the global default value for unspecified decoding options when logical decoding starts.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

Currently, the following logical decoding options are supported: **parallel-decode-num**, **parallel-queue-size**, **max-txn-in-memory**, **max-reorderbuffer-in-memory**, and **exclude-users**.

**Value range:** a string of key=value characters separated by commas (,), for example, '**parallel-decode-num=4,parallel-queue-size=128,exclude-users=userA**'. An empty string indicates that the default value hardcoded by the program is used.

**Default value:** empty

---

### NOTICE

The SIGHUP parameter does not affect the started logic decoding process. The options specified by this parameter are used as the default settings for subsequent logic decoding startup, and the settings specified in the startup command are preferentially used.

The **exclude-users** option is different from the logic decoding startup option. You are not allowed to specify multiple blacklisted users.

---

## logical\_sender\_timeout

**Parameter description:** Specifies the maximum waiting time for the sender to wait for the receiver to receive logical logs.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647. The unit is ms.

**Default value:** 30s

## enable\_wal\_shipping\_compression

**Parameter description:** Specifies whether to enable cross-database instance log compression in streaming DR mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

This parameter applies only to a pair of walsender and walreceiver for cross-database instance transmission in streaming DR and is configured on the primary database instance.

---

**Value range:** Boolean

- **true** indicates that cross-database instance log compression is enabled for streaming DR.
- **false** indicates that cross-database instance log compression is disabled for streaming DR.

**Default value:** false

## repl\_auth\_mode

**Parameter description:** Specifies the validation mode for primary/standby replication and standby node rebuilding.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- If UUID validation is enabled on the primary node and a non-null **repl\_uuid** validation code is configured, UUID validation must also be enabled on the standby node and the same **repl\_uuid** validation code must be configured on the standby node. Otherwise, requests for log replication between the primary and standby nodes and standby node rebuilding will be rejected by the primary node.
- The SIGHUP parameter can dynamically load new values. The modification does not affect the established primary/standby connection and takes effect for subsequent primary/standby replication requests and primary/standby rebuilding requests.
- It supports the standby node rebuild validation under the Quorum and DCF protocols and the primary/standby replication validation under the Quorum protocol. It does not support primary/standby replication validation under the DCF protocol.
- Authentication between primary and standby database instances is not supported, including primary and standby Dorado instances and DR instances.
- The UUID validation function is used to prevent data crosstalk and pollution caused by incorrect connection between the primary and standby nodes. It is not used for security purposes.
- This parameter cannot be automatically synchronized between the primary and standby nodes.

---

**Value range:** enumerated values

- **off:** indicates that UUID validation is disabled.
- **default:** indicates that UUID validation is disabled.
- **uuid:** indicates that UUID validation is enabled.

**Default value:** default

## repl\_uuid

**Parameter description:** Specifies the UUID used for primary/standby UUID validation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**NOTICE**

- If UUID validation is enabled on the primary node and a non-null **repl\_uuid** validation code is configured, UUID validation must also be enabled on the standby node and the same **repl\_uuid** validation code must be configured on the standby node. Otherwise, requests for log replication between the primary and standby nodes and standby node rebuilding will be rejected by the primary node.
- The SIGHUP parameter can dynamically load new values. The modification does not affect the established primary/standby connection and takes effect for subsequent primary/standby replication requests and primary/standby rebuilding requests.
- It supports the standby node rebuild validation under the Quorum and DCF protocols and the primary/standby replication validation under the Quorum protocol. It does not support primary/standby replication validation under the DCF protocol.
- Authentication between primary and standby database instances is not supported, including primary and standby Dorado instances and DR instances.
- The UUID validation function is used to prevent data crosstalk and pollution caused by incorrect connection between the primary and standby nodes. It is not used for security purposes.
- This parameter cannot be automatically synchronized between the primary and standby nodes.

**Value range:** a string. The value is a string of 0 to 63 case-insensitive letters and digits. It is converted to lowercase letters for storage. An empty string indicates that UUID validation is disabled.

**Default value:** empty

## replconninfo1

**Parameter description:** Specifies the information about the first node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the first node is configured.

**Default value:** empty

## replconninfo2

**Parameter description:** Specifies the information about the second node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the second node is configured.

**Default value:** empty

## replconninfo3

**Parameter description:** Specifies the information about the third node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the third node is configured.

**Default value:** empty

## replconninfo4

**Parameter description:** Specifies the information about the fourth node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the fourth node is configured.

**Default value:** empty

## replconninfo5

**Parameter description:** Specifies the information about the fifth node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the fifth node is configured.

**Default value:** empty

## replconninfo6

**Parameter description:** Specifies the information about the sixth node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the sixth node is configured.

**Default value:** empty

## replconninfo7

**Parameter description:** Specifies the information about the seventh node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the seventh node is configured.

**Default value:** empty

## replconninfo8

**Parameter description:** Specifies the information about the eighth node to be listened on and authenticated by the current server.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the eighth node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo1

**Parameter description:** Specifies the information about the local first node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the first node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo2

**Parameter description:** Specifies the information about the local second node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the second node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo3

**Parameter description:** Specifies the information about the local third node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the third node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo4

**Parameter description:** Specifies the information about the local fourth node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the fourth node is configured.

**Default value:** empty



## cross\_cluster\_replconninfo5

**Parameter description:** Specifies the information about the local fifth node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the fifth node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo6

**Parameter description:** Specifies the information about the local sixth node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the sixth node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo7

**Parameter description:** Specifies the information about the local seventh node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the seventh node is configured.

**Default value:** empty

## cross\_cluster\_replconninfo8

**Parameter description:** Specifies the information about the local eighth node to be listened on and authenticated across database instances.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the eighth node is configured.

**Default value:** empty

## available\_zone

**Parameter description:** Specifies the region where the local node is located.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that no information about the node is configured.

**Default value:** empty

## enable\_availablezone

**Parameter description:** Specifies whether the local cascaded standby node can connect to standby nodes across AZs.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **true** indicates that the cascaded standby node can only connect to standby nodes in the same AZ.
- **false** indicates that the cascaded standby node can connect to standby nodes across AZs.

**Default value:** false

## max\_keep\_log\_seg

**Parameter description:** Stream control parameter. In logical replication, physical logs are parsed and converted into logical logs locally on the DN. When the number of physical log files that are not parsed is greater than the value of this parameter, stream control is triggered. The value **0** indicates that the stream control function is disabled.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647

**Default value:** 0

### 17.3.7.2 Primary Server

## synchronous\_standby\_names

**Parameter description:** Specifies a comma-separated list of names of potential standby nodes that support synchronous replication.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

#### NOTICE

- The current synchronous standby node is on the top of the list. If the current synchronous standby node is disconnected, it will be replaced immediately with the next-highest-priority standby node. Name of the next-highest-priority standby node is added to the list.
- The standby node name can be specified by setting the environment variable *PGAPPNAME*.

**Value range:** a string If this parameter is set to \*, the name of any standby node that provides synchronous replication is matched. The value can be configured in the following format:

- ANY *num\_sync* (*standby\_name* [, ...]) [, ANY *num\_sync* (*standby\_name* [, ...])]
- [FIRST] *num\_sync* (*standby\_name* [, ...])

- *standby\_name* [, ...]

 **NOTE**

- *num\_sync* indicates the number of synchronous standby nodes that need to wait for responses before the transaction is committed, *standby\_name* indicates the name of the standby node, and **FIRST** and **ANY** specify the policies for selecting standby nodes for synchronous replication from the listed servers.
- **ANY N (dn\_instanceld1, dn\_instanceld2,...)** indicates that any *N* host names in the brackets are selected as the name list of standby nodes for synchronous replication. For example, **ANY 1 (dn\_instanceld1, dn\_instanceld2)** indicates that any one of **dn\_instanceld1** and **dn\_instanceld2** is used as the standby node for synchronous replication.
- **FIRST N (dn\_instanceld1, dn\_instanceld2, ...)** indicates that the first *N* primary node names in the brackets are selected as the standby node name list for synchronous replication based on the priority. For example, **FIRST 1 (dn\_instanceld1, dn\_instanceld2)** indicates that **dn\_instanceld1** is selected as the standby node for synchronous replication.
- The meanings of *dn\_instanceld1*, *dn\_instanceld2*, ... are the same as those of **FIRST 1 (dn\_instanceld1, dn\_instanceld2, ...)**.

If you use the *gs\_guc* tool to set this parameter, perform the following operations:

```
gs_guc reload -Z datanode -N @NODE_NAME@ -D @DN_PATH@ -c "synchronous_standby_names='ANY
NODE 1(dn_instanceld1, dn_instanceld2)';"
```

or

```
gs_guc reload -Z datanode -N @NODE_NAME@ -D @DN_PATH@ -c "synchronous_standby_names='ANY
1(AZ1, AZ2)';"
```

**Default value:** \*

## most\_available\_sync

**Parameter description:** Specifies whether transactions on the primary node are not blocked due to faults on synchronous standby nodes. For example, if one of the two synchronous standby nodes is faulty and the other is normal, the primary node waits for the normal synchronous standby node instead of being blocked by the faulty synchronous standby node.

For another example, when the quorum protocol is executed, one primary node and three synchronous standby nodes are configured. **ANY 2 (node1, node2, and node3)** is configured. When node1 and node3 are faulty and node2 is normal, host services are not blocked.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the primary node is not blocked when all synchronous standby nodes are faulty.
- **off** indicates that the primary node is blocked when all synchronous standby nodes are faulty.

**Default value:** off

## keep\_sync\_window

**Parameter description:** Specifies the delay for entering the maximum availability mode.

- If **most\_available\_sync** is set to **on**, when synchronous standby nodes are faulty in primary/standby scenarios and the number of configured synchronous standby nodes is insufficient (for details, see the meaning of **synchronous\_standby\_name**), setting **keep\_sync\_window** will retain the maximum protection mode within the time window specified by **keep\_sync\_window**. That is, committing transactions on the primary node is blocked, delay the primary node to enter the maximum availability mode.
- If synchronous standby nodes recover from faults and the number of synchronous standby nodes meets the configuration requirements, transactions are not blocked.
- You are advised to set **keep\_sync\_window** to 5s. This prevents the monitoring system from incorrectly reporting network instability.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is s.

- The value **0** indicates that the **keep\_sync\_window** is not set, that is, the maximum availability mode is entered directly.
- Other values indicate the size of the timeout window.

**Default value:** 0

### NOTE

Setting this parameter may affect the RPO. If the primary node is faulty within the configured timeout window, the data generated from the time when the primary node is blocked to the time when the primary node is faulty may be lost.

## enable\_stream\_replication

**Parameter description:** Specifies whether data and logs are synchronized between the primary and standby nodes.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

- This parameter is used for performance testing in scenarios where data synchronization to standby nodes is enabled and where it is disabled. After this parameter is set to **off**, tests for abnormal scenarios such as switchovers and fault rectifications cannot be performed.
- This parameter is a restricted parameter, and you are advised not to set it to **off** in normal service scenarios.

---

**Value range:** Boolean

- **on** indicates that data and log synchronization between primary and standby nodes is enabled.

- **off** indicates that data and log synchronization between the primary and standby nodes is disabled.

**Default value:** on

## enable\_mix\_replication

**Parameter description:** Specifies how WAL files and data are replicated between primary and standby nodes.

This parameter is an INTERNAL parameter. Its default value is **off** and cannot be modified.

---

### NOTICE

- This parameter cannot be modified in normal service scenarios. That is, mixed replication of the WAL files and data pages is disabled.
- 

**Value range:** Boolean

- **on** indicates that the WAL file and data page mixed replication mode is enabled.
- **off** indicates that the WAL file and data page mixed replication mode is disabled.

**Default value:** off

## vacuum\_defer\_cleanup\_age

**Parameter description:** Specifies the number of transactions by which **VACUUM** will defer the cleanup of invalid row-store table records, so that **VACUUM** and **VACUUM FULL** do not clean up deleted tuples immediately.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000. **0** means no delay.

**Default value:** 0

## data\_replicate\_buffer\_size

**Parameter description:** Specifies the amount of memory used by queues when the sender sends data pages to the receiver. The value of this parameter affects the buffer size used during the replication from the primary node to the standby node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 4096 to 1072693248. The unit is KB.

**Default value:** 16 MB (that is, 16384 KB)

## walsender\_max\_send\_size

**Parameter description:** Specifies the size of the WAL or Sender buffers on the primary node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 8 to *INT\_MAX*. The unit is KB.

**Default value:** 8192 KB (8 MB)

## enable\_data\_replicate

**Parameter description:** Specifies how data is synchronized between primary and standby nodes when the data is imported to a row-store table.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the primary and standby nodes synchronize data using data pages when the data is imported to a row-store table. When **replication\_type** is set to **1**, this parameter cannot be set to **on**. If this parameter is set to **on** using the GUC tool, its value will be forcibly changed to **off**.
- **off** indicates that the primary and standby servers synchronize data using Xlogs when the data is imported to a row-store table.

**Default value:** off

## ha\_module\_debug

**Parameter description:** Specifies the replication status log of a specific data block during data replication.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the status of each data block is recorded in logs during data replication.
- **off** indicates that the status of each data block is not recorded in logs during data replication.

**Default value:** off

## catchup2normal\_wait\_time

**Parameter description:** Specifies the maximum duration for the standby node to catch up with the primary node when **most\_available\_sync** is enabled in primary/standby scenarios. The value of this parameter is an estimate and may be different from the actual value.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to 10000. The unit is ms.

- The value **-1** indicates that the primary node is blocked until the data catchup on the standby node is complete.

- The value **0** indicates that the primary node is not blocked during the data catchup on the standby node.
- Other values indicate the maximum duration that the primary node is blocked during the data catchup on the standby node. For example, if this parameter is set to **5000**, the primary node is blocked until the data catchup on the standby node is complete in 5s.

**Default value:** -1

## check\_sync\_standby

**Parameter description:** Specifies whether to enable the standby node check function. After the **synchronous\_standby\_names** parameter is correctly configured in the primary/standby scenario, if the synchronous standby node is faulty, the write service on the primary node reports a write failure. This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** on or off

- **on** indicates that the standby node check is enabled.
- **off** indicates that the standby node check is disabled.

**Default value:** off

### NOTE

- This parameter cannot be synchronized in job work and autonomous transactions. Otherwise, the check may not take effect.
- If the standby node check is not configured for a specified user or session and the standby node is faulty when the forcible synchronization commit mode is enabled, the write operation on a table causes the query on the same table by another user or in another session to hang. In this case, you need to recover the standby node or manually terminate the hung client.
- The standby node check function cannot be enabled in scenarios (such as VACUUM ANALYZE) where non-write operations trigger log writing. If the standby node does not meet the requirements for synchronizing configurations to the standby node, services will be hung in this scenario. In this case, you need to manually terminate the services.

## sync\_config\_strategy

**Parameter description:** Specifies the policy for synchronizing configuration files between the primary node and standby node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **all\_node:** If this parameter is set to **all\_node** for the primary node, the primary node is allowed to proactively synchronize configuration files to all standby nodes. If this parameter is set to **all\_node** for a standby node, the standby node is allowed to send synchronization requests to its primary node.
- **only\_sync\_node:** If this parameter is set to **only\_sync\_node** for the primary node, the primary node is only allowed to proactively synchronize configuration files to all standby nodes. If this parameter is set to **only\_sync\_node** for a standby node, the standby node is allowed to send synchronization requests to its primary node.

- **none\_node**: If this parameter is set to **none\_node** for the primary node, the primary node is not allowed to proactively synchronize configuration files to all standby nodes. If this parameter is set to **none\_node** for a standby node, the standby node is not allowed to send synchronization requests to its primary node.

**Default value:** all\_node

## hadr\_recovery\_time\_target

**Parameter description:** Specifies whether the standby database instance completes log writing and replay in streaming DR mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600. The unit is s.

0 indicates that log flow control is disabled. A value from 1 to 3600 indicates that a standby node can write and replay logs within the period specified by the value of **hadr\_recovery\_time\_target**. This ensures that the logs can be written and replayed within the period specified by the value of **hadr\_recovery\_time\_target** and the standby database instance can be promoted to primary quickly. If this parameter is set to a small value, the performance of the primary node is affected. If it is set to a large value, the log flow is not effectively controlled.

**Default value:** 0

## hadr\_recovery\_point\_target

**Parameter description:** Specifies the RPO time allowed for the standby database instance to flush logs to disks in streaming DR mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600. The unit is s.

0 indicates that log flow control is disabled. A value from 1 to 3600 indicates that the standby node can flush logs to disks within the period specified by the value of **hadr\_recovery\_point\_target**. This ensures that the log difference between the primary and standby database instances is controlled within the period specified by the value of **hadr\_recovery\_point\_target** during the switchover and the standby database instance can be promoted to primary. If this parameter is set to a small value, the performance of the primary node is affected. If it is set to a large value, the log flow is not effectively controlled.

**Default value:** 0

## hadr\_super\_user\_record\_path

**Parameter description:** Specifies the path for storing encrypted files of the **hadr\_disaster** user in the standby database instance in streaming DR mode. This is a SIGHUP parameter.

Set it based on instructions in [Table 17-1](#).



**Modification suggestion:** The value is automatically set by the streaming DR password transfer tool and does not need to be manually added.

**Value range:** a string

**Default value:** NULL

---

**NOTICE**

- In a database instance that contains a primary node and a standby node, the primary node is a sender relative to the standby node and the standby node is a receiver relative to the primary node.
  - The sender actively synchronizes the configuration file to the receiver, and the receiver requests the sender to synchronize the configuration file, which are two independent events, so that the configuration files are synchronized. If you do not want to synchronize configuration files, set this parameter to **none\_node** on the receiver. If the sender is a standby node, set this parameter to **none\_node** only. If the sender is a primary node, set this parameter to **none\_node** when the primary node does not synchronize with any standby node; or set this parameter to **only\_sync\_node** when the primary node synchronizes with synchronous standby nodes only and does not synchronize with asynchronous standby nodes.
  - To be specific, the sender sends a configuration file which directly overwrites the corresponding parameter in the configuration file of the receiver. After the policy for synchronizing configuration files is set, even if you modify configuration parameters of the receiver, the modification does not take effect because the sender immediately overwrites these parameters.
  - The following configuration parameters are not synchronized even if the policy for synchronizing configuration files is set: **application\_name**, **audit\_directory**, **available\_zone**, **comm\_control\_port**, **comm\_sctp\_port**, **listen\_addresses**, **log\_directory**, **port**, **replconninfo1**, **replconninfo2**, **replconninfo3**, **replconninfo4**, **replconninfo5**, **replconninfo6**, **replconninfo7**, **replconninfo8**, **replconninfo9**, **replconninfo10**, **replconninfo11**, **replconninfo12**, **replconninfo13**, **replconninfo14**, **replconninfo15**, **replconninfo16**, **replconninfo17**, **replconninfo18**, **ssl**, **ssl\_ca\_file**, **ssl\_cert\_file**, **ssl\_ciphers**, **ssl\_crl\_file**, **ssl\_key\_file**, **ssl\_renegotiation\_limit**, **ssl\_cert\_notify\_time**, **synchronous\_standby\_names**, **local\_bind\_address**, **perf\_directory**, **query\_log\_directory**, **asp\_log\_directory**, **streaming\_router\_port**, **enable\_upsert\_to\_merge**, **recovery\_min\_apply\_delay**, and **sync\_config\_strategy**.
- 

### 17.3.7.3 Standby Server

#### hot\_standby

**Parameter description:** Specifies whether to allow connections and queries on a standby node during its recovery.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**NOTICE**

- If this parameter is set to **on**, **wal\_level** must be set to **hot\_standby** or higher. Otherwise, the database startup fails.
- In an HA system, **hot\_standby** cannot be set to **off**, because this setting can affect other features of the HA system.
- If the **hot\_standby** parameter was disabled and the **wal\_level** parameter was set to a value lower than **hot\_standby**, perform the following operations to ensure that the logs to be replayed on the standby node can be queried on the standby node before enabling the **hot\_standby** parameter again:
  1. Change the **wal\_level** value of the primary and standby nodes to **hot\_standby** or higher, and restart the instances for the change to take effect.
  2. Perform the checkpoint operation on the primary node and query the **pg\_stat\_get\_wal\_senders()** function to ensure that the **receiver\_replay\_location** value of each standby node is the same as the **sender\_flush\_location** value of the primary node. Ensure that the value adjustment of **wal\_level** is synchronized to standby nodes and takes effect, and standby nodes do not need to replay low-level logs.
  3. Set the **hot\_standby** parameter of the primary and standby nodes to **on**, and restart the instances for the setting to take effect.

**Value range:** Boolean

- **on** indicates that connections and queries are allowed on the standby node during the recovery.
- **off** indicates that connections and queries are not allowed on the standby node during the recovery.

**Default value:** on

## max\_standby\_archive\_delay

**Parameter description:** Specifies the wait period before queries on a standby node are canceled when the queries conflict with WAL processing and archiving in hot standby mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**NOTICE**

-1 indicates that the standby node waits until the conflicting queries are complete.

**Value range:** an integer ranging from -1 to *INT\_MAX*. The unit is ms.

**Default value:** 3s (3000 ms)

## max\_standby\_streaming\_delay

**Parameter description:** Specifies the wait period before queries on a standby node are canceled when the queries conflict with WAL data receiving through streaming replication in hot standby mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#). If this parameter is set to a large value or the service load is heavy, an error may be reported for waiting for transaction replay and flushing to disks.

---

### NOTICE

-1 indicates that the standby node waits until the conflicting queries are complete.

---

**Value range:** an integer ranging from -1 to *INT\_MAX*. The unit is ms.

**Default value:** 3s (3000 ms)

## wal\_receiver\_status\_interval

**Parameter description:** Specifies the maximum interval for notifying the primary node of the WAL Receiver status.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is ms.

**Default value:** 5s (5000 ms)

---

### NOTICE

If this parameter is set to 0, the standby node does not send information, such as the log receiving location, to the primary node. As a result, the transaction commit on the primary node may be blocked, and the switchover may fail. In normal service scenarios, you are not advised to set this parameter to 0.

---

## hot\_standby\_feedback

**Parameter description:** Specifies whether a standby node is allowed to send the result of a query performed on it to the primary node, preventing a query conflict.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the ID of the oldest transaction active on standby nodes will be sent to the primary node.
- **off** indicates that the ID of the oldest transaction active on standby nodes will not be sent to the primary node.

**Default value:** off

#### NOTICE

If this parameter is set to **on**, VACUUM on the primary node will not clean up tuples modified in transactions later than the oldest transaction active on standby nodes.

Therefore, the performance of the primary node will be affected. If replay conflicts with query on the standby node and a query error is reported, you are advised to increase the value of **max\_standby\_streaming\_delay**.

### wal\_receiver\_timeout

**Parameter description:** Specifies the maximum wait period for a standby node to receive data from the primary node.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is ms.

**Default value:** 6s (6000 ms)

### wal\_receiver\_connect\_timeout

**Parameter description:** Specifies the timeout period for a standby node to connect to the primary node.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*/1000. The unit is s.

**Default value:** 2s

### wal\_receiver\_connect\_retries

**Parameter description:** Specifies the maximum attempts that a standby node connects to the primary node

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to *INT\_MAX*

**Default value:** 1

### wal\_receiver\_buffer\_size

**Parameter description:** Specifies the memory buffer size for the standby nodes to store the received Xlog files.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 4096 to 1047552. The unit is KB.

**Default value:** 64 MB (65536 KB)

## primary\_slotname

**Parameter description:** Specifies the slot name of the primary node corresponding to a standby node. This parameter is used for the mechanisms to verify the primary-standby relationship and delete WALs.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** empty

## max\_logical\_replication\_workers

**Parameter description:** Specifies the maximum number of apply worker threads on the subscriber side.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 262143

**Default value:** 4

## 17.3.8 Query Planning

This section describes the method configuration, cost constants, planning algorithm, and some configuration parameters for the optimizer.

### NOTE

Two parameters are involved in the optimizer:

- *INT\_MAX* indicates the maximum value of the INT data type. The value is **2147483647**.
- *DBL\_MAX* indicates the maximum value of the FLOAT data type.

In addition to customer services, global query planning parameters also affect database O&M and monitoring services, such as WDR generation, scale-out, redistribution, and data import and export.

### 17.3.8.1 Optimizer Method Configuration

These configuration parameters provide a crude method of influencing the query plans chosen by the query optimizer. If the default plan chosen by the optimizer for a particular query is not optimal, a temporary solution is to use one of these configuration parameters to force the optimizer to choose a different plan. Better ways include adjusting the optimizer cost constants, manually running **ANALYZE**, increasing the value of the **default\_statistics\_target** parameter, or increasing the amount of the statistics collected in specific columns using **ALTER TABLE SET STATISTICS**.

## enable\_broadcast

**Parameter description:** Controls whether the query optimizer uses the broadcast distribution method when it evaluates the cost of stream.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

 **NOTE**

This parameter does not take effect in the current version.

## enable\_bitmapscan

**Parameter description:** Specifies the query optimizer's use of bitmap-scan plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

## force\_bitmapand

**Parameter description:** Specifies the query optimizer's use of BitmapAnd plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** off

## enable\_hashagg

**Parameter description:** Specifies the query optimizer's use of Hash aggregation plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

## enable\_hashjoin

**Parameter description:** Specifies the query optimizer's use of hash-join plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** on

## enable\_indexscan

**Parameter description:** Specifies the query optimizer's use of index-scan plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** on

## enable\_indexonlyscan

**Parameter description:** Specifies the query optimizer's use of index-only-scan plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** on

## enable\_material

**Parameter description:** Specifies the query optimizer's use of materialization.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** on

## enable\_mergejoin

**Parameter description:** Specifies the query optimizer's use of merge-join plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.

- **off**: disabled.

**Default value:** off

## enable\_nestloop

**Parameter description:** Specifies whether the query optimizer uses the nested-loop join plan type to fully scan internal tables. If this variable is disabled, the optimizer preferentially selects another method when another method exists.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** off

## enable\_index\_nestloop

**Parameter description:** Specifies whether the query optimizer uses the nested-loop join plan type to scan the parameterized indexes of internal tables.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

## enable\_seqscan

**Parameter description:** Specifies the query optimizer's use of sequential scan plan types. If this variable is disabled, the optimizer preferentially selects another method when another method exists.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

## enable\_sort

**Parameter description:** Specifies the query optimizer's choice of sort methods. If this variable is disabled, the optimizer preferentially selects another method when another method exists.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean



- **on**: enabled.
- **off**: disabled.

**Default value:** on

## enable\_tidscan

**Parameter description:** Specifies the query optimizer's use of Tuple ID (TID) scan plan types.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** on

## enable\_kill\_query

**Parameter description:** In CASCADE mode, when a user is deleted, all the objects belonging to the user are deleted. This parameter specifies whether the queries of the objects belonging to the user can be unlocked when the user is deleted.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the unlocking is allowed.
- **off** indicates that the unlocking is not allowed.

**Default value:** off

## max\_recursive\_times

**Parameter description:** Specifies the maximum number of **WITH RECURSIVE** iterations.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

**Default value:** 200

## enable\_change\_hjcost

**Parameter description:** Specifies whether the optimizer excludes internal table running costs when selecting the Hash Join cost path. If it is set to **on**, tables with a few records and high running costs are more possible to be selected.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: enabled.
- **off**: disabled.

**Default value:** off

## **enable\_absolute\_tablespace**

**Parameter description:** Specifies whether the tablespace can use the absolute path.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that an absolute path can be used.
- **off** indicates that an absolute path cannot be used.

**Default value:** on

## **enable\_valuepartition\_pruning**

**Parameter description:** Specifies whether the DFS partitioned table is dynamically or statically optimized.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the DFS partitioned table is dynamically or statically optimized.
- **off** indicates that the DFS partitioned table is not dynamically or statically optimized.

**Default value:** on

## **qrw\_inlist2join\_optmode**

**Parameter description:** Specifies whether to enable inlist-to-join (inlist2join) query rewriting.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- **disable** indicates that the inlist2join query rewriting is disabled.
- **cost\_base** indicates that the cost-based inlist2join query rewriting is enabled.
- **rule\_base** indicates that the forcible rule-based inlist2join query rewriting is enabled.
- A positive integer indicates the threshold of inlist2join query rewriting. If the number of list elements in the IN clause is greater than the threshold, the rewriting is performed.

**Default value:** cost\_base

## **skew\_option**

**Parameter description:** Specifies whether an optimization policy is used.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- **off** indicates that the policy is disabled.
- **normal** indicates that a radical policy is used. All possible skews are optimized.
- **lazy** indicates that a conservative policy is used. Uncertain skews are ignored.

**Default value:** normal

## default\_limit\_rows

**Parameter description:** Specifies the default estimated number of limit rows for generating genericplan. If this parameter is set to a positive value, the positive value is used as the estimated number of limit rows. If the positive value is a decimal, the value is rounded up automatically. If this parameter is set to a negative value, the negative value is converted to a percentage and used as default estimated value, that is, -5 indicates 5%.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating-point number ranging from -100 to *DBL\_MAX*

**Default value:** -10

## check\_implicit\_conversions

**Parameter description:** Specifies whether to check that candidate index paths are generated for index columns with implicit type conversion in a query.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the system checks whether candidate index paths are generated for index columns with implicit type conversion in a query.
- **off** indicates that a check will not be performed.

**Default value:** off

## cost\_weight\_index

**Parameter description:** Specifies the cost weight of **index\_scan**.

- If this parameter is set to **1**, no adjustment is performed.
- If this parameter is set to a value less than 1, the cost of **index\_scan** is reduced and **index\_scan** is more likely to be selected by the optimizer.
- If this parameter is set to a value greater than 1, the cost of **index\_scan** increases.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating-point number ranging from 1e-10 to 1e+10.

**Default value:** 1

### 17.3.8.2 Optimizer Cost Constants

This section describes the optimizer cost constants. The cost variables described here are measured on an arbitrary scale. Only their relative values matter, therefore scaling them all up or down by the same factor will result in no change in the optimizer's choices. By default, these cost variables are based on the cost of sequential page fetches, that is, **seq\_page\_cost** is conventionally set to **1.0** and the other cost variables are set with reference to the parameter. However, you can use a different scale, such as actual execution time in milliseconds.

#### seq\_page\_cost

**Parameter description:** Specifies the optimizer's estimated cost of a disk page fetch that is part of a series of sequential fetches.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 1

#### random\_page\_cost

**Parameter description:** Specifies the optimizer's estimated cost of an out-of-sequence disk page fetch.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

Although the server allows you to set **random\_page\_cost** to a value less than that of **seq\_page\_cost**, it is not physically sensitive to do so. However, setting them equal makes sense if the database is entirely cached in RAM, because in that case there is no penalty for fetching pages out of sequence. Also, in a heavily-cached database you should lower both values relative to the CPU parameters, since the cost of fetching a page already in RAM is much smaller than it would normally be.

---

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 4

#### NOTE

- This value can be overwritten for tables and indexes in a particular tablespace by setting the tablespace parameter of the same name.
- Reducing this value relative to **seq\_page\_cost** will cause the system to prefer index scans and raising it will make index scans relatively more expensive. You can increase or decrease both values together to change the disk I/O costs relative to CPU costs.

#### cpu\_tuple\_cost

**Parameter description:** Specifies the optimizer's estimated cost of processing each row during a query.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 0.01

## cpu\_index\_tuple\_cost

**Parameter description:** Specifies the optimizer's estimated cost of processing each index entry during an index scan.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 0.005

## cpu\_operator\_cost

**Parameter description:** Specifies the optimizer's estimated cost of processing each operator or function executed during a query.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 0.0025

## effective\_cache\_size

**Parameter description:** Specifies the optimizer's assumption about the effective size of the disk cache that is available to a single query.

When setting this parameter you should consider both GaussDB's shared buffers and the kernel's disk cache. Also, take into account the expected number of concurrent queries on different tables, since they will have to share the available space.

This parameter has no effect on the size of shared memory allocated by GaussDB. It is used only for estimation purposes and does not reserve kernel disk cache. The value is in the unit of disk page. Usually the size of each page is 8192 bytes.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647. The unit is 8 KB.

A value greater than the default one makes index scans more likely to be used, whereas a value less than the default one makes sequential scans more likely to be used.

**Default value:**

**180 GB** (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, 96-core CPU/1024-GB memory); **135 GB** (96-core CPU/768-GB memory); **90 GB** (64-core CPU/512-GB memory); **80 GB** (60-core CPU/480-GB memory); **40 GB** (32-core CPU/256-GB memory); **18 GB** (16-core CPU/128-GB memory); **8 GB** (8-core CPU/64-GB memory); )

## allocate\_mem\_cost

**Parameter description:** Specifies the query optimizer's estimated cost of creating a hash table for memory space using hash join. This parameter is used for optimization when the hash join estimation is inaccurate.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to *DBL\_MAX*

**Default value:** 0

### 17.3.8.3 Genetic Query Optimizer

This section describes parameters related to genetic query optimizer. The genetic query optimizer (GEQO) is an algorithm that plans queries by using heuristic searching. This algorithm reduces planning time for complex queries and the costs of producing plans are sometimes inferior to those found by the normal exhaustive-search algorithm.

## geqo

**Parameter description:** Specifies whether to enable the genetic query optimization.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

It is best not to turn it off in execution. **geqo\_threshold** provides more subtle control of GEQO.

---

**Value range:** Boolean

- **on** indicates that the genetic query optimization is enabled.
- **off** indicates that the genetic query optimization is disabled.

**Default value:** on

## geqo\_threshold

**Parameter description:** Specifies the number of **FROM** items. Genetic query optimization is used to plan queries when the number of statements executed is greater than this value.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

- For simpler queries, it is best to use the regular, exhaustive-search planner; but for queries with many tables, it is better to use GEQO to manage the queries.
  - A **FULL OUTER JOIN** construct counts as only one **FROM** item.
-

**Value range:** an integer ranging from 2 to *INT\_MAX*

**Default value:** 12

## geqo\_effort

**Parameter description:** Controls the trade-off between planning time and query plan quality in GEQO.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

**geqo\_effort** does not do anything directly. This parameter is only used to compute the default values for the other variables that influence GEQO behavior. If you prefer, you can manually set the other parameters instead.

---

**Value range:** an integer ranging from 1 to 10

---

### NOTICE

Larger values increase the time spent in query planning, but also increase the probability that an efficient query plan is chosen.

---

**Default value:** 5

## geqo\_pool\_size

**Parameter description:** Controls the pool size used by GEQO, that is, the number of individuals in the genetic population.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

---

### NOTICE

The value of this parameter must be at least **2**, and useful values are typically from **100** to **1000**. If this parameter is set to **0**, GaussDB selects a proper value based on **geqo\_effort** and the number of tables.

---

**Default value:** 0

## geqo\_generations

**Parameter description:** Specifies the number of iterations of the GEQO.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

---

**NOTICE**

The value of this parameter must be at least **1**, and useful values are typically from **100** to **1000**. If it is set to **0**, a suitable value is chosen based on **geqo\_pool\_size**.

---

**Default value:** 0

### geqo\_selection\_bias

**Parameter description:** Specifies the selection bias used by GEQO. The selection bias is the selective pressure within the population.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 1.5 to 2.0

**Default value:** 2

### geqo\_seed

**Parameter description:** Specifies the initial value of the random number generator used by GEQO to select random paths through the join order search space.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0.0 to 1.0

---

**NOTICE**

Varying the value changes the set of join paths explored, and may result in a better or worse best path being found.

---

**Default value:** 0

## 17.3.8.4 Other Optimizer Options

### explain\_dna\_file

**Parameter description:** Sets [explain\\_perf\\_mode](#) to **run** to export object files in CSV format.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

The value of this parameter must be an absolute path plus a file name with the extension **.csv**.

---

**Value range:** a string



**Default value:** empty

## **explain\_perf\_mode**

**Parameter description:** Specifies the display format of the **explain** command.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** **normal**, **pretty**, **summary**, and **run**

- **normal** indicates that the default printing format is used.
- **pretty** indicates a new format improved by using GaussDB. The new format contains a plan node ID, directly and effectively analyzing performance.
- **summary** indicates that the analysis result on this information is printed in addition to the printed information specified by **pretty**.
- **run** indicates that the system exports the printed information specified by **summary** as a CSV file for further analysis.

**Default value:** **pretty** (In the current version, only the value **normal** takes effect. If the value is not **normal**, **normal** is still displayed.)

 **NOTE**

In pretty mode, only one plan is displayed. After you run **CREATE RULE** to create a rule, the SQL statement executed may generate multiple plans. In this case, you are advised to use the normal mode.

The following is an example:

```
openGauss=# CREATE TABLE another_table (id int, name text);
CREATE TABLE
openGauss=# CREATE TABLE my_table (id int, name text);
CREATE TABLE
openGauss=# CREATE RULE my_rule AS ON INSERT TO my_table
openGauss=# WHERE NEW.id > 5
openGauss=# DO INSTEAD (INSERT INTO another_table VALUES (NEW.id, 'Some Data'));
CREATE RULE
openGauss=# explain INSERT INTO my_table VALUES (5, 'Test Name');
QUERY PLAN

Insert on my_table (cost=0.00..0.01 rows=1 width=0)
-> Result (cost=0.00..0.01 rows=1 width=0)

Insert on another_table (cost=0.00..0.01 rows=1 width=0)
-> Result (cost=0.00..0.01 rows=1 width=0)
 One-Time Filter: false
(6 rows)

openGauss=# set explain_perf_mode=pretty;
SET
openGauss=# explain INSERT INTO my_table VALUES (5, 'Test Name');
id | operation | E-rows | E-width | E-costs
-----+-----+-----+-----+-----
1 | -> Insert on another_table | 1 | 0 | 0.000..0.010
2 | -> Result | 1 | 0 | 0.000..0.010
(2 rows)

Predicate Information (identified by plan id)

2 --Result
 One-Time Filter: false
(2 rows)

openGauss=# drop table my_table;
DROP TABLE
openGauss=# drop table another_table;
DROP TABLE
```

## analysis\_options

**Parameter description:** Specifies whether to enable function options in the corresponding options to use the corresponding location functions, including data verification and performance statistics. For details, see the options in the value range.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- **HASH\_CONFLICT** indicates that the log in the **pg\_log** directory of the database node process displays the statistics of the hash table, including the hash table size, hash link length, and hash conflict.
- **STREAM\_DATA\_CHECK** indicates that a CRC check is performed on data before and after network data transmission.

**Default value:** **ALL,on(),off(HASH\_CONFLICT,STREAM\_DATA\_CHECK)**, which indicates that no location function is enabled.

## cost\_param

**Parameter description:** Specifies use of different estimation methods in specific customer scenarios, allowing estimated values approximating to onsite values. This parameter can control various methods simultaneously by performing AND (&) on the bit of each method. A method is selected if the result value is not 0.

When **cost\_param & 1** is set to a value other than 0, an improved mechanism is used for connecting the selection rate of non-equi-joins. This method is more accurate for estimating the selection rate of joins between two identical tables. At present, **cost\_param & 1=0** is not used. That is, a better formula is selected for calculation.

When **cost\_param & 2** is set to a value other than 0, the selection rate is estimated based on multiple filter criteria. The lowest selection rate among all filter criteria, but not the product of the selection rates for two tables under a specific filter criterion, is used as the total selection rate. This method is more accurate when a close correlation exists between the columns to be filtered.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

**Default value:** 0

## enable\_partitionwise

**Parameter description:** Specifies whether to select an intelligent algorithm for joining partition tables.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that an intelligent algorithm is selected.
- **off** indicates that an intelligent algorithm is not selected.

**Default value:** off

## rewrite\_rule

**Parameter description:** Specifies the optional query rewriting rules that have been enabled. Some query rewrite rules are optional. Enabling them cannot always improve the query efficiency. In a specific customer scenario, you can set the query rewriting rules through this GUC parameter to achieve optimal query efficiency.

This parameter can control the combination of query rewriting rules, for example, there are over one override rules: rule1, rule2, rule3, and rule4. You can perform the following settings:

```
set rewrite_rule=rule1; -- Enable query rewriting rule rule1
set rewrite_rule=rule2, rule3; -- Enable the query rewriting rules rule2 and rule3
set rewrite_rule=none; -- Disable all optional query rewriting rules
```

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- **none**: Does not use any optional query rewriting rules.
- **lazyagg**: Uses the Lazy Agg query rewriting rules for eliminating aggregation operations in subqueries.
- **magicset**: The Magic Set query rewriting rules are used to associate subqueries which have aggregation operators with the main query in advance to reduce repeated scanning of sublinks.
- **uniquecheck**: Uses the Unique Check query rewriting rules. Optimize the subquery statements in target columns without agg and check whether the number of returned rows is 1.
- **intargetlist**: Uses the In Target List query rewriting rules (subquery optimization in the target column).
- **predpushnormal**: Uses the Predicate Push query rewriting rule (push the predicate condition to the subquery).
- **predpushforce**: Uses the Predicate Push query rewriting rules. Push down predicate conditions to subqueries and use indexes as much as possible for acceleration.
- **predpush**: Selects the optimal plan based on the cost in **predpushnormal** and **predpushforce**.
- **disable\_pullup\_expr\_sublink**: The optimizer is not allowed to pull up sublinks of the expr\_sublink type. For details about sublink classification and pull-up principles, see "SQL Optimization" > "Typical SQL Optimization Methods" > "Optimizing Subqueries" in *Developer Guide*.
- **disable\_pullup\_not\_in\_sublink**: The optimizer is not allowed to pull up sublinks related to NOT IN. For details about sublink classification and pull-up principles, see "SQL Optimization" > "Typical SQL Optimization Methods" > "Optimizing Subqueries" in *Developer Guide*.
- **enable\_sublink\_pullup\_rownum**: The optimizer can be used to promote subjoins when the SQL statement contains the ROWNUM pseudocolumn.

**Default value:** magicset

 **NOTE**

In the current version, the **partialpush** and **disablerep** parameters can be set but do not take effect.

## enable\_pbe\_optimization

**Parameter description:** Specifies whether the optimizer optimizes the query plan for statements executed in Parse Bind Execute (PBE) mode.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the optimizer optimizes the query plan of the PBE statement.
- **off** indicates that the optimization is not used.

**Default value:** on

## enable\_global\_stats

**Parameter description:** Specifies the current statistics collection mode, which can be global statistics collection or single-node statistics collection. By default, the global statistics collection mode is used. If this parameter is set to **off**, the statistics of the first node in the database are collected by default. In this case, the quality of the generated query plan may be affected. However, the information collection performance is optimal. Therefore, exercise caution when setting this parameter.

This is a **SUSET** parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** or **true** indicates the global statistics mode.
- **off** or **false** indicates the database node statistics.

**Default value:** on

### NOTE

This parameter does not take effect in centralized mode.

## enable\_opfusion

**Parameter description:** Specifies whether to optimize simple addition, deletion, modification, and query operations.

This is a **USERSET** parameter. Set it based on instructions in [Table 17-1](#).

The restrictions on simple query are as follows:

- Only indexscan and indexonlyscan are supported, and the filter criteria of all WHERE statements are on indexes.
- Only single tables can be added, deleted, modified, and queried. Join and using are not supported.
- Only row-store tables are supported. Partitioned tables and tables with triggers are not supported.
- Information statistics features such as active sql and QPS are not supported.
- Tables that are being scaled out or in are not supported.
- The system column cannot be queried or modified.

- Only simple **SELECT** statements are supported. For example:

```
SELECT c3 FROM t1 WHERE c1 = ? and c2 =10;
```

Only columns in the target table can be queried. Columns **c1** and **c2** are index columns, which can be followed by constants or parameters. You can use **for update**.

- Only simple **INSERT** statements are supported. For example:

```
INSERT INTO t1 VALUES (?,10,?);
```

Only one **VALUES** is supported. The type in **VALUES** can be a constant or a parameter. **RETURNING** is not supported.

- Only simple **DELETE** statements are supported. For example:

```
DELETE FROM t1 WHERE c1 = ? and c2 = 10;
```

Columns **c1** and **c2** are index columns, which can be followed by constants or parameters.

- Only simple **UPDATE** statements are supported. For example:

```
UPDATE t1 SET c3 = c3+? WHERE c1 = ? and c2 = 10;
```

The values modified in column **c3** can be constants, parameters, or a simple expression. Columns **c1** and **c2** are index columns, which can be followed by constants or parameters.

**Value range:** Boolean

- on:** enabled.
- off:** disabled.

**Default value:** on

## sql\_beta\_feature

**Parameter description:** Specifies the SQL engine's optional beta features to be enabled, including optimization of row count estimation and query equivalence estimation.

These optional features provide optimization for specific scenarios, but performance deterioration may occur in some scenarios for which testing is not performed. In a specific customer scenario, you can set the query rewriting rules through this GUC parameter to achieve optimal query efficiency.

This parameter determines the combination of the SQL engine's beta features, for example, feature1, feature2, feature3, and feature4. You can perform the following settings:

```
-- Enable beta feature feature1 of the SQL engine.
set sql_beta_feature=feature1;
-- Enable beta features feature2 and feature3 of the SQL engine.
set sql_beta_feature=feature2,feature3;
-- Disable all optional beta features of the SQL engine.
set sql_beta_feature=none;
```

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- none:** None of the beta optimizer features is used.
- sel\_semi\_poisson:** Uses poisson distribution to calibrate the equivalent semi-join and anti-join selection rates.
- sel\_expr\_instr:** Uses the matching row count to provide more accurate estimation for **instr(col, 'const') > 0, = 0, = 1**.
- param\_path\_gen:** Generates more possible parameterized paths.
- rand\_cost\_opt:** Optimizes the random read cost of tables that have a small amount of data.
- param\_path\_opt:** Uses the bloat ratio of the table to optimize the analysis information of indexes.
- page\_est\_opt:** optimizes the **relpages** estimation for the analysis information of table indexes.
- no\_unique\_index\_first:** Disables optimization of the primary key index scanning path first.
- join\_sel\_with\_cast\_func:** Type conversion functions are supported when the number of join rows is estimated.

- **canonical\_pathkey**: The regular path key is generated in advance. (**pathkey**: a set of ordered key values of data.)
- **index\_cost\_with\_leaf\_pages\_only**: Considers index leaf nodes when the index cost is estimated.
- **a\_style\_coerce**: enables the Decode type conversion rule to be compatible with ORA. For details, see case processing in O-compatible mode in "SQL Reference" > "Type Conversion" > "UNION, CASE, and Related Constructs" in *Developer Guide*.
- **predpush\_same\_level**: Enables the **predpush** hint to control parameterized paths at the same layer.
- **enable\_plsql\_smp**: Enables parallel execution of queries in stored procedures. Currently, only one query can be executed in parallel at a time, and no parallel execution plan is generated for cursor-related operations, autonomous transactions, and queries in exceptions.
- **disable\_bitmap\_cost\_with\_lossy\_pages**: Disables the computation of the cost of lossy pages in the bitmap path cost.

**Default value:**

"sel\_semi\_poisson,sel\_expr\_instr,rand\_cost\_opt,param\_path\_opt,page\_est\_opt"

## default\_statistics\_target

**Parameter description:** Specifies the default statistics target for table columns without a column-specific target set via **ALTER TABLE SET STATISTICS**. The number of rows sampled during statistics collection is affected.

If this parameter is set to a positive number, the number of rows sampled in the statistics histogram is **default\_statistics\_target** x 300. If the parameter is set to a negative number, it indicates the percentage of statistics collected. The negative number converts to its corresponding percentage, for example, **-5** means 5%. The number of sampled rows is the total number of rows multiplied by 5%. This parameter affects only the target number of sampled rows in the statistics. The actual number of sampled rows is also affected by the memory parameter [work\\_mem](#).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -100 to 10000

---

**NOTICE**

- A larger positive number than the default value increases the time required to do **ANALYZE**, but might improve the quality of the optimizer's estimates.
  - Changing settings of this parameter may result in performance deterioration. If query performance deteriorates, you can:
    1. Restore to the default statistics.
    2. Use hints to force the optimizer to use the optimal query plan. For details, see "SQL Optimization" > "Hint-based Tuning" in *Developer Guide*.
  - If this parameter is set to a negative value, the number of samples is equal to or greater than 2% of the total data volume, and the number of records in user tables is less than 1.6 million, the time taken by running **ANALYZE** will be longer than when this parameter uses its default value.
  - If this parameter is set to a negative value, the auto-analyze function is disabled.
- 

**Default value:** 100

## constraint\_exclusion

**Parameter description:** Specifies the query optimizer's use of table constraints to optimize queries.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **on** indicates that constraints for all tables are examined.
- **off** indicates that constraints for any table are not examined.
- **partition** indicates that only constraints for inheritance child tables and **UNION ALL** subqueries are examined.

---

**NOTICE**

When **constraint\_exclusion** is set to **on**, the optimizer compares query conditions with the table's **CHECK** constraints, and omits scanning tables for which the conditions contradict the constraints.

---

**Default value:** partition

 **NOTE**

Currently, **constraint\_exclusion** is enabled by default only for cases that are often used to implement table partitioning. If this parameter is enabled for all tables, extra planning is imposed on simple queries, which has no benefits. If you have no partitioned tables, set it to **off**.

## cursor\_tuple\_fraction

**Parameter description:** Specifies the optimizer's estimated fraction of a cursor's rows that are retrieved.



This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating-point number ranging from 0.0 to 1.0

---

#### NOTICE

Smaller values of this setting bias the optimizer towards using **fast start** plans for cursors, which will retrieve the first few rows quickly while perhaps taking a long time to fetch all rows. Larger values put more emphasis on the total estimated time. At the maximum setting of **1.0**, cursors are planned exactly like regular queries, considering only the total estimated time and how soon the first rows might be delivered.

---

**Default value:** 0.1

## from\_collapse\_limit

**Parameter description:** Specifies whether the optimizer merges sub-queries into upper queries based on the resulting FROM list. The optimizer merges sub-queries into upper queries if the resulting FROM list would have no more than this many items.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to *INT\_MAX*

---

#### NOTICE

Smaller values reduce planning time but may lead to inferior execution plans.

---

**Default value:** 8

## join\_collapse\_limit

**Parameter description:** Specifies whether the optimizer rewrites **JOIN** constructs (except **FULL JOIN**) into lists of **FROM** items based on the number of the items in the result list.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to *INT\_MAX*

---

#### NOTICE

- Setting this parameter to **1** prevents join reordering. As a result, the join order specified in the query will be the actual order in which the relations are joined. The query optimizer does not always choose the optimal join order. Therefore, advanced users can temporarily set this variable to **1**, and then specify the join order they desire explicitly.
  - Smaller values reduce planning time but lead to inferior execution plans.
-

**Default value:** 8

## plan\_mode\_seed

**Parameter description:** This is a commissioning parameter. Currently, it supports only **OPTIMIZE\_PLAN** and **RANDOM\_PLAN**. The value **0** (for **OPTIMIZE\_PLAN**) indicates the optimized plan using the dynamic planning algorithm. Other values are for **RANDOM\_PLAN**, which indicates that the plan is randomly generated. **-1** indicates that users do not specify the value of the seed identifier. In this case, the optimizer generates a random integer from **1** to **2147483647** and a random execution plan based on the generated integer. A value from **1** to **2147483647** is regarded as the seed identifier, based on which the optimizer generates a random execution plan.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to 2147483647

**Default value:** 0

---

### NOTICE

- If this parameter is set to a random execution plan, the optimizer generates a random execution plan that may not be the optimal one. Therefore, to guarantee the query performance, the default value **0** is recommended during upgrade, scale-out, scale-in, and O&M.
  - If this parameter is not set to **0**, the specified plan hint will not be used.
- 

## hashagg\_table\_size

**Parameter description:** Specifies the hash table size during the execution of the HASH JOIN operation.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to  $INT\_MAX/2$

**Default value:** 0

## enable\_bloom\_filter

**Parameter description:** Specifies whether the BloomFilter optimization is used.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the BloomFilter optimization can be used.
- **off** indicates that the BloomFilter optimization cannot be used.

**Default value:** on

## enable\_extrapolation\_stats

**Parameter description:** Specifies whether the extrapolation logic is used for data of DATE type based on historical statistics. The logic can increase the accuracy of

estimation for tables whose statistics are not collected in time, but will possibly provide an overlarge estimation due to incorrect extrapolation. Enable the logic only in scenarios where the data of DATE type is periodically inserted.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the extrapolation logic is used for data of DATE type based on historical statistics.
- **off** indicates that the extrapolation logic is not used for data of DATE type based on historical statistics.

**Default value:** off

## autoanalyze

**Parameter description:** Specifies whether to automatically collect statistics on tables that have no statistics when a plan is generated. **autoanalyze** cannot be used for foreign or temporary tables. To collect statistics, manually perform the ANALYZE operation. If an exception occurs in the database during the execution of autoanalyze on a table, after the database is recovered, the system may still prompt you to collect the statistics of the table when you run the statement again. In this case, manually perform the ANALYZE operation on the table to synchronize statistics.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the table statistics are automatically collected.
- **off** indicates that the table statistics are not automatically collected.

**Default value:** off

### NOTE

This parameter is irrelevant to autoanalyze in the autovacuum thread. This parameter does not take effect in centralized mode.

## enable\_analyze\_check

**Parameter description:** Checks whether statistics were collected about tables whose **reltuples** and **relpages** are displayed as 0 in **pg\_class** during plan generation.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the tables will be checked.
- **off** indicates that the tables will not be checked.

**Default value:** off

## enable\_sonic\_hashagg

**Parameter description:** Specifies whether to use the hash aggregation operator designed for column-oriented hash tables when certain constraints are met.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the hash aggregation operator designed for column-oriented hash tables is used when certain constraints are met.
- **off** indicates that the hash aggregation operator designed for column-oriented hash tables is not used.

### NOTE

- If **enable\_sonic\_hashagg** is enabled and the Hash Agg operator designed based on the column-oriented hash table is used when the query meets the constraint condition, the memory usage of the Hash Agg operator can be reduced. However, in scenarios where the code generation technology can significantly improve performance, the performance of the operator may deteriorate.
- If **enable\_sonic\_hashagg** is enabled and the hash aggregation operator designed based on the column-oriented hash table is used when the query meets the constraint condition, the operator is displayed as Sonic Hash Aggregation in the execution plan and execution information of Explain Analyze/Performance; when the query does not meet the constraint condition, the operator is displayed as Hash Aggregation. For details, see "SQL Optimization" > "Introduction to the SQL Execution Plan" > "Description" in *Developer Guide*.

**Default value:** on

## enable\_sonic\_hashjoin

**Parameter description:** Specifies whether to use the hash join operator designed for column-oriented hash tables when certain constraints are met.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the hash join operator designed for column-oriented hash tables is used when certain constraints are met.
- **off** indicates that the hash join operator designed for column-oriented hash tables is not used.

### NOTE

- Currently, the parameter can be used only for Inner Join.
- If **enable\_sonic\_hashjoin** is enabled, the memory usage of query using the Hash Inner operator can be reduced. However, in scenarios where the code generation technology can significantly improve performance, the performance of the operator may deteriorate.
- If **enable\_sonic\_hashjoin** is enabled and the hash join operator designed based on the column-oriented hash table is used when the query meets the constraint condition, the operator is displayed as **Sonic Hash Join** in the execution plan and execution information of Explain Analyze/Performance; when the query does not meet the constraint condition, the operator is displayed as **Hash Join**. For details, see "SQL Optimization" > "Introduction to the SQL Execution Plan" > "Description" in *Developer Guide*.

**Default value:** on

## enable\_sonic\_optspill

**Parameter description:** Specifies whether to optimize the number of files to be written to disks for the Hash Join operator designed for column-oriented hash tables. If this parameter is set to **on**, the number of files written to disks does not increase significantly when the Hash Join operator writes a large number of files to disks.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the optimization is enabled.
- **off** indicates that the optimization is disabled.

**Default value:** on

## plan\_cache\_mode

**Parameter description:** Specifies the policy for generating an execution plan in the **prepare** statement.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **auto** indicates that the **custom plan** or **generic plan** is selected by default.
- **force\_generic\_plan** indicates that the generic plan is forcibly used (soft parsing). **generic plan** is a plan generated after you run the **prepare** statement. The plan policy binds parameters to the plan when you run the **execute** statement and execute the plan. The advantage of this scheme is that repeated optimizer overheads can be avoided in each execution. The disadvantage is that the plan may not be optimal when data skew occurs for the bound parameters and may result in poor plan execution performance.
- **force\_custom\_plan** indicates that the custom plan is forcibly used (hard parse). **custom plan** is a plan generated after you run the **prepare** statement where parameters in the **execute** statement is embedded. The **custom plan** generates a plan based on specific parameters in the **execute** statement. This scheme generates a preferred plan based on specific parameters each time and has good execution performance. The disadvantage is that the plan needs to be regenerated before each execution, resulting in a large amount of repeated optimizer overhead.

### NOTE

This parameter is valid only for the **prepare** statement. It is used when the parameterized field in the **prepare** statement has severe data skew.

**Default value:** auto

## enable\_hypo\_index

**Parameter description:** Determines whether the optimizer considers virtual indexes when executing the **EXPLAIN** command.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** off

## enable\_auto\_explain

**Parameter description:** Specifies whether to enable the function of automatically printing execution plans. This parameter is used to locate slow stored procedures or slow queries and is valid only for the currently connected primary database node.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean. The value **true** indicates that the function is enabled, and the value **false** indicates that the function is disabled.

**Default value:** false

## auto\_explain\_level

**Parameter description:** Specifies the log level for automatically printing execution plans.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated type. The value can be **LOG** or **NOTICE**. **LOG** indicates that the execution plan is printed in logs. **NOTICE** indicates that the execution plan is printed in notification mode.

**Default value:** LOG

## auto\_explain\_log\_min\_duration

**Parameter description:** Specifies how long execution plans are automatically printed for. Plans can be printed only when the time required to execute the plans is greater than the value of **auto\_explain\_log\_min\_duration**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647. The unit is ms.

- **0:** All executed plans are generated.
- **3000:** All execution plans will be generated after the execution of a statement takes more than 3000 ms.

**Default value:** 0

## query\_dop

**Parameter description:** Specifies the user-defined degree of parallelism (DOP). This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 64. If the fixed SMP function is enabled, the system uses the fixed DOP.

 **NOTE**

After enabling concurrent queries, ensure you have sufficient CPU, memory, and network to achieve the optimal performance.

**Default value:** 1

## **enable\_startwith\_debug**

**Parameter description:** Specifies whether to enable the **start with** or **connect by** parameter for debugging. If this parameter is enabled, information about all tail columns related to the **start with** feature is displayed.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean. The value **true** indicates that the function is enabled, and the value **false** indicates that the function is disabled.

**Default value:** false

## **enable\_global\_plancache**

**Parameter description:** Specifies whether to share the cache of the PBE query execution plan. Enabling this function can reduce the memory usage of database nodes in high concurrency scenarios.

When **enable\_global\_plancache** is enabled, the default value of **local\_syscache\_threshold** is equal to or greater than 16 MB to ensure that GPC takes effect. If the value of **local\_syscache\_threshold** is less than 16 MB, set it to 16 MB. If the value is greater than 16 MB, do not change it.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the execution plan of the PBE query is shared in the cache.
- **off** indicates no sharing.

**Default value:** off

## **gpc\_clean\_timeout**

**Parameter description:** When **enable\_global\_plancache** is set to **on**, if a plan in the shared plan list is not used within the period specified by **gpc\_clean\_timeout**, the plan will be deleted. This parameter is used to control the retention period of a shared plan that is not used.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 300 to 86400. The unit is s.

**Default value:** 1800, that is, 30 minutes

## **17.3.9 Error Reporting and Logging**

### 17.3.9.1 Logging Destination

#### log\_destination

**Parameter description:** GaussDB supports several methods of logging server messages. Set this parameter to a list of desired log destinations separated by commas. (For example, **log\_destination** can be set to "**stderr, csvlog**".)

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

The valid values are **stderr**, **csvlog**, **syslog**, and **eventlog**.

- **stderr** indicates that logs are printed to the screen.
- **csvlog** indicates that logs are output in comma separated value (CSV) format. The prerequisite for generating logs in CSV format is that **logging\_collector** must be set to **on**. For details, see [Using CSV Log Output](#).
- **syslog** indicates that logs are recorded using the syslog of the OS. GaussDB can record logs using syslog from **LOCAL0** to **LOCAL7**. For details, see [syslog\\_facility](#). To record logs using syslog, add the following information to syslog daemon's configuration file:  
`local0.* /var/log/bmm`

**Default value:** **stderr**

#### logging\_collector

**Parameter description:** Specifies whether to enable the logger process to collect logs. This process captures log messages sent to **stderr** or **csvlog** and redirects them into log files.

This method is more effective than recording logs to syslog because some types of messages cannot be displayed in syslog output, such as messages indicating the loading failures of dynamic link libraries and error messages generated by scripts.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

It is possible to log to **stderr** without using the logging collector and the log messages will go to where the server's **stderr** is directed. However, this method is only suitable for low log volumes due to difficulties in rotating log files.

---

**Value range:** Boolean

- **on** indicates that the log collection is enabled.
- **off** indicates that the log collection is disabled.

**Default value:** **on**

#### log\_directory

**Parameter description:** Specifies the directory for storing log files when **logging\_collector** is set to **on**. The value can be an absolute path, or relative to



the data directory. The **log\_directory** parameter can be dynamically modified using the **gs\_guc reload** command. Only users with the **sysadmin** permission can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- If this parameter is set to an invalid path, the database cannot be started.
- If you modify the **log\_directory** parameter using the **gs\_guc reload** command, and the specified path is valid, the log files are output to this new path. If the specified path is invalid, the log files are output to the valid path set last time and the database operation is not affected. The invalid value is still written into the configuration file.
- In the sandbox environment, the path cannot contain **/var/chroot**. For example, if the absolute path of log is **/var/chroot/var/lib/log/Ruby/pg\_log/cn\_log**, you only need to set the path to **/var/lib/log/Ruby/pg\_log/cn\_log**.

---

 **NOTE**

- Valid path: Users have read and write permissions on the path.
- Invalid path: Users do not have read or write permissions on an invalid path.

**Value range:** a string

**Default value:** specified during installation

## log\_filename

**Parameter description:** Specifies the names of generated log files when **logging\_collector** is set to **on**. The value is treated as a strftime pattern, so %-escapes can be used to specify time-varying file names. Only users with the **sysadmin** permission can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

- You are advised to use %-escapes to specify the log file names for efficient management of log files.
- If **log\_destination** is set to **csvlog**, log files are output in CSV format with timestamped names, for example, **server\_log.1093827753.csv**.

---

**Value range:** a string

**Default value:** postgresql-%Y-%m-%d\_%H%M%S.log

## log\_file\_mode

**Parameter description:** Specifies the permissions of log files when **logging\_collector** is set to **on**. The parameter value is usually a number in the format acceptable to the **chmod** and **umask** system calls.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

- Before setting this parameter, set **log\_directory** to store the logs to a directory other than the data directory.
  - Do not make the log files world-readable because they might contain sensitive data.
- 

**Value range:** an octal integer ranging from 0000 to 0777 (that is, 0 to 511 in the decimal format)

#### NOTE

- **0600** indicates that log files are readable and writable only to the server administrator.
- **0640** indicates that log files are readable and writable to members of the administrator's group.

**Default value:** 0600

## log\_truncate\_on\_rotation

**Parameter description:** Specifies the writing mode of the log files when **logging\_collector** is set to **on**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

A setting example is as follows:

Assume that you want logs to be kept for 7 days, a log file generated each day to be named **server\_log.Mon** on Monday, **server\_log.Tue** on Tuesday, and so forth, and this week's log files to be overwritten by next week's log files. Then you can set **log\_filename** to **server\_log.%a**, **log\_truncate\_on\_rotation** to **on**, and **log\_rotation\_age** to **1440** (indicating that the valid duration of the log file is 24 hours).

**Value range:** Boolean

- **on** indicates that GaussDB overwrites the existing log files of the same name on the server.
- **off** indicates that GaussDB appends the logging messages to the existing log files of the same name on the server.

**Default value:** off

## log\_rotation\_age

**Parameter description:** Specifies the interval for creating a log file when **logging\_collector** is set to **on**. If the duration from the time when the last log file was created to the current time is greater than the value of **log\_rotation\_age**, a new log file will be generated.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 35791394. The unit is min. **0** indicates that the time-based creation of new log files is disabled.

**Default value:** 1440

## log\_rotation\_size

**Parameter description:** Specifies the maximum size of a server log file when **logging\_collector** is set to **on**. If the total size of messages in a log file exceeds the specified value, a log file will be generated.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to  $INT\_MAX/1024$ . The unit is KB.

**0** indicates that the capacity-based creation of new log files is disabled.

It is recommended that the unit of the value be MB or bigger, so that log files can be of proper size.

**Default value:** 20 MB

## syslog\_facility

**Parameter description:** Specifies the syslog facility to be used when **log\_destination** is set to **syslog**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values. Valid values are **local0**, **local1**, **local2**, **local3**, **local4**, **local5**, **local6**, and **local7**.

**Default value:** local0

## syslog\_ident

**Parameter description:** Specifies the identifier of GaussDB messages in syslog logs when **log\_destination** is set to **syslog**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** postgres

## event\_source

**Parameter description:** This parameter takes effect only in a Windows environment and is not supported in GaussDB. It specifies the identifier of GaussDB messages in logs when **log\_destination** is set to **eventlog**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** PostgreSQL

## 17.3.9.2 Logging Time

### client\_min\_messages

**Parameter description:** Specifies which level of messages will be sent to the client. Each level covers all the levels following it. The lower the level is, the fewer messages are sent.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

A same value for **client\_min\_messages** and **log\_min\_messages** does not indicate the same level.

---

**Value range:** enumerated values. Valid values are **debug**, **debug5**, **debug4**, **debug3**, **debug2**, **debug1**, **info**, **log**, **notice**, **warning**, **error**, **fatal**, and **panic**. For details about the parameter, see [Table 17-5](#). If the configured level is higher than **error**, for example, **fatal** or **panic**, the system changes the level to **error** by default.

**Default value:** notice

### log\_min\_messages

**Parameter description:** Specifies which level of messages will be written into the server log. Each level covers all the levels following it. The lower the level is, the fewer messages will be written into the log.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

---

#### NOTICE

A same value for **client\_min\_messages** and **log\_min\_messages** does not indicate the same level. For some log information, after this parameter is enabled, you also need to set **logging\_module** to enable log printing for the corresponding module.

---

**Value range:** enumerated values. Valid values are **debug**, **debug5**, **debug4**, **debug3**, **debug2**, **debug1**, **info**, **log**, **notice**, **warning**, **error**, **fatal**, and **panic**. For details about the parameter, see [Table 17-5](#).

**Default value:** warning

### log\_min\_error\_statement

**Parameter description:** Controls which SQL statements that cause an error condition are recorded in the server log.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values. Valid values are **debug**, **debug5**, **debug4**, **debug3**, **debug2**, **debug1**, **info**, **log**, **notice**, **warning**, **error**, **fatal**, and **panic**. For details about the parameter, see [Table 17-5](#).

---

 NOTE

- The default is **error**, indicating that statements causing errors, log messages, fatal errors, or panics will be logged.
- **panic** indicates that this feature is disabled.

**Default value:** error

## log\_min\_duration\_statement

**Parameter description:** Specifies the threshold for logging the duration of a completed statement. If a statement runs for a period equal to or greater than the specified value, its duration will be logged.

Setting **log\_min\_duration\_statement** makes it easy to trace query statements that need to be optimized. For clients using extended query protocols, the time required for parsing, binding, and executing steps are logged independently.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

---

NOTICE

When using this option together with [log\\_statement](#), the text of statements that are logged because of **log\_statement** will not be repeated in the duration log message. If you are not using **syslog**, it is recommended that you log the process ID (PID) or session ID using [log\\_line\\_prefix](#) so that you can link the statement message to the later duration message.

---

**Value range:** an integer ranging from -1 to 2147483647. The unit is ms.

- If this parameter is set to **250**, all SQL statements that run for 250 ms or longer will be logged.
- **0** indicates that the execution durations of all the statements are logged.
- **-1** indicates that the duration logging is disabled.

**Default value:** 3s (that is, 3000 ms)

## backtrace\_min\_messages

**Parameter description:** Prints the function's stack information to the server's log file if the information generated is equal to or greater than the level specified by this parameter.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

---

NOTICE

This parameter is used to locate problems on-site. Frequent stack printing will affect the system's overhead and stability. Therefore, you are advised not to set **backtrace\_min\_messages** to a level other than fatal and panic when locating problems.

---

**Value range:** enumerated values

Valid values are **debug**, **debug5**, **debug4**, **debug3**, **debug2**, **debug1**, **info**, **log**, **notice**, **warning**, **error**, **fatal**, and **panic**. For details about the parameter, see [Table 17-5](#).

**Default value:** panic

[Table 17-5](#) explains message severities used by GaussDB. If logging output is sent to **syslog** or **eventlog**, the severities are translated as shown in the table. (Note that the translation takes effect only in a Windows environment where GaussDB does not involve this parameter.)

**Table 17-5** Message severity levels

| Severity   | Description                                                                                                                                    | System Log | Event Log       |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------|------------|-----------------|
| debug[1-5] | Provides detailed debug information.                                                                                                           | DEBUG      | INFORMATIO<br>N |
| log        | Reports information of interest to administrators, for example, checkpoint activity.                                                           | INFO       | INFORMATIO<br>N |
| info       | Provides information implicitly requested by users, for example, output from <b>VACUUM VERBOSE</b> .                                           | INFO       | INFORMATIO<br>N |
| notice     | Provides information that might be helpful to users, for example, truncation of long identifiers and index created as part of the primary key. | NOTICE     | INFORMATIO<br>N |
| warning    | Provides warnings of likely problems, for example, <b>COMMIT</b> outside a transaction block.                                                  | NOTICE     | WARNING         |
| error      | Reports an error that causes a command to terminate.                                                                                           | WARNING    | ERROR           |
| fatal      | Reports the reason that causes a session to terminate.                                                                                         | ERR        | ERROR           |
| panic      | Reports an error that caused all database sessions to terminate.                                                                               | CRIT       | ERROR           |

### 17.3.9.3 Logging Content

#### debug\_print\_parse

**Parameter description:** Specifies whether to print parsing tree results.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the printing is enabled.
- **off** indicates that the printing is disabled.

**Default value:** off

#### debug\_print\_rewritten

**Parameter description:** Specifies whether to print query rewriting results.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the printing is enabled.
- **off** indicates that the printing is disabled.

**Default value:** off

#### debug\_print\_plan

**Parameter description:** Specifies whether to print the query execution plan to logs.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the printing is enabled.
- **off** indicates that the printing is disabled.

**Default value:** off

#### NOTICE

- Debugging information about **debug\_print\_parse**, **debug\_print\_rewritten**, and **debug\_print\_plan** are printed only when the log level is set to **log** or higher. When these parameters are set to **on**, their debugging information will be recorded in server logs and will not be sent to client logs. You can change the log level by setting **client\_min\_messages** and **log\_min\_messages**.
- Do not invoke the **gs\_encrypt\_aes128** and **gs\_decrypt\_aes128** functions when **debug\_print\_plan** is set to **on**, preventing the risk of sensitive information disclosure. You are advised to filter parameter information of the **gs\_encrypt\_aes128** and **gs\_decrypt\_aes128** functions in the log files generated when **debug\_print\_plan** is set to **on** before providing the log files to external maintenance engineers for fault locating. After you finish using the logs, delete them as soon as possible.

## debug\_pretty\_print

**Parameter description:** Indents the logs produced by **debug\_print\_parse**, **debug\_print\_rewritten**, and **debug\_print\_plan**. The output format is more readable but much longer than that generated when this parameter is set to **off**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the indentation is enabled.
- **off** indicates that the indentation is disabled.

**Default value:** on

## log\_checkpoints

**Parameter description:** Specifies whether the statistics on checkpoints and restart points are recorded in the server logs. When this parameter is set to **on**, statistics on checkpoints and restart points are recorded in the log messages, including the number of buffers written and the time spent in writing them.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the statistics on checkpoints and restart points are recorded in the server logs.
- **off** indicates that the statistics on checkpoints and restart points are not recorded in the server logs

**Default value:** off

## log\_connections

**Parameter description:** Specifies whether to record connection request information of the client.

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).



#### NOTICE

Some client programs, such as gsql, attempt to connect twice while determining if a password is required. In this case, duplicate "connection receive" messages do not necessarily indicate a problem.

**Value range:** Boolean

- **on** indicates that the request information is recorded.
- **off** indicates that the request information is not recorded.

**Default value:** off

## log\_disconnections

**Parameter description:** Specifies whether to record disconnection request information of the client.

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the request information is recorded.
- **off** indicates that the request information is not recorded.

**Default value:** off

## log\_duration

**Parameter description:** Specifies whether to record the duration of every completed SQL statement. For clients using extended query protocols, the time required for parsing, binding, and executing steps are logged independently.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **off:** Compared with this option, [log\\_min\\_duration\\_statement](#) forcibly records the query text.
- If this parameter is set to **on** and [log\\_min\\_duration\\_statement](#) is set to a positive value, the duration of each completed statement is logged but the query text is included only for statements exceeding the threshold. This behavior can be used for gathering statistics in high-load situation.

**Default value:** off

## log\_error\_verbosity

**Parameter description:** Specifies the amount of detail written in the server log for each message that is logged.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **terse** indicates that the output excludes the DETAIL, HINT, QUERY, and CONTEXT error information.

- **verbose** indicates that the output includes the SQLSTATE error code, the source code file name, function name, and number of the line in which the error occurs.
- **default** indicates that the output includes the DETAIL, HINT, QUERY, and CONTEXT error information, and excludes the SQLSTATE error code, the source code file name, function name, and number of the line in which the error occurs.

**Default value:** default

## log\_hostname

**Parameter description:** By default, connection log messages only show the IP address of the connecting host. The host name can be recorded when this parameter is set to **on**. It may take some time to parse the host name. Therefore, the database performance may be affected.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the host name is simultaneously recorded.
- **off** indicates that the host name is not simultaneously recorded.

**Default value:** off

## log\_line\_prefix

**Parameter description:** Specifies the prefix format of each log information. A prefix is a printf-style string that is output at the beginning of each line of the log. The "escape sequences" which begin with **%** are replaced with status information as listed in [Table 17-6](#).

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Table 17-6** Escape characters

| Escape Character | Effect                                                                                                                                           |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| %a               | Application name                                                                                                                                 |
| %u               | Username                                                                                                                                         |
| %d               | Database name                                                                                                                                    |
| %r               | Remote host name or IP address and remote port. If <b>log_hostname</b> is set to <b>off</b> , only the IP address and remote port are displayed. |
| %h               | Remote host name or IP address. If <b>log_hostname</b> is set to <b>off</b> , only the IP address is displayed.                                  |
| %p               | Thread ID                                                                                                                                        |
| %t               | Time stamp without milliseconds                                                                                                                  |

| Escape Character | Effect                                                                                                                                                                                                                         |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| %m               | Timestamp with milliseconds                                                                                                                                                                                                    |
| %n               | Node from which an error is reported                                                                                                                                                                                           |
| %i               | Command tag: type of command executed in the current session                                                                                                                                                                   |
| %e               | SQLSTATE error code                                                                                                                                                                                                            |
| %c               | Session ID: For details, see the note below the table.                                                                                                                                                                         |
| %l               | Number of the log line for each session or thread, starting at 1                                                                                                                                                               |
| %s               | Process startup time                                                                                                                                                                                                           |
| %v               | Virtual transaction ID (backendID/ localXID)                                                                                                                                                                                   |
| %x               | Transaction ID ( <b>0</b> indicates that no transaction ID is assigned)                                                                                                                                                        |
| %q               | Produces no output. If the current thread is a backend thread, this escape sequence is ignored and subsequent escape sequences are processed. Otherwise, this escape sequence and subsequent escape sequences are all ignored. |
| %S               | Session ID                                                                                                                                                                                                                     |
| %T               | Trace ID                                                                                                                                                                                                                       |
| %%               | The character %                                                                                                                                                                                                                |

 **NOTE**

The %c escape character prints a session ID consisting of two 4-byte hexadecimal numbers separated by a period (.). The numbers are the process startup time and the process ID. Therefore, %c can also be used as a space saving way of printing those items. For example, run the following query to generate the session ID from **pg\_stat\_activity**:

```
SELECT to_hex(EXTRACT(EPOCH FROM backend_start)::integer) || '.' ||
 to_hex(pid)
FROM pg_stat_activity;
```

- If you set a nonempty value for **log\_line\_prefix**, you should usually make its last character be a space, to provide visual separation from the rest of the log line. A punctuation character can be used, too.
- Syslog generates its own timestamp and process ID information. Therefore, you do not need to include those escapes characters when you are logging in to syslog.

**Value range:** a string

**Default value:** %m %n %u %d %h %p %S %x %a

 **NOTE**

**%m %n %u %d %h %p %S %x %a** indicates the session start timestamp, error reporting node, username, database name, remote host name or IP address, thread ID, session ID, transaction ID, and application name.

## log\_lock\_waits

**Parameter description:** If the time for which a session waits to acquire a lock is longer than the value of [deadlock\\_timeout](#), this parameter specifies whether to record this message in the database. This is useful in determining if lock waits are causing poor performance.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the information is recorded.
- **off** indicates that the information is not recorded.

**Default value:** off

## log\_statement

**Parameter description:** Specifies which SQL statements are recorded. For clients using extended query protocols, logging occurs when an Execute message is received, and values of the Bind parameters are included (with any embedded single quotation marks doubled).

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

Statements that contain simple syntax errors are not logged even if **log\_statement** is set to **all**, because the log message is emitted only after basic parsing has been completed to determine the statement type. If an extended query protocol is used, statements that fail before the execution phase (during parse analysis or planning) are not logged, either. Set **log\_min\_error\_statement** to **ERROR** or lower to log such statements.

---

**Value range:** enumerated values

- **none** indicates that no statement is recorded.
- **ddl** indicates that all data definition statements, such as CREATE, ALTER, and DROP, are recorded.
- **mod** indicates that all DDL statements and data modification statements, such as INSERT, UPDATE, DELETE, TRUNCATE, and COPY FROM, are recorded.
- **all** indicates that all statements, including the PREPARE, EXECUTE, and EXPLAIN ANALYZE statements, are recorded.

**Default value:** none

## log\_temp\_files

**Parameter description:** Specifies whether to record the deletion information of temporary files. Temporary files can be created for sorting, hashing, and storing temporary querying results. If the recording is enabled, a log entry is generated for each temporary file when it is deleted.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to 2147483647. The unit is KB.

- A positive value indicates that the deletion information of temporary files whose size is larger than the specified value of **log\_temp\_files** is recorded.
- **0** indicates that the delete information of all temporary files is recorded.
- **-1** indicates that the delete information of any temporary files is not recorded.

**Default value:** -1

## log\_timezone

**Parameter description:** Specifies the time zone used for timestamps written in the server log. Different from **TimeZone**, this parameter takes effect for all sessions in the database.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. You can obtain it by querying the PG\_TIMEZONE\_NAMES view.

**Default value:** Set this parameter based on the OS time zone.

### NOTE

The default value will be changed when **gs\_initdb** is used to set system environments.

## logging\_module

**Parameter description:** Specifies whether module logs are output on the server. This parameter is a session-level parameter, and you are advised not to use the gs\_guc tool to set it.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** **off**. All the module logs are not output on the server. You can view the logs by running **SHOW logging\_module**.

```
ALL,on(),off(COMMAND,GUC,GSCLEAN,SLRU,MEM_CTL,AUTOVAC,CACHE,ADIO,SSL,GDS,TBLSPC,WLM,OBS,I
NDEX,EXECUTOR,OPFUSION,GPC,GSC,VEC_EXECUTOR,STREAM,OPT,OPT_REWRITE,OPT_JOIN,OPT_AGG,OPT
_CHOICE,OPT_SUBPLAN,OPT_SETOP,OPT_SKEW,OPT_PLANNER,UDF,COOP_ANALYZE,WLMCP,ACCELERATE,M
OT,PLANHINT,PARQUET,PGSTAT,
SNAPSHOT,XACT,HANDLE,CLOG,EC,REMOTE,CN_RETRY,PLSQL,TEXTSEARCH,SEQ,REDO,FUNCTION,PARSER,I
NSTR,WDR,SNAPSHOT,INCR_CHKPT,INCR_BG_WRITER,DBL_WRT,RTO_RPO,
HEARTBEAT,COMM_IPC,COMM_PARAM,TIMESERIES,SCHEMA,SEGMENT_PAGE,LIGHTPROXY,HOTKEY,THREA
D_POOL,OPT_AI,WALRECEIVER,USTORE,UNDO,TIMECAPSULE,GEN_COL,DCF,PLDEBUGGER,ADVIS
OR,SEC,SEC_FE,SEC_LEGER,SEC_POLICY,SEC_SDD,SEC_TDE,COMM_PROXY,COMM_POOLER,VACUUM,JOB,SPI,
NEST_COMPILE,RESOWNER,LOGICAL_DECODE,GPRC,DISASTER_READ,REPSYNC,ENCODING_CHECK)
```

**Setting method:** Run **show logging\_module** to view which modules are controllable. For example, the query output result is as follows:

```
openGauss=# show logging_module;
logging_module
```

```
ALL,on(),off(COMMAND,GUC,GSCLEAN,SLRU,MEM_CTL,AUTOVAC,CACHE,ADIO,SSL,GDS,TBLSPC,WLM,OBS,I
NDEX,EXECUTOR,OPFUSION,GPC,GSC,VEC_EXECUTOR,STREAM,OPT,OPT_REWRITE,OPT_JOIN,OPT_AGG,OPT
```

```
_CHOICE,OPT_SUBPLAN,OPT_SETOP,OPT_SKEW,OPT_PLANNER,UDF,COOP_ANALYZE,WLMCP,ACCELERATE,M
OT,PLANHINT,PARQUET,PGSTAT,
SNAPSHOT,XACT,HANDLE,CLOG,EC,REMOTE,CN_RETRY,PLSQL,TEXTSEARCH,SEQ,REDO,FUNCTION,PARSER,I
NSTR,WDR_SNAPSHOT,INCR_CHKPT,INCR_BG_WRITER,DBL_WRT,RTO_RPO,
HEARTBEAT,COMM_IPC,COMM_PARAM,TIMESERIES,SCHEMA,SEGMENT_PAGE,LIGHTPROXY,HOTKEY,THREA
D_POOL,OPT_AI,WALRECEIVER,USTORE,UNDO,TIMECAPSULE,GEN_COL,DCF,PLDEBUGGER,ADVIS
OR,SEC,SEC_FE,SEC_LEGER,SEC_POLICY,SEC_SDD,SEC_TDE,COMM_PROXY,COMM_POOLER,VACUUM,JOB,SPI,
NEST_COMPILE,RESOWNER,LOGICAL_DECODE,GPRC,DISASTER_READ,REPSYNC,ENCODING_CHECK)
(1 row)
```

Controllable modules are identified by uppercase letters, and the special ID **ALL** is used for setting all module logs. You can control the output of module logs by setting **logging\_module** to **on** or **off**. Enable log output for SSL:

```
openGauss=# set logging_module='on(SSL)';
SET
openGauss=# show
logging_module;
 logging_module

(1 row)

ALL,on(SSL),off(COMMAND,GUC,GSCLEAN,SLRU,MEM_CTL,AUTOVAC,CACHE,ADIO,GDS,TBLSPC,WLM,OBS,I
NDEX,EXECUTOR,OPFUSION,GPC,GSC,VEC_EXECUTOR,STREAM,OPT,OPT_REWRITE,OPT_JOIN,OPT_AGG,OPT
_CHOICE,OPT_SUBPLAN,OPT_SETOP,OPT_SKEW,OPT_PLANNER,UDF,COOP_ANALYZE,WLMCP,ACCELERATE,M
OT,PLANHINT,PARQUET,PGSTAT,
SNAPSHOT,XACT,HANDLE,CLOG,EC,REMOTE,CN_RETRY,PLSQL,TEXTSEARCH,SEQ,REDO,FUNCTION,PARSER,I
NSTR,WDR_SNAPSHOT,INCR_CHKPT,INCR_BG_WRITER,DBL_WRT,RTO_RPO,H
EARTBEAT,COMM_IPC,COMM_PARAM,TIMESERIES,SCHEMA,SEGMENT_PAGE,LIGHTPROXY,HOTKEY,THREA
D_POOL,OPT_AI,WALRECEIVER,USTORE,UNDO,TIMECAPSULE,GEN_COL,DCF,PLDEBUGGER,ADVISO
R,SEC,SEC_FE,SEC_LEGER,SEC_POLICY,SEC_SDD,SEC_TDE,COMM_PROXY,COMM_POOLER,VACUUM,JOB,SPI,N
EST_COMPILE,RESOWNER,LOGICAL_DECODE,GPRC,DISASTER_READ,REPSYNC,ENCODING_CHECK)
(1 row)
```

SSL log output is enabled.

The **ALL** identifier can be used to quickly enable or disable log output for all modules.

```
openGauss=# set logging_module='off(ALL)';
SET
openGauss=# show
logging_module;
 logging_module

(1 row)

ALL,on(),off(COMMAND,GUC,GSCLEAN,SLRU,MEM_CTL,AUTOVAC,CACHE,ADIO,SSL,GDS,TBLSPC,WLM,OBS,I
NDEX,EXECUTOR,OPFUSION,GPC,GSC,VEC_EXECUTOR,STREAM,OPT,OPT_REWRITE,OPT_JOIN,OPT_AGG,OPT
_CHOICE,OPT_SUBPLAN,OPT_SETOP,OPT_SKEW,OPT_PLANNER,UDF,COOP_ANALYZE,WLMCP,ACCELERATE,M
OT,PLANHINT,PARQUET,PGSTAT,
SNAPSHOT,XACT,HANDLE,CLOG,EC,REMOTE,CN_RETRY,PLSQL,TEXTSEARCH,SEQ,REDO,FUNCTION,PARSER,I
NSTR,WDR_SNAPSHOT,INCR_CHKPT,INCR_BG_WRITER,DBL_WRT,RTO_RPO,
HEARTBEAT,COMM_IPC,COMM_PARAM,TIMESERIES,SCHEMA,SEGMENT_PAGE,LIGHTPROXY,HOTKEY,THREA
D_POOL,OPT_AI,WALRECEIVER,USTORE,UNDO,TIMECAPSULE,GEN_COL,DCF,PLDEBUGGER,ADVIS
OR,SEC,SEC_FE,SEC_LEGER,SEC_POLICY,SEC_SDD,SEC_TDE,COMM_PROXY,COMM_POOLER,VACUUM,JOB,SPI,
NEST_COMPILE,RESOWNER,LOGICAL_DECODE,GPRC,DISASTER_READ,REPSYNC,ENCODING_CHECK)
(1 row)

openGauss=# set logging_module='on(ALL)';
SET
openGauss=# show
logging_module;
 logging_module

(1 row)

ALL,on(COMMAND,GUC,GSCLEAN,SLRU,MEM_CTL,AUTOVAC,CACHE,ADIO,SSL,GDS,TBLSPC,WLM,OBS,INDE
X,EXECUTOR,OPFUSION,GPC,GSC,VEC_EXECUTOR,STREAM,OPT
```

```
,OPT_REWRITE,OPT_JOIN,OPT_AGG,OPT_CHOICE,OPT_SUBPLAN,OPT_SETOP,OPT_SKEW,OPT_PLANNER,UDF,
COOP_ANALYZE,WLMCP,ACCELERATE,MOT,PLANHINT,PARQUET,PGSTAT,
SNAPSHOT,XACT,HANDLE,CLOG,EC,REMOTE,CN_RETRY,PLSQL,TEXTSEARCH,SEQ,REDO,FUNCTION,PARSER,I
NSTR,WDR_SNAPSHOT,INCR_CHKPT,INCR_BG_WRITER,DBL_WRT,RTO_RPO,HEARTB
EAT,COMM_IPC,COMM_PARAM,TIMESERIES,SCHEMA,SEGMENT_PAGE,LIGHTPROXY,HOTKEY,THREAD_POOL,
OPT_AI,WALRECEIVER,USTORE,UNDO,TIMECAPSULE,GEN_COL,DCF,PLDEBUGGER,ADVISOR,SEC
,SEC_FE,SEC_LEGER,SEC_POLICY,SEC_SDD,SEC_TDE,COMM_PROXY,COMM_POOLER,VACUUM,JOB,SPI,NEST_C
OMPILE,RESOWNER,LOGICAL_DECODE,GPRC,DISASTER_READ,REPSYNC,ENCODING_CHECK),off()
(1 row)
```

**Dependency:** The value of this parameter depends on the settings of **log\_min\_messages**.

## opfusion\_debug\_mode

**Parameter description:** Checks whether simple queries are optimized for debugging. If this parameter is set to **log**, you can view the specific reasons why queries are not optimized in the database node execution plans.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **off** indicates that reasons why queries are not optimized are not included.
- **log** indicates that reasons why queries are not optimized are included in the database node execution plan.

### NOTICE

To view the reasons why queries are not optimized in the log, set **opfusion\_debug\_mode** to **log**, **log\_min\_messages** to **debug4**, and **logging\_module** to **on(OPFUSION)**. Note that a large number of log messages may be generated. Therefore, execute only a small number of jobs during debugging.

**Default value:** off

## enable\_debug\_vacuum

**Parameter description:** Specifies whether to allow output of some VACUUM-related logs for problem locating. This parameter is used only by developers. Common users are advised not to use it.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** or **true** indicates that generation of ADIO logs is allowed.
- **off** or **false** indicates that generation of ADIO logs is not allowed.

**Default value:** off

### 17.3.9.4 Using CSV Log Output

#### Prerequisites

- The **log\_destination** parameter is set to **csvlog**.
- The **logging\_collector** parameter is set to **on**.

#### Definition of csvlog

Log lines are emitted in comma separated values (CSV) format.

An example table definition for storing CSV-format log output is shown as follows:

```
CREATE TABLE gaussdb_log
(
 log_time timestamp(3) with time zone,
 node_name text,
 user_name text,
 database_name text,
 process_id bigint,
 connection_from text,
 "session_id" text,
 session_line_num bigint,
 command_tag text,
 session_start_time timestamp with time zone,
 virtual_transaction_id text,
 transaction_id bigint,
 query_id bigint,
 module text,
 error_severity text,
 sql_state_code text,
 message text,
 detail text,
 hint text,
 internal_query text,
 internal_query_pos integer,
 context text,
 query text,
 query_pos integer,
 location text,
 application_name text
);
```

For details, see [Table 17-7](#).

**Table 17-7** Meaning of each csvlog field

| Column    | Description               | Column         | Description                     |
|-----------|---------------------------|----------------|---------------------------------|
| log_time  | Timestamp in milliseconds | module         | Module to which the log belongs |
| node_name | Node name                 | error_severity | ERRORSTATE code                 |
| user_name | Username                  | sql_state_code | SQLSTATE code                   |



| Column                 | Description                     | Column             | Description                                                                                                   |
|------------------------|---------------------------------|--------------------|---------------------------------------------------------------------------------------------------------------|
| database_name          | Database name                   | message            | Error message                                                                                                 |
| process_id             | Process ID                      | detail             | Detailed error message                                                                                        |
| connection_from        | Port number of the client host  | hint               | Prompt message                                                                                                |
| session_id             | Session ID                      | internal_query     | Internal query (This field is used to query the information leading to errors if any.)                        |
| session_line_num       | Number of lines in each session | internal_query_pos | Pointer for an internal query                                                                                 |
| command_tag            | Command tag                     | context            | Environment                                                                                                   |
| session_start_time     | Start time of a session         | query              | Character count at the position where errors occur                                                            |
| virtual_transaction_id | Regular transaction             | query_pos          | Pointer at the position where errors occur                                                                    |
| transaction_id         | Transaction ID                  | location           | Position where errors occur in the GaussDB source code if <b>log_error_verbosity</b> is set to <b>verbose</b> |
| query_id               | Query ID                        | application_name   | Application name                                                                                              |

Run the following command to import a log file to this table:

```
COPY gaussdb_log FROM '/opt/data/pg_log/logfile.csv' WITH csv;
```

#### NOTE

The log name (**logfile.csv**) here needs to be replaced with the name of a log generated.

## Simplifying Input

Simplify importing CSV log files by performing the following operations:

- Set **log\_filename** and **log\_rotation\_age** to provide a consistent, predictable naming solution for log files. By doing this, you can predict when an individual log file is complete and ready to be imported.
- Set **log\_rotation\_size** to **0** to disable size-based log rollback, as it makes the log file name difficult to predict.

- Set **log\_truncate\_on\_rotation** to **on** so that old log data cannot be mixed with the new one in the same file.

## 17.3.10 Alarm Detection

During the running of the database, error scenarios can be detected so that users are informed of the errors in time. You can view the **system\_alarm** log written by the alarm in the *\$GAUSSLOG/cm* directory.

### enable\_alarm

**Parameter description:** Specifies whether to enable the alarm detection thread to detect fault scenarios that may occur in the database.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the alarm detection thread is enabled.
- **off** indicates that the alarm detection thread is disabled.

**Default value:** on

 **NOTE**

This parameter takes effect only on DNs.

### connection\_alarm\_rate

**Parameter description:** Specifies the ratio restriction on the maximum number of allowed parallel connections to the database. The maximum number of concurrent connections to the database is **max\_connections** x **connection\_alarm\_rate**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0.0 to 1.0

**Default value:** 0.9

### alarm\_report\_interval

**Parameter description:** specifies the interval at which an alarm is reported.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. The unit is s.

**Default value:** 10

### alarm\_component

**Parameter description:** Certain alarms are suppressed during alarm reporting. That is, the same alarm will not be repeatedly reported by an instance within the period specified by **alarm\_report\_interval**. Its default value is **10s**. In this case, the parameter specifies the location of the alarm component that is used to process alarm information. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- If **--alarm-type** in the **gs\_preinstall** script is set to **5**, no third-party component is connected and alarms are written into the **system\_alarm** log. In this case, the value of **alarm\_component** is **/opt/huawei/snas/bin/snas\_cm\_cmd**.
- If **--alarm-type** in the **gs\_preinstall** script is set to **1**, a third-party component is connected. In this case, the value of **alarm\_component** is the absolute path of the executable program of the third-party component.

**Default value:** **/opt/huawei/snas/bin/snas\_cm\_cmd**

## table\_skewness\_warning\_threshold

**Parameter description:** Specifies the threshold for triggering a table skew alarm.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a floating point number ranging from 0 to 1

**Default value:** 1

## table\_skewness\_warning\_rows

**Parameter description:** Specifies the number of rows for triggering a table skew alarm.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

**Default value:** 100000

## 17.3.11 Statistics During the Database Running

### 17.3.11.1 Query and Index Statistics Collector

The query and index statistics collector is used to collect statistics during database running. The statistics include the times of inserting and updating a table and index, the number of disk blocks and tuples, and the time required for the last cleanup and analysis on each table. The statistics can be viewed by querying system view families **pg\_stats** and **pg\_statistic**. The following parameters are used to set the statistics collection feature in the server scope.

## track\_activities

**Parameter description:** Collects statistics about the commands that are being executed in session. For a stored procedure, if this parameter is enabled, you can view the PERFORM statement, stored procedure calling statement, SQL statement, and OPEN CURSOR statement that are being executed in the stored procedure in the **pg\_stat\_activity** view.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the statistics collection function is enabled.
- **off** indicates that the statistics collection function is disabled.

**Default value:** on

## track\_counts

**Parameter description:** Collects statistics about database activities.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the statistics collection function is enabled.
- **off** indicates that the statistics collection function is disabled.

### NOTE

Database statistics are required when the autovacuum process checks for databases that need to be vacuumed. Therefore, the default value is set to **on**.

**Default value:** on

## track\_io\_timing

**Parameter description:** Collects statistics about I/O timing in the database. The I/O timing statistics can be queried by using the **pg\_stat\_database** parameter.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- If this parameter is set to **on**, the collection function is enabled. In this case, the collector repeatedly queries the operating system at the current time. As a result, large number of costs may occur on some platforms. Therefore, the default value is set to **off**.
- **off** indicates that the statistics collection function is disabled.

**Default value:** off

## track\_functions

**Parameter description:** Collects statistics of the number and duration of function invocations.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

When the SQL functions are set to inline functions queried by the invoking, these SQL functions cannot be traced no matter these functions are set or not.

---

**Value range:** enumerated values

- **pl** indicates that only procedural language functions are traced.
- **all** indicates that SQL language functions area traced.
- **none** indicates that the function tracing function is disabled.

**Default value:** none

## track\_activity\_query\_size

**Parameter description:** Specifies byte counts of the current running commands used to trace each active session.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 100 to 102400

**Default value:** 1024

## track\_procedure\_sql

**Parameter description:** Specifies whether the SQL statements that are being executed in the stored procedure are printed in the **query** column in the **pg\_stat\_activity** system catalog.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on**: indicates that when a stored procedure is called, the statements that are being executed by the stored procedure are printed in the query column of **pg\_stat\_activity**.
- **off**: indicates that when a stored procedure is called, only statements for invoking the stored procedure are printed in the **query** column of **pg\_stat\_activity**.

**Default value:** on

## update\_process\_title

**Parameter description:** Collects statistics updated with a process name each time the server receives a new SQL statement.

The process name can be viewed by running the **ps** command.

This is an INTERNAL parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the statistics collection function is enabled.
- **off** indicates that the statistics collection function is disabled.

**Default value:** off

## stats\_temp\_directory

**Parameter description:** Specifies the directory for storing temporary statistics. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

If a RAM-based file system directory is used, the actual I/O cost can be lowered and the performance can be improved.

---

**Value range:** a string

**Default value:** pg\_stat\_tmp

## track\_thread\_wait\_status\_interval

**Parameter description:** Specifies the interval of collecting the thread status information.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 to 1 day. The unit is min.

**Default value:** 30 min

## enable\_save\_datachanged\_timestamp

**Parameter description:** Specifies whether to record the time when **INSERT**, **UPDATE**, **DELETE**, or **EXCHANGE/TRUNCATE/DROP PARTITION** is performed on table data.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the time when an operation is performed on table data will be recorded.
- **off** indicates that the time when an operation is performed on table data will not be recorded.

**Default value:** on

## track\_sql\_count

**Parameter description:** Collects statistics about the statements (**SELECT**, **INSERT**, **UPDATE**, **MERGE INTO**, and **DELETE**) that are being executed in a session.

In the x86-based centralized deployment scenario, the hardware configuration specifications are 32-core CPU and 256 GB memory. When the Benchmark SQL 5.0 tool is used to test performance, the performance fluctuates by about 0.8% by enabling or disabling this parameter.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the statistics collection function is enabled.
- **off** indicates that the statistics collection function is disabled.

**Default value:** on

 **NOTE**

- The **track\_sql\_count** parameter is restricted by the **track\_activities** parameter when the **gs\_sql\_count** or **pgxc\_sql\_count** view is queried.
  - If **track\_activities** is set to **on** and **track\_sql\_count** is set to **off**, a warning message indicating that **track\_sql\_count** is disabled will be displayed in logs when the **gs\_sql\_count** view is queried.
  - If both **track\_activities** and **track\_sql\_count** are set to **off**, two warning messages indicating that **track\_activities** is disabled and **track\_sql\_count** is disabled will be displayed in logs when the views are queried.
  - If **track\_activities** is set to **off** and **track\_sql\_count** is set to **on**, a warning message indicating that **track\_activities** is disabled will be displayed in logs when the views are queried.
- If **track\_sql\_count** is set to **off**, querying the **gs\_sql\_count** or **pgxc\_sql\_count** view returns 0.

### 17.3.11.2 Performance Statistics

During the running of the database, the lock access, disk I/O operation, and invalid message processing are involved. All these operations are the bottleneck of the database performance. The performance statistics provided by GaussDB can facilitate the performance fault location.

### Generating Performance Statistics Logs

**Parameter description:** For each query, the following four parameters record the performance statistics of corresponding modules in the server log:

- The **log\_parser\_stats** parameter records the performance statistics of a parser in the server log.
- The **log\_planner\_stats** parameter records the performance statistics of a query optimizer in the server log.
- The **log\_executor\_stats** parameter records the performance statistics of an executor in the server log.
- The **log\_statement\_stats** parameter records the performance statistics of the whole statement in the server log.

All these parameters can only provide assistant analysis for administrators, which are similar to the `getrusage()` of the Linux OS.

These are SUSE parameters. Set them based on instructions in [Table 17-1](#).

---

**NOTICE**

- The **log\_statement\_stats** records the total statement statistics whereas other parameters record statistics only about their corresponding modules.
- The **log\_statement\_stats** parameter cannot be enabled together with any parameter recording statistics about a module.

---

**Value range:** Boolean

- **on** indicates that performance statistics are recorded.
- **off** indicates that performance statistics are not recorded.

**Default value:** off

## 17.3.12 Automatic Vacuuming

The **autovacuum** process automatically runs the **VACUUM** and **ANALYZE** statements to recycle the record space marked as deleted and update statistics about the table.

### autovacuum

**Parameter description:** Specifies whether to start the **autovacuum** process in the database. Ensure that the **track\_counts** parameter is set to **on** before starting the automatic vacuuming process.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

#### NOTE

- Set the **autovacuum** parameter to **on** if you want to start the automatic vacuuming of abnormal two-phase transactions when the system recovers from faults.
- If **autovacuum** is set to **on** and **autovacuum\_max\_workers** to 0, the autovacuum process is started only when the system recovers from faults to clean up abnormal two-phase transactions.
- If **autovacuum** is set to **on** and the value of **autovacuum\_max\_workers** is greater than 0, the system will automatically vacuum the two-phase transactions and processes after recovering from faults.

**Value range:** Boolean

- **on** indicates that the **autovacuum** process is started.
- **off** indicates that the **autovacuum** process is not started.

**Default value:** on

### autovacuum\_mode

**Parameter description:** Specifies whether the autoanalyze or autovacuum function is started. This parameter is valid only when **autovacuum** is set to **on**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** enumerated values

- **analyze** indicates that only autoanalyze is performed.
- **vacuum** indicates that only autovacuum is performed.
- **mix** indicates that both autoanalyze and autovacuum are performed.
- **none** indicates that neither of them is performed.

**Default value:** mix



## autoanalyze\_timeout

**Parameter description:** Specifies the timeout period of autoanalyze. If the duration of autoanalyze on a table exceeds the value of **autoanalyze\_timeout**, the autoanalyze is automatically canceled.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 0 to 2147483. The unit is s.

**Default value:** 5 min (300s)

## autovacuum\_io\_limits

**Parameter description:** Specifies the upper limit of I/Os triggered by the autovacuum process per second.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer. The value can be **-1** or a number ranging from 0 to 1073741823. **-1** indicates that the default cgroup is used.

**Default value:** -1

## log\_autovacuum\_min\_duration

**Parameter description:** Records each step performed by the autovacuum process to the server log when the execution time of the autovacuum process is equal to or greater than a certain value. This parameter helps track the autovacuum behavior.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

A setting example is as follows:

Set the **log\_autovacuum\_min\_duration** parameter to 250 ms to record the actions of autovacuum if it runs for 250 ms or longer.

**Value range:** an integer ranging from -1 to 2147483647. The unit is ms.

- **0** indicates that all autovacuum actions are recorded in the log.
- **-1** indicates that all autovacuum actions are not recorded in the log.
- A value greater than 0 indicates that a message is recorded when an autovacuum action is skipped due to a lock conflict.

**Default value:** -1

## autovacuum\_max\_workers

**Parameter description:** Specifies the maximum number of autovacuum worker threads that can run at the same time. The upper limit of this parameter is related to the values of **max\_connections** and **job\_queue\_processes**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer. The minimum value is **0**, indicating that autovacuum is not enabled. The theoretical maximum value is **262143**, but the actual maximum value is a dynamic value calculated by the following formula: 262143 -

**max\_connections - job\_queue\_processes** - Number of auxiliary threads - Number of autovacuum launcher threads - 1. The number of auxiliary threads and the number of autovacuum launcher threads are specified by two macros. Their default values are **20** and **2** respectively.

**Default value:** 3

## autovacuum\_naptime

**Parameter description:** Specifies the interval between activity rounds for the autovacuum process.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 1 to 2147483. The unit is s.

**Default value:** 10 min (600s)

## autovacuum\_vacuum\_threshold

**Parameter description:** Specifies the threshold for triggering the **VACUUM** operation. When the number of deleted or updated records in a table exceeds the specified threshold, the VACUUM operation is executed on this table.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 0 to 2147483647.

**Default value:** 50

## autovacuum\_analyze\_threshold

**Parameter description:** Specifies the threshold for triggering the **ANALYZE** operation. When the number of deleted, inserted, or updated records in a table exceeds the specified threshold, the **ANALYZE** operation is executed on this table.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 0 to 2147483647.

**Default value:** 50

## autovacuum\_vacuum\_scale\_factor

**Parameter description:** Specifies a fraction of the table size added to the **autovacuum\_vacuum\_threshold** parameter when deciding whether to vacuum a table.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** a floating-point number ranging from 0.0 to 100.0

**Default value:** 0.2

## autovacuum\_analyze\_scale\_factor

**Parameter description:** Specifies a fraction of the table size added to the **autovacuum\_analyze\_threshold** parameter when deciding whether to analyze a table.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** a floating-point number ranging from 0.0 to 100.0

**Default value:** 0.1

## autovacuum\_freeze\_max\_age

**Parameter description:** Specifies the maximum age (in transactions) that a table's **pg\_class.relfrozenxid** field can attain before a VACUUM operation is performed.

- The old files under the subdirectory of **pg\_clog/** can also be deleted by the **VACUUM** operation.
- Even if the **autovacuum** process is not started, the system will invoke the process to prevent transaction ID wraparound.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 100000 to 576460752303423487

**Default value:** 4000000000

## autovacuum\_vacuum\_cost\_delay

**Parameter description:** Specifies the value of the cost delay used in the **autovacuum** operation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from -1 to 100. The unit is ms. -1 indicates that the normal vacuum cost delay is used.

**Default value:** 20ms

## autovacuum\_vacuum\_cost\_limit

**Parameter description:** Sets the value of the cost limit used in the **autovacuum** operation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from -1 to 10000 -1 indicates that the normal vacuum cost limit is used.

**Default value:** -1

## defer\_csn\_cleanup\_time

**Parameter description:** Specifies the local recycling interval, in ms.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*

**Default value:** 5s (5000 ms)

## 17.3.13 Default Settings of Client Connection

### 17.3.13.1 Statement Behavior

This section describes related default parameters involved in the execution of SQL statements.

#### search\_path

**Parameter description:** Specifies the order in which schemas are searched when an object is referenced with no schema specified. The value of this parameter consists of one or more schema names. Different schema names are separated by commas (,).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

- If the schema of temporary tables exists in the current session, the schema can be listed in the search path by using the alias **pg\_temp**, for example, '**pg\_temp,public**'. The schema of temporary tables has the highest search priority and is always searched before all the other schemas specified in **pg\_catalog** and **search\_path**. Therefore, do not explicitly specify **pg\_temp** to be searched after other schemas in **search\_path**. This setting will not take effect and an error message will be displayed. If the alias **pg\_temp** is used, the temporary schema will be searched only for tables, views, and data types, and not for functions or operators.
- The system catalog schema, **pg\_catalog**, has the second highest search priority and is the first to be searched among all the schemas, excluding **pg\_temp**, specified in **search\_path**. Therefore, do not explicitly specify **pg\_catalog** to be searched after other schemas in **search\_path**. This setting will not take effect and an error message will be displayed.
- When an object is created without a specific target schema, the object will be placed in the first valid schema listed in **search\_path**. An error is reported if the search path is empty.
- The current effective value of the search path can be examined through the SQL function **current\_schema**. This is different from examining the value of **search\_path**, because the **current\_schema** function displays the first valid schema name in **search\_path**.

**Value range:** a string

 NOTE

- When this parameter is set to "**\$user**", **public**, shared use of a database (where no users have private schemas, and all share use of public), private per-user schemas and combinations of them are supported. Other effects can be obtained by modifying the default search path setting, either globally or per-user.
- When this parameter is set to a null string (""), the system automatically converts it into a pair of double quotation marks ("").
- If the content contains double quotation marks, the system considers them as insecure characters and converts each double quotation mark into a pair of double quotation marks.

**Default value:** "**\$user**",**public**

 NOTE

**\$user** indicates the name of the schema with the same name as the current session user. If the schema does not exist, **\$user** will be ignored.

## current\_schema

**Parameter description:** Specifies the current schema.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** "**\$user**",**public**

 NOTE

- **\$user** indicates the name of the schema with the same name as the current session user. If the schema does not exist, **\$user** will be ignored.
- If you need to obtain a schema during GaussDB development, use the value of **search\_path** because the schema is determined by **search\_path**. To be compatible with database A, **current\_schema** is used only to modify the value of **search\_path**.

## default\_tablespace

**Parameter description:** Specifies the default tablespace of the created objects (tables and indexes) when a **CREATE** command does not explicitly specify a tablespace.

- The value of this parameter is either the name of a tablespace, or an empty string that indicates the use of the default tablespace of the current database. If a non-default tablespace is specified, users must have CREATE privilege for it. Otherwise, creation attempts will fail.
- This parameter is not used for temporary tables. For them, the [temp\\_tablespaces](#) is used instead.
- This parameter is not used when users create databases. By default, a new database inherits its tablespace setting from the template database.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. An empty string indicates that the default tablespace is used.

**Default value:** empty

## temp\_tablespaces

**Parameter description:** Specifies one or more tablespaces to which temporary objects (temporary tables and their indexes) will be created when a CREATE command does not explicitly specify a tablespace. Temporary files for sorting large data sets are created in these tablespaces.

The value of this parameter can be a list of names of tablespaces. When there is more than one name in the list, GaussDB chooses a random tablespace from the list upon the creation of a temporary object each time. However, within a transaction, successively created temporary objects are placed in successive tablespaces in the list. If the element selected from the list is an empty string, GaussDB will automatically use the default tablespace of the current database instead.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string An empty string indicates that all temporary objects are created only in the default tablespace of the current database. For details, see [default\\_tablespace](#).

**Default value:** empty

## check\_function\_bodies

**Parameter description:** Specifies whether to enable validation of the function body string during the execution of **CREATE FUNCTION**. Verification is occasionally disabled to avoid problems, such as forward references when you restore function definitions from a dump. After the function is enabled, the word syntax of the PL/SQL in the stored procedure is verified, including the data type, statement, and expression. The SQL statements in the stored procedure are not checked in the CREATE phase. Instead, they are checked during running.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that validation of the function body string is enabled during the execution of **CREATE FUNCTION**.
- **off** indicates that validation of the function body string is disabled during the execution of **CREATE FUNCTION**.

**Default value:** on

## default\_transaction\_isolation

**Parameter description:** Specifies the default isolation level of each transaction.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

### NOTE

The current version does not support the setting of the default transaction isolation level. The default value is **read committed**. Do not change the value.

**Value range:** enumerated values

- **read committed** indicates that the data read by a transaction is committed at the moment it is read.
- **repeatable read** indicates that the data that has been read by the current transaction cannot be modified by other transactions until the current transaction completes, thereby preventing unrepeatable reads.
- **serializable**: Currently, this isolation level is not supported in GaussDB. It is equivalent to **repeatable read**.

**Default value:** read committed

## default\_transaction\_read\_only

**Parameter description:** Specifies whether each new transaction is in read-only state.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).



### CAUTION

If this parameter is set to **on**, the DML and write transactions cannot be executed.

---

**Value range:** Boolean

- **on** indicates that the transaction is in read-only state.
- **off** indicates that the transaction is in read/write state.

**Default value:** off

## default\_transaction\_deferrable

**Parameter description:** Specifies the default deferrable status of each new transaction. It currently has no effect on read-only transactions or those running at isolation levels lower than serializable.

GaussDB does not support the serializable isolation level. Therefore, the parameter takes no effect.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that a transaction is delayed by default.
- **off** indicates that a transaction is not delayed by default.

**Default value:** off

## session\_replication\_role

**Parameter description:** Specifies the behavior of replication-related triggers and rules for the current session.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

Setting this parameter will discard all the cached execution plans.

---

**Value range:** enumerated values

- **origin** indicates that the system copies operations such as insert, delete, and update from the current session.
- **replica** indicates that the system copies operations such as insert, delete, and update from other places to the current session.
- **local** indicates that the system will detect the role that has logged in to the database when using the function to copy operations and will perform related operations.

**Default value:** origin

## statement\_timeout

**Parameter description:** If the statement execution time (starting from the time the server receives the command) is longer than the duration specified by the parameter, error information is displayed and the statement exits.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#). The default value is **0**, indicating that the parameter does not take effect.

**Value range:** an integer ranging from 0 to 2147483647. The unit is ms.

**Default value:** 0

## vacuum\_freeze\_min\_age

**Parameter description:** Specifies whether VACUUM replaces the **xmin** column of a record with **FrozenXID** when scanning a table (in the same transaction).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 576460752303423487

 **NOTE**

Although you can set this parameter to any value, VACUUM will limit the effective value to half the value of [autovacuum\\_freeze\\_max\\_age](#) by default.

**Default value:** 2000000000

## vacuum\_freeze\_table\_age

**Parameter description:** Specifies when VACUUM scans the whole table and freezes old tuples. VACUUM performs a full table scan if the difference between the current transaction ID and the value of **pg\_class.relfrozenxid64** is greater than the specified time.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 576460752303423487



 NOTE

Although you can set this parameter to any value, **VACUUM** will limit the effective value to 95% of **autovacuum\_freeze\_max\_age** by default. Therefore, a periodic manual **VACUUM** has a chance to run before an anti-wraparound autovacuum is launched for the table.

**Default value:** 4000000000

## bytea\_output

**Parameter description:** Specifies the output format for values of the bytea type.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **hex** indicates that the binary data is converted to hexadecimal format.
- **escape** indicates that the traditional PostgreSQL format is used. It takes the approach of representing a binary string as a sequence of ASCII characters, while converting those bytes that cannot be represented as an ASCII character into special escape sequences.

**Default value:** hex

## xmlbinary

**Parameter description:** Specifies how binary values are to be encoded in XML.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- base64
- hex

**Default value:** base64

## xmloption

**Parameter description:** Specifies whether DOCUMENT or CONTENT is implicit when converting between XML and string values.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **document** indicates an HTML document.
- **content** indicates a common string.

**Default value:** content

## max\_compile\_functions

**Parameter description:** Specifies the maximum number of function compilation results stored in the server. Excessive functions and compilation results of stored procedures may occupy large memory space. Setting this parameter to a proper value can reduce the memory usage and improve system performance.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647

**Default value:** 1000

### 17.3.13.2 Locale and Formatting

This section describes parameters related to the time format setting.

#### DateStyle

**Parameter description:** Specifies the display format for date and time values, as well as the rules for interpreting ambiguous date input values.

This variable contains two independent components: the output format specifications (ISO, Postgres, SQL, or German) and the input/output order of year/month/day (DMY, MDY, or YMD). The two components can be set separately or together. The keywords Euro and European are synonyms for DMY. The keywords US, NonEuro, and NonEuropean are synonyms for MDY.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** ISO, MDY

#### NOTE

**gs\_initdb** will initialize this parameter so that its value is the same as that of [lc\\_time](#).

**Setting Suggestions:** The ISO format is recommended. Postgres, SQL, and German use abbreviations for time zones, such as **EST**, **WST**, and **CST**. These abbreviations can be ambiguous. For example, **CST** can represent Central Standard Time (USA) UT-6:00, Central Standard Time (Australia) UT+9:30, and China Standard Time UT+8:00. This may lead to incorrect time zone conversion and cause errors.

#### IntervalStyle

**Parameter description:** Specifies the display format for interval values.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **sql\_standard** indicates that output matching SQL standards will be generated.
- **postgres** indicates that output matching PostgreSQL 8.4 will be generated when the [DateStyle](#) parameter is set to **ISO**.
- **postgres\_verbose** indicates that output matching PostgreSQL 8.4 will be generated when the [DateStyle](#) parameter is set to **non\_ISO**.
- **iso\_8601** indicates that output matching the time interval "format with designators" defined in ISO 8601 will be generated.
- **a** indicates the output result that matches the numtodsinterval function.

---

**NOTICE**

The **IntervalStyle** parameter also affects the interpretation of ambiguous interval input.

---

**Default value:** postgres

## TimeZone

**Parameter description:** Specifies the time zone for displaying and interpreting timestamps.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. You can obtain it by querying the **pg\_timezone\_names** view.

**Default value:** PRC

 **NOTE**

**gs\_initdb** will set a time zone value that is consistent with the system environment.

## timezone\_abbreviations

**Parameter description:** Specifies the time zone abbreviations that will be accepted by the server.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. You can obtain it by querying the **pg\_timezone\_names** view.

**Default value:** Default

 **NOTE**

**Default** indicates an abbreviation that works in most of the world, which is applicable to most cases. There are also other abbreviations, such as **Australia** and **India** that can be defined for a particular installation. For other time zone abbreviations, you need to set them in the corresponding configuration files before creating the database.

## extra\_float\_digits

**Parameter description:** Adjusts the number of digits displayed for floating-point values, including float4, float8, and geometric data types. The parameter value is added to the standard number of digits (FLT\_DIG or DBL\_DIG as appropriate).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -15 to 3

 NOTE

- This parameter can be set to **3** to include partially-significant digits. It is especially useful for dumping float data that needs to be restored exactly.
- This parameter can also be set to a negative value to suppress unwanted digits.

**Default value:** 0

## client\_encoding

**Parameter description:** Specifies the client-side encoding (character set).

Set this parameter based on the situation of the front-end services. Try to keep the encoding consistent on the client and server to improve efficiency.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** encoding compatible with PostgreSQL. **UTF8** indicates that the database encoding is used.

 NOTE

- You can run the **locale -a** command to check the system-supported locales and the corresponding encodings, and select one as required.
- By default, **gs\_initdb** will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.
- To use consistent encoding for communication within the database, you are advised to retain the default value of **client\_encoding**. Modification to this parameter in the **postgresql.conf** file (by using the **gs\_guc** tool, for example) does not take effect.

**Default value:** UTF8

**Recommended value:** SQL\_ASCII or UTF8

## lc\_messages

**Parameter description:** Specifies the language in which messages are displayed.

- Acceptable values are system-related.
- On some systems, this locale category does not exist. Setting this variable will still work, but there will be no effect. In addition, translated messages for the desired language may not exist. In this case, you can still see the English messages.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

 NOTE

- You can run the **locale -a** command to check the system-supported locales and the corresponding encodings, and select one as required.
- By default, **gs\_initdb** will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.

**Default value:** C

## lc\_monetary

**Parameter description:** Specifies the display format of monetary values. It affects the output of functions such as **to\_char**. Acceptable values are system-related.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

### NOTE

- You can run the **locale -a** command to check the system-supported locales and the corresponding encodings, and select one as required.
- By default, **gs\_initdb** will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.

**Default value:** C

## lc\_numeric

**Parameter description:** Specifies the display format of numbers. It affects the output of functions such as **to\_char**. Acceptable values are system-related.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

### NOTE

- You can run the **locale -a** command to check the system-supported locales and the corresponding encodings, and select one as required.
- By default, **gs\_initdb** will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.

**Default value:** C

## lc\_time

**Parameter description:** Specifies the display format of time and locale. It affects the output of functions such as **to\_char**. Acceptable values are system-related.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

### NOTE

- You can run the **locale -a** command to check the system-supported locales and the corresponding encodings, and select one as required.
- By default, **gs\_initdb** will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.

**Default value:** C

### 17.3.13.3 Other Default Parameters

This section describes the default database loading parameters.

#### dynamic\_library\_path

**Parameter description:** Specifies the path that the system will search for a shared database file that is dynamically loadable. When a dynamically loadable module needs to be opened and the file name specified in the **CREATE FUNCTION** or **LOAD** command does not have a directory component, the system will search this path for the required file. Only the sysadmin user can access this parameter.

The value of **dynamic\_library\_path** must be a list of absolute paths separated by colons (:). When the name of a path starts with the special variable \$libdir, the variable will be replaced with the directory in which the module provided by the GaussDB is installed. For example:

```
dynamic_library_path = '/usr/local/lib/gaussdb:/opt/testgs/lib:$libdir'
```

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

#### NOTE

If the value of this parameter is set to an empty character string, the automatic path search is turned off.

**Default value:** \$libdir

#### local\_preload\_libraries

**Parameter description:** Specifies one or more shared libraries that are to be preloaded at connection start. If multiple libraries are to be loaded, separate their names with commas (,). All library names are converted to lower case unless double-quoted.

- Any user can change this option. Therefore, library files that can be loaded are restricted to those saved in the **plugins** subdirectory of the standard library installation directory. It is the database administrator's responsibility to ensure that libraries in this directory are all safe. Entries in **local\_preload\_libraries** can specify the library directory explicitly, for example, **\$libdir/plugins/mylib**, or just specify the library name, for example, **mylib**. (**mylib** is equivalent to **\$libdir/plugins/mylib**.)
- Unlike **shared\_preload\_libraries**, there are no differences in performance between loading a module at session start or doing this during the session. The intent of this feature is to allow debugging or performance-measurement libraries to be loaded into specific sessions without an explicit **LOAD** command. For example, debugging can be enabled under a given username by setting this parameter to **ALTER USER SET**.
- If a specified library is not found, the connection attempt will fail.
- Every GaussDB-supported library has a "magic block" that is checked to guarantee compatibility. For this reason, non-GaussDB-supported libraries cannot be loaded in this way.

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** empty

## 17.3.14 Lock Management

In GaussDB, a deadlock may occur when concurrently executed transactions compete for resources. This section describes parameters used for managing transaction locks.

### deadlock\_timeout

**Parameter description:** Specifies the time, in milliseconds, to wait on a lock before checking whether there is a deadlock condition. When the applied lock exceeds the preset value, the system will check whether a deadlock occurs. This parameter takes effect only for common locks.

- The check for deadlock is relatively expensive. Therefore, the server does not check it when waiting for a lock every time. Deadlocks do not frequently occur when the system is running. Therefore, the system just needs to wait on the lock for a while before checking for a deadlock. Increasing this value reduces the time wasted in needless deadlock checks, but slows down reporting of real deadlock errors. On a heavily loaded server, you may need to raise it. The value you have set needs to exceed the transaction time. By doing this, the possibility that a lock will be checked for deadlocks before it is released will be reduced.
- When [log\\_lock\\_waits](#) is set to **on**, **deadlock\_timeout** determines a waiting time to write the lock waiting time information during query execution to logs. To study the lock delay, you can set **deadlock\_timeout** to a value smaller than the normal value.

This is a **SUSET** parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647. The unit is ms.

**Default value:** 1s

### lockwait\_timeout

**Parameter description:** Specifies the timeout for attempts to acquire a lock. If the time spent in waiting for a lock exceeds the specified time, an error is reported. This parameter takes effect only for common locks.

This is a **USERSET** parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is ms.

**Default value:** 20 min

### update\_lockwait\_timeout

**Parameter description:** Specifies the maximum duration that a lock waits for concurrent updates on a row to complete when the concurrent update feature is enabled. If the time spent in waiting for a lock exceeds the specified time, an error is reported. This parameter takes effect only for common locks.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647. The unit is ms.

**Default value:** 2 min (120000 ms)

## max\_locks\_per\_transaction

**Parameter description:** Determines the average number of object locks allocated to each transaction.

**Parameter type:** integer

**Unit:** none

**Value range:** 10 to 2147483647

**Default value:** 256

**Setting method:** This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Setting suggestions:**

- The maximum number of hash tables that can be locked by a shared lock at any time is calculated based on an assumption: **max\_locks\_per\_transaction** x (**max\_connections** + **max\_prepared\_transactions**) and **max\_locks\_per\_transaction** ≥ Number of concurrent service transactions x Number of object locks added by each service transaction/(**max\_connections** + **max\_prepared\_transactions**). In this case an upper limit is determined for this parameter.
- Within the specified range, objects can be locked simultaneously at any time. You may need to increase the value of this parameter if many different tables are modified in a single transaction.
- Increasing the value of this parameter may cause GaussDB to request more System V-shared memory than the OS's default configuration allows, leading to database startup failure.
- When running a standby node, you must set this parameter to a value that is no less than that on the primary node. Otherwise, queries will not be allowed on the standby node.

## max\_pred\_locks\_per\_transaction

**Parameter description:** Specifies the average number of predicate locks allocated for each transaction.

- The size of the shared predicate lock table is calculated under the condition that a maximum of *N* independent objects need to be locked at any time. *N* = **max\_pred\_locks\_per\_transaction** x (**max\_connections** + **max\_prepared\_transactions**). Within the specified range, objects can be locked simultaneously at any time. You may need to increase the value if many different tables are modified in a single transaction. This parameter can only be set at server start.
- Increasing the value of this parameter may cause GaussDB to request more System V-shared memory than the OS's default configuration allows.



This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to *INT\_MAX*

**Default value:** 64

## gs\_clean\_timeout

**Parameter description:** Specifies the average interval for clearing temporary tables on the primary node.

- When the database connection is terminated abnormally, temporary tables may exist. In this case, you need to call the **gs\_clean** tool to clear the temporary tables in the database.
- If this parameter is set to a larger value, the time for clearing GaussDB temporary tables may be prolonged.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483. The unit is s.

**Default value:** 1 min

## partition\_lock\_upgrade\_timeout

**Parameter description:** Specifies the timeout for attempts to upgrade an exclusive lock (read allowed) to an access exclusive lock (read/write blocked) on a partitioned table during the execution of some query statements. If there are concurrent read transactions running, the lock upgrade will need to wait. This parameter sets the waiting timeout for lock upgrade attempts.

- When you do **MERGE PARTITION** and **CLUSTER PARTITION** on a partitioned table, temporary tables are used for data rearrangement and file exchange. To concurrently perform as many operations as possible on the partitions, exclusive locks are acquired for the partitions during data rearrangement and access exclusive locks are acquired during file exchange.
- Generally, a partition waits until it acquires a lock, or a timeout occurs if the partition waits for a period longer than the value specified by the [lockwait\\_timeout](#) parameter.
- When doing **MERGE PARTITION** or **CLUSTER PARTITION** on a partitioned table, an access exclusive lock needs to be acquired during file exchange. If the lock fails to be acquired, the acquisition is retried at an interval of 50 ms until timeout occurs. The **partition\_lock\_upgrade\_timeout** parameter specifies the time to wait before the lock acquisition attempt times out.
- If this parameter is set to **-1**, the lock upgrade never times out. The lock upgrade is continuously retried until it succeeds.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from -1 to 3000. The unit is s.

**Default value:** 1800

## fault\_mon\_timeout

**Parameter description:** Specifies the period for detecting lightweight deadlocks. This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1440. The unit is minute.

**Default value:** 5min

## enable\_online\_ddl\_waitlock

**Parameter description:** Specifies whether to block DDL operations to wait for the release of database locks, such as **pg\_advisory\_lock**. This parameter is mainly used in online OM operations and you are not advised to modify the settings.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that DDL operations will be blocked to wait for the lock release.
- **off** indicates that DDL operations will not be blocked.

**Default value:** off

## xloginsert\_locks

**Parameter description:** Specifies the number of locks on concurrent write-ahead logging. This parameter is used to improve the efficiency of writing write-ahead logs.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 1000

**Default value:** 16

## num\_internal\_lock\_partitions

**Parameter description:** Specifies the number of internal lightweight lock partitions. It is mainly used for performance optimization in various scenarios. The content is organized in the KV format of keywords and numbers. Different types of locks are separated by commas (.). The sequence does not affect the setting result. For example, **CLOG\_PART=256,CSNLOG\_PART=512** is equivalent to **CSNLOG\_PART=512,CLOG\_PART=256**. If you set the same keyword multiple times, only the latest setting takes effect. For example, if you set **CLOG\_PART** to **256** and **CLOG\_PART** to **2**, the value of **CLOG\_PART** is **2**. If no keyword is set, the default value is used. The usage description, maximum value, minimum value, and default value of each lock type are as follows:

- **CLOG\_PART:** number of Clog file controllers. Increasing the value of this parameter improves the Clog writing efficiency and transaction submission performance, but increases the memory usage. Decreasing the value of this parameter reduces the memory usage, but may increase the conflict of writing Clogs and affect the performance. The value ranges from 1 to 256.
- **CSNLOG\_PART:** number of CSNLOG file controllers. Increasing the value of this parameter improves the CSNLOG log writing efficiency and transaction

submission performance, but increases the memory usage. Decreasing the value of this parameter reduces the memory usage, but may increase the conflict of writing CSNLOG logs and affect the performance. The value ranges from 1 to 512.

- **LOG2\_LOCKTABLE\_PART**: two logarithms of the number of common table lock partitions. Increasing the value can improve the concurrency of obtaining locks in the normal process, but may increase the time required for transferring and clearing locks. When waiting events occur in **LockMgrLock**, you can increase the value to improve the performance. The minimum value is 4, that is, the number of lock partitions is 16. The maximum value is 16, that is, the number of lock partitions is 65536.
- **TWOPHASE\_PART**: number of partitions of the two-phase transaction lock. Increasing the value can increase the number of concurrent two-phase transaction commits. The value ranges from 1 to 64.
- **FASTPATH\_PART**: maximum number of locks that each thread can obtain without using the main lock table. When a partitioned table is read, updated, inserted, or deleted and the waiting event is LockMgrLock, you can increase the value of this parameter to prevent **LockMgrLock** from being obtained and improve performance. It is recommended that the value be equal to or greater than **Number of partitions x (1 + Number of local indexes) + Number of global indexes + 10**. Increasing the value will increase the memory usage. The value ranges from 20 to 10000.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:**

- **CLOG\_PART**: 256
- **CSNLOG\_PART**: 512
- **LOG2\_LOCKTABLE\_PART**: 4
- **TWOPHASE\_PART**: 1
- **FASTPATH\_PART**: 20

## 17.3.15 Version and Platform Compatibility

### 17.3.15.1 Compatibility with Earlier Versions

This section describes the parameters that control the backward compatibility and external compatibility of GaussDB. A backward compatible database supports applications of earlier versions. This section describes parameters used for controlling backward compatibility of a database.

#### array\_nulls

**Parameter description:** Controls whether the array input parser recognizes unquoted NULL as a null array element.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that null values can be entered in arrays.
- **off** indicates backward compatibility with the old behavior. Arrays containing the value **NULL** can still be created when this parameter is set to **off**.

**Default value:** on

## backslash\_quote

**Parameter description:** Controls whether a single quotation mark can be represented by \ in a string text.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

### NOTICE

When the string text meets the SQL standards, \ has no other meanings. This parameter only affects the handling of non-standard-conforming string texts, including escape string syntax (E'...').

**Value range:** enumerated values

- **on** indicates that the use of \ is always allowed.
- **off** indicates that the use of \ is rejected.
- **safe\_encoding** indicates that the use of \ is allowed only when client encoding does not allow **ASCII** \ within a multibyte character.

**Default value:** safe\_encoding

## escape\_string\_warning

**Parameter description:** Specifies whether to issue a warning when a backslash (\) is used as an escape in an ordinary character string.

- Applications that wish to use a backslash (\) as an escape need to be modified to use escape string syntax (E'...'). This is because the default behavior of ordinary character strings treats the backslash as an ordinary character in each SQL standard.
- This variable can be enabled to help locate codes that need to be changed.
- If E'...' is used as an escape, logs may be incomplete in some scenarios.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** on

## lo\_compat\_privileges

**Parameter description:** Specifies whether to enable backward compatibility for the privilege check of large objects.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**on** indicates that the privilege check is disabled when users read or modify large objects. This setting is compatible with versions earlier than PostgreSQL 9.0.

**Default value:** off

## quote\_all\_identifiers

**Parameter description:** Specifies whether to forcibly quote all identifiers even if they are not keywords when the database generates SQL. This will affect the output of **EXPLAIN** and the results of functions, such as `pg_get_viewdef`. For details, see the `--quote-all-identifiers` parameter of `gs_dump`.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the forcible quoting is enabled.
- **off** indicates that the forcible quoting is disabled.

**Default value:** off

## sql\_inheritance

**Parameter description:** Controls the inheritance semantics. This parameter specifies the access policy of descendant tables. **off** indicates that subtables cannot be accessed by commands. That is, the **ONLY** keyword is used by default. This setting is compatible with earlier versions.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that subtables can be accessed.
- **off** indicates that subtables cannot be accessed.

**Default value:** on

## standard\_conforming\_strings

**Parameter description:** Controls whether ordinary string texts ('...') treat backslashes as ordinary texts as specified in the SQL standard.

- Applications can check this parameter to determine how string texts will be processed.
- It is recommended that characters be escaped by using the escape string syntax (E'...').

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that backslashes are treated as ordinary texts.
- **off** indicates that backslashes are not treated as ordinary texts.

**Default value:** on

## synchronize\_seqscans

**Parameter description:** Controls sequential scans of tables to synchronize with each other, so that concurrent scans read the same data block at about the same time and share the I/O workload.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that a scan may start in the middle of the table and then "wrap around" the end to cover all rows to synchronize with the activity of scans already in progress. This may result in unpredictable changes in the row ordering returned by queries that have no ORDER BY clause.
- **off** indicates that the scan always starts from the table heading.

**Default value:** on

## enable\_beta\_features

**Parameter description:** Specifies whether to enable some features that are not officially released and are used only for POC verification. Exercise caution when enabling these extended features because they may cause errors in some scenarios.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the features are enabled for forward compatibility. Note that enabling them may cause errors in certain scenarios.
- **off** indicates that the features are disabled.

**Default value:** off

## default\_with\_oids

**Parameter description:** Specifies whether **CREATE TABLE** and **CREATE TABLE AS** include an **OID** field in newly-created tables if neither **WITH OIDS** nor **WITHOUT OIDS** is specified. It also determines whether OIDs will be included in tables created by **SELECT INTO**.

It is not recommended that OIDs be used in user tables. Therefore, this parameter is set to **off** by default. When OIDs are required for a particular table, **WITH OIDS** needs to be specified during the table creation.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that **CREATE TABLE** and **CREATE TABLE AS** can include an **OID** field in newly-created tables.
- **off** indicates that **CREATE TABLE** and **CREATE TABLE AS** cannot include any **OID** field in newly-created tables.

**Default value:** off

### 17.3.15.2 Platform and Client Compatibility

Many platforms use the database system. External compatibility of the database system provides a lot of convenience for platforms.

#### convert\_string\_to\_digit

**Parameter description:** Specifies the implicit conversion priority, which determines whether to preferentially convert strings into numbers.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that strings are preferentially converted into numbers.
- **off** indicates that strings are not preferentially converted into numbers.

**Default value:** on

---

#### NOTICE

Adjusting this parameter will change the internal data type conversion rule and cause unexpected behaviors. Exercise caution when performing this operation.

---

#### nls\_timestamp\_format

**Parameter description:** Specifies the default timestamp format.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** DD-Mon-YYYY HH:MI:SS.FF AM

#### max\_function\_args

**Parameter description:** Specifies the maximum number of parameters allowed for a function.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Value range:** an integer

**Default value:** 8192

#### transform\_null\_equals

**Parameter description:** Specifies whether expressions of the form `expr = NULL` (or `NULL = expr`) are treated as `expr IS NULL`. They return true if `expr` evaluates to the **NULL** value, and false otherwise.

- The correct SQL-standard-compliant behavior of **expr = NULL** is to always return **NULL** (unknown).

- Filtered forms in Microsoft Access generate queries that appear to use **expr = NULL** to test for null values. If you turn this option on, you can use this interface to access the database.

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that expressions of the form **expr = NULL** (or **NULL = expr**) are treated as **expr IS NULL**.
- **off** indicates that **expr = NULL** always returns **NULL** (unknown).

**Default value:** off

 **NOTE**

New users are always confused about the semantics of expressions involving **NULL** values. Therefore, **off** is used as the default value.

## support\_extended\_features

**Parameter description:** Specifies whether extended database features are supported.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that extended database features are supported.
- **off** indicates that extended database features are not supported.

**Default value:** off

## sql\_compatibility

**Parameter description:** Specifies the type of mainstream database with which the SQL syntax and statement behavior of the database is compatible. This parameter is an INTERNAL parameter. It can be viewed but cannot be modified.

**Value range:** enumerated type

- **A** indicates that the database is compatible with the Oracle database.
- **B** indicates that the database is compatible with the MySQL database.
- **C** indicates that the database is compatible with the Teradata database.
- **PG** indicates that the database is compatible with the PostgreSQL database.

**Default value:** A

---

**NOTICE**

- This parameter can be set only when you run the **CREATE DATABASE** command to create a database.
  - In the database, this parameter must be set to a specific value. It can be set to **A** or **B** and cannot be changed randomly. Otherwise, the setting is not consistent with the database behavior.
-



## behavior\_compat\_options

**Parameter description:** Specifies database compatibility behavior. Multiple items are separated by commas (,).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** ""

### NOTE

- Currently, only items in [Table 17-8](#) are supported.
- Multiple items are separated by commas (,), for example, **set behavior\_compat\_options='end\_month\_calculate,display\_leading\_zero';**

**Table 17-8** Compatibility configuration items

| Configuration Item   | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| display_leading_zero | <p>Specifies how floating point numbers are displayed. It controls the display of zeros before the decimal point of all character string types (such as char, character, nchar, varchar, character varying, varchar2, nvarchar2, text, and clob) and any precision types (such as float4, float8, and numeric) in the numeric type. In addition, the length of the calculated number is displayed synchronously.</p> <ul style="list-style-type: none"><li>If this item is not specified, for a decimal number between -1 and 1, the 0 before the decimal point is not displayed. For example:<br/>openGauss=# select 0.1231243 as a, 0.1231243::numeric as b,0.1231243::integer(10,3) as c, length(0.1242343) as d;<br/>a   b   c   d<br/>-----+-----+-----+---<br/>.1231243   .1231243   .123   8<br/>(1 row)</li><li>If this item is specified, for a decimal number between -1 and 1, the 0 before the decimal point is displayed. For example:<br/>openGauss=# select 0.1231243 as a, 0.1231243::numeric as b,0.1231243::integer(10,3) as c, length(0.1242343) as d;<br/>a   b   c   d<br/>-----+-----+-----+---<br/>0.1231243   0.1231243   0.123   9<br/>(1 row)</li></ul> |

| Configuration Item        | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| end_month_calculate       | <p>Specifies the calculation logic of the add_months function. Assume that the two parameters of the add_months function are <b>param1</b> and <b>param2</b>, and that the month of <b>param1</b> and <b>param2</b> is <b>result</b>.</p> <ul style="list-style-type: none"> <li>If this item is not specified, and the <b>Day</b> of <b>param1</b> indicates the last day of a month shorter than <b>result</b>, the <b>Day</b> in the calculation result will equal that in <b>param1</b>. For example:<br/> <pre>openGauss=# select add_months('2018-02-28',3) from sys_dummy; add_months ----- 2018-05-28 00:00:00 (1 row)</pre> </li> <li>If this item is specified, and the <b>Day</b> of <b>param1</b> indicates the last day of a month shorter than <b>result</b>, the <b>Day</b> in the calculation result will equal that in <b>result</b>. For example:<br/> <pre>openGauss=# select add_months('2018-02-28',3) from sys_dummy; add_months ----- 2018-05-31 00:00:00 (1 row)</pre> </li> </ul> |
| compat_analyze_sample     | <p>Specifies the sampling behavior of the <b>ANALYZE</b> operation. If this item is specified, the sample collected by the ANALYZE operation will be limited to around 30,000 records, controlling database node memory consumption and maintaining the stability of ANALYZE.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| bind_schema_tablespace    | <p>Binds a schema with the tablespace with the same name. If a tablespace name is the same as <i>sche_name</i>, <b>default_tablespace</b> will also be set to <i>sche_name</i> if <b>search_path</b> is set to <i>sche_name</i>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| bind_procedure_searchpath | <p>Specifies the search path of the database object in a stored procedure for which no schema name is specified. If no schema name is specified for a stored procedure, the search is performed in the schema to which the stored procedure belongs. If the stored procedure is not found, the following operations are performed:</p> <ul style="list-style-type: none"> <li>If this item is not specified, the system reports an error and exits.</li> <li>If this item is specified, the search continues based on the settings of <b>search_path</b>. If the issue persists, the system reports an error and exits.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                         |

| Configuration Item              | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| correct_to_number               | <p>Specifies the compatibility of the to_number() result.</p> <p>If this item is not set, the result of the to_number() function is the same as that in the A database by default.</p> <pre>openGauss=# select " AS to_number_14, to_number('34,50','999,99'); ERROR:  invalid data. CONTEXT:  referenced column: to_number</pre> <p>If this item is set, the result of the to_number() function is the same as that of pg11.</p> <pre>openGauss=# select " AS to_number_14, to_number('34,50','999,99'); to_number_14   to_number -----+-----                 3450 (1 row)</pre>                                                                       |
| unbind_divide_bound             | <p>Specifies the range check on the result of integer division.</p> <p>If this item is not set, the range of the division result is verified. For example, an out-of-bounds error is reported because the output result of <i>INT_MIN</i>/(-1) is greater than <i>INT_MAX</i>.</p> <pre>openGauss=# select (-2147483648)::int4 / (-1)::int4; ERROR:  integer out of range</pre> <p>If this item is set, the range of the division result does not need to be verified. For example, the output result of <i>INT_MIN</i>/(-1) is <i>INT_MAX</i>+1.</p> <pre>openGauss=# select (-2147483648)::int4 / (-1)::int4; ?column? ----- 2147483648 (1 row)</pre> |
| convert_string_digit_to_numeric | <p>Determines whether to convert numeric constants of the character string type to those of the numeric type before these two types are compared.</p> <pre>openGauss=# create table test1 (c1 int, c2 varchar); openGauss=# insert into test1 values (2, '1.1'); openGauss=# set behavior_compat_options=""; openGauss=# select * from test1 where c2 &gt; 1; ERROR:  invalid input syntax for type bigint: "1.1"</pre> <p>openGauss=# set behavior_compat_options='convert_string_digit_to_numeric';</p> <pre>openGauss=# select * from test1 where c2 &gt; 1; c1   c2 ---+--- 2   1.1 (1 row)</pre>                                                   |

| Configuration Item     | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| return_null_string     | <p>Specifies how to display the empty result (empty string '') of the <b>lpad()</b> and <b>rpadd()</b> functions.</p> <ul style="list-style-type: none"><li>If this item is not specified, the empty string is displayed as <b>NULL</b>.</li></ul> <pre>openGauss=# select length(lpad('123',0,'*')) from sys_dummy;<br/>length<br/>-----<br/>(1 row)</pre> <ul style="list-style-type: none"><li>If this item is specified, the empty string is displayed as '0'.</li></ul> <pre>openGauss=# select length(lpad('123',0,'*')) from sys_dummy;<br/>length<br/>-----<br/>0<br/>(1 row)</pre>                                                                                                                                                                                                                                                                                                                                                                                      |
| compat_concat_variadic | <p>Specifies the compatibility of variadic results of the <b>concat()</b> and <b>concat_ws()</b> functions. The B database does not have the variadic type. Therefore, this option has no impact on the B database.</p> <p>If this item is not set and the concat function parameter is of the variadic type, the results of the A and C databases in compatibility mode are the same by default.</p> <pre>openGauss=# select concat(variadic NULL::int[]) is NULL;<br/>?column?<br/>-----<br/>t<br/>(1 row)</pre> <p>If this item is set and the concat function parameter is of the variadic type, different result formats of the A and C databases in compatibility mode are retained.</p> <p>--In the A database:</p> <pre>openGauss=# select concat(variadic NULL::int[]) is NULL;<br/>?column?<br/>-----<br/>t<br/>(1 row)</pre> <p>--In the C database:</p> <pre>openGauss=# select concat(variadic NULL::int[]) is NULL;<br/>?column?<br/>-----<br/>f<br/>(1 row)</pre> |

| Configuration Item        | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| merge_update_multi        | <p>When MERGE INTO... When WHEN MATCHED THEN UPDATE (see "SQL Reference" &gt; "SQL Syntax" &gt; "MERGE INTO" in <i>Developer Guide</i>) and INSERT ... ON DUPLICATE KEY UPDATE (see "SQL Reference" &gt; "SQL Syntax" &gt; "INSERT" in <i>Developer Guide</i>) are used, it controls the UPDATE behavior if a piece of target data in the target table conflicts with multiple pieces of source data.</p> <p>If this item is specified and the preceding scenario exists, the system performs multiple UPDATE operations on the conflicting row. If this item is not specified and the preceding scenario exists, an error is reported, that is, the MERGE or INSERT operation fails.</p>                                                                                                                                             |
| plstmt_implicit_savepoint | <p>Determines whether the execution of an UPDATE statement in a stored procedure has an independent subtransaction.</p> <p>If this parameter is set, the implicit savepoint is enabled before executing each UPDATE statement in the stored procedure, and the subtransaction is rolled back to the latest savepoint in the EXCEPTION block by default, ensuring that only the modification of failed statements is rolled back. This option is used to be compatible with the <b>EXCEPTION</b> behavior of the O database.</p>                                                                                                                                                                                                                                                                                                       |
| hide_tailing_zero         | <p>Configuration item for numeric display. If this parameter is not set, numeric values are displayed based on the specified precision. If this parameter is set, all numeric values are output with trailing zeros (after a decimal point) hidden, including the scenario where the format precision is explicitly specified.</p> <p>For example:</p> <pre>openGauss=# set behavior_compat_options='hide_tailing_zero'; openGauss=# select cast(123.123 as numeric(15,10)) as a, to_char(cast(123.123 as numeric(15,10)), '999D999999');    a     to_char -----+----- 123.123   123.123 (1 row)  openGauss=# set behavior_compat_options=""; openGauss=# select cast(123.123 as numeric(15,10)) as a, to_char(cast(123.123 as numeric(15,10)), '999D999999');    a     to_char -----+----- 123.1230000000   123.123000 (1 row)</pre> |

| Configuration Item | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rownum_type_compat | <p>Specifies the ROWNUM type. The default value is <b>BIGINT</b>. After this parameter is specified, the value is changed to <b>NUMERIC</b>.</p> <pre>openGauss=# set behavior_compat_options=""; openGauss=# create table tb_test(c1 int,c2 varchar2,c3 varchar2); openGauss=# insert into tb_test values(1,'a','b'); openGauss=# create or replace view v_test as select rownum from tb_test; openGauss=# \d+ v_test       View "public.v_test"   Column   Type     Modifiers   Storage   Description -----+-----+-----+-----+-----  rownum   bigint              plain     View definition: SELECT ROWNUM AS "rownum" FROM tb_test;</pre> <pre>openGauss=# set behavior_compat_options = 'rownum_type_compat'; openGauss=# create or replace view v_test1 as select rownum from tb_test; openGauss=# \d+ v_test1       View "public.v_test1"   Column   Type     Modifiers   Storage   Description -----+-----+-----+-----+-----  rownum   numeric              main      View definition: SELECT ROWNUM AS "rownum" FROM tb_test;</pre>                                                                                                                                                                                                                                                                                                                                           |
| aformat_null_test  | <p>Specifies the logic for checking whether <b>rowtype</b> is not null. When this parameter is set, if one column in a row is not empty, <b>true</b> is returned for checking whether <b>rowtype</b> is not null. When this parameter is not set, if all columns in a row are not empty, true is returned for checking whether <b>rowtype</b> is not null. This parameter has no influence on checking whether <b>rowtype</b> is not null.</p> <pre>openGauss=# set behavior_compat_options='aformat_null_test'; openGauss=# select r, r is null as isnull, r is not null as isnotnull from (values (1,row(1,2)), (1,row(null,null)), (1,null), (null,row(1,2)), (null,row(null,null)), (null,null) ) r(a,b);       r        isnull   isnotnull -----+-----+----- (1,(1,2))   f        t (1,(,))     f        t (1,)        f        t (,(1,2))    f        t (,(,))      f        t (,)         t        f (6 rows)</pre> <pre>openGauss=# set behavior_compat_options=""; openGauss=# select r, r is null as isnull, r is not null as isnotnull from (values (1,row(1,2)), (1,row(null,null)), (1,null), (null,row(1,2)), (null,row(null,null)), (null,null) ) r(a,b);       r        isnull   isnotnull -----+-----+----- (1,(1,2))   f        t (1,(,))     f        t (1,)        f        f (,(1,2))    f        f (,(,))      f        f (,)         t        f (6 rows)</pre> |

| Configuration Item              | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| aformat_regexp_match            | <p>Determines the matching behavior of regular expression functions.</p> <p>When this parameter is set and <b>sql_compatibility</b> is set to <b>A</b> or <b>B</b>, the options supported by the <b>flags</b> parameter of the regular expression are changed as follows:</p> <ol style="list-style-type: none"><li>1. By default, the character '\n' cannot be matched.</li><li>2. If <b>flags</b> contains the <b>n</b> option, the period (.) can match the '\n' character.</li><li>3. The <b>regexp_replace(source, pattern replacement)</b> function replaces all matching substrings.</li><li>4. <b>regexp_replace(source, pattern, replacement, flags)</b> returns null when the value of <b>flags</b> is '' or null.</li></ol> <p>Otherwise, the meanings of the options supported by the <b>flags</b> parameter of the regular expression are as follows:</p> <ol style="list-style-type: none"><li>1. By default, the character '\n' can be matched.</li><li>2. The <b>n</b> option in <b>flags</b> indicates that the multi-line matching mode is used.</li><li>3. The <b>regexp_replace(source, pattern replacement)</b> function replaces only the first matched substring.</li><li>4. If the value of <b>flags</b> is '' or null, the return value of <b>regexp_replace(source, pattern, replacement, flags)</b> is the character string after replacement.</li></ol> |
| compat_cursor                   | <p>Determines the compatibility behavior of implicit cursor states. If this parameter is set and the O compatibility mode is used, the effective scope of implicit cursor states (<b>SQL %FOUND</b>, <b>SQL%NOTFOUND</b>, <b>SQL%ISOPEN</b>, and <b>SQL %ROWCOUNT</b>) are extended from only the currently executed function to all subfunctions invoked by this function.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| proc_outparam_override          | <p>Determines the overloading of output parameters of a stored procedure. After this parameter is enabled, the stored procedure can be properly created and invoked even if only the output parameters of the stored procedure are different. Currently, this parameter can be used only when gsql and JDBC are used to connect to the database. If this parameter is enabled for other tools to connect to the database, stored procedures with the output parameters cannot be invoked.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| proc_implicit_for_loop_variable | <p>Determines the behavior of the <b>FOR_LOOP</b> query statement in a stored procedure. When this parameter is set, if <b>rec</b> has been defined in the <b>FOR rec IN query LOOP</b> statement, the defined <b>rec</b> variable is not reused and a new variable is created. Otherwise, the defined <b>rec</b> variable is reused and no new variable is created.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

| Configuration Item            | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| allow_procedure_compile_check | Controls the compilation check of the <b>SELECT</b> and <b>OPEN CURSOR</b> statements in a stored procedure. If this parameter is set, when the <b>SELECT, OPEN CURSOR FOR, CURSOR %rowtype</b> , or <b>for rec in</b> statement is executed in a stored procedure, the stored procedure cannot be created if the queried table does not exist, and the compilation check of the trigger function is not supported. If the queried table exists, the stored procedure is successfully created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| char_coerce_compat            | <p>Determines the behavior when char(n) types are converted to other variable-length string types. If this parameter is not set, spaces at the end are omitted when the char(n) type is converted to other variable-length string types. After this parameter is set, spaces at the end are not omitted during conversion. In addition, if the length of the char(n) type exceeds the length of other variable-length string types, an error is reported. This parameter is valid only when the <b>sql_compatibility</b> parameter is set to <b>A</b>. After this parameter is enabled, spaces at the end are not omitted in implicit conversion, explicit conversion, or conversion by calling the <b>text(bpchar)</b> function.</p> <pre>openGauss=# set behavior_compat_options=""; openGauss=# create table tab_1(col1 varchar(3)); openGauss=# create table tab_2(col2 char(3)); openGauss=# insert into tab_2 values(' '); openGauss=# insert into tab_1 select col2 from tab_2; openGauss=# select * from tab_1 where col1 is null; col1 ----- (1 row) openGauss=# select * from tab_1 where col1=' '; col1 ----- (0 rows) openGauss=# delete from tab_1; openGauss=# set behavior_compat_options = 'char_coerce_compat'; openGauss=# insert into tab_1 select col2 from tab_2; openGauss=# select * from tab_1 where col1 is null; col1 ----- (0 rows) openGauss=# select * from tab_1 where col1=' '; col1 ----- (1 row)</pre> |



| Configuration Item         | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| truncate_numeric_tail_zero | <p>Configuration item for numeric display. If this parameter is not specified, numeric values are displayed based on the default precision. If this parameter is specified, the last 0 after the decimal point is hidden in all numeric output scenarios except <b>to_char(numeric, format)</b>. For example:</p> <pre>openGauss=#set behavior_compat_options='truncate_numeric_tail_zero'; openGauss=#select cast(123.123 as numeric(15,10)) as a, to_char(cast(123.123 as numeric(15,10)), '999D999999');  a   to_char -----+----- 123.123   123.123000 (1 row) openGauss=#set behavior_compat_options=""; openGauss=#select cast(123.123 as numeric(15,10)) as a, to_char(cast(123.123 as numeric(15,10)), '999D999999');  a   to_char -----+----- 123.1230000000   123.123000 (1 row)</pre> |
| plsql_security_definer     | After this parameter is enabled, the definer permission is used by default when a stored procedure is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| array_count_compat         | Controls the array.count function. If the parameter is enabled, the function returns <b>0</b> . Otherwise, the function returns <b>NULL</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| disable_emptystr2null      | If this parameter is enabled, the function of converting empty strings to null by default is disabled for the following character types: text, clob, blob, raw, bytea, varchar, nvarchar2, bpchar, char, name, byteawithoutorderwithqualcol, and byteawithoutordercol. This parameter is reserved for emergency. Do not set it unless necessary.                                                                                                                                                                                                                                                                                                                                                                                                                                                |

| Configuration Item         | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| proc_uncheck_default_param | <p>When a function is called, the system does not check whether the default parameter is omitted.</p> <ul style="list-style-type: none"> <li>If this item is not set and a function with default parameters is invoked, input parameters are added to the function from left to right. If inputs of non-default parameters are missing, an error is reported. For example:<br/> <pre> openGauss=# create or replace function test(f1 int, f2 int default 20, f3 int, f4 int default 40, f5 int default 50) return int openGauss-# as openGauss\$\$ begin openGauss\$\$ raise info 'f1:%',f1; openGauss\$\$ raise info 'f2:%',f2; openGauss\$\$ raise info 'f3:%',f3; openGauss\$\$ raise info 'f4:%',f4; openGauss\$\$ raise info 'f5:%',f5; openGauss\$\$ return 1; openGauss\$\$ end; openGauss\$\$ / CREATE FUNCTION openGauss=# select test(1,2); ERROR: function test(integer, integer) does not exist LINE 1: select test(1,2);                 ^ HINT: No function matches the given name and argument types. You might need to add explicit type casts. CONTEXT: referenced column: test </pre> </li> <li>If this item is set and a function with default parameters is invoked, input parameters are added to the function from left to right. The number of defaulted inputs depends on the number of default parameters. If an input of a non-default parameter is missing, the previous default value is used to fill this parameter. For example:<br/> <pre> openGauss=# create or replace function test(f1 int, f2 int default 20, f3 int, f4 int default 40, f5 int default 50) return int openGauss-# as openGauss\$\$ begin openGauss\$\$ raise info 'f1:%',f1; openGauss\$\$ raise info 'f2:%',f2; openGauss\$\$ raise info 'f3:%',f3; openGauss\$\$ raise info 'f4:%',f4; openGauss\$\$ raise info 'f5:%',f5; openGauss\$\$ return 1; openGauss\$\$ end; openGauss\$\$ / CREATE FUNCTION openGauss=# select test(1,2); INFO: f1:1 CONTEXT: referenced column: test INFO: f2:2 CONTEXT: referenced column: test INFO: f3:20 CONTEXT: referenced column: test INFO: f4:40 CONTEXT: referenced column: test INFO: f5:50 CONTEXT: referenced column: test test ----- </pre> </li> </ul> |

| Configuration Item       | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                          | <p>1<br/>(1 row)</p> <p>As shown above, f3 is filled with an incorrect default value.</p> <p><b>WARNING</b><br/>In this scenario, a non-default parameter is filled with the previous default value.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| plsql_rollback_keep_user | <p>Determines whether ROLLBACK and ROLLBACK TO SAVEPOINT in PL/SQL rolls back the current user. If this parameter is enabled, ROLLBACK in the PL/SQL does not change the current user.</p> <p>Example:</p> <pre>openGauss=# CREATE USER plsql_rollback1 password 'huawei@123'; openGauss=# CREATE USER plsql_rollback2 password 'huawei@123'; openGauss=# GRANT plsql_rollback1 to plsql_rollback2; openGauss=# CREATE OR REPLACE procedure plsql_rollback1.p1 () authid definer openGauss=# as openGauss\$\$ va int; openGauss\$\$ begin openGauss\$\$ raise info 'current usr:%', current_user; openGauss\$\$ rollback; openGauss\$\$ raise info 'current usr:%', current_user; openGauss\$\$ end; openGauss\$\$ / CREATE PROCEDURE openGauss=# SET session AUTHORIZATION plsql_rollback2 PASSWORD 'huawei@123'; SET openGauss=&gt; SET behavior_compat_options = 'plsql_rollback_keep_user'; SET openGauss=&gt; CALL plsql_rollback1.p1 (); INFO: current usr:plsql_rollback1 INFO: current usr:plsql_rollback1 p1 ----</pre> <p>(1 row)</p> <p><b>CAUTION</b><br/>This parameter is valid only in the A compatibility mode.</p> |

| Configuration Item | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| dynamic_sql_check  | <p>After the parameter is enabled, the number of template parameters in the template SQL statement is compared with the number of variables in the USING clause. If they are inconsistent, an error is reported.</p> <pre>openGauss=# create table tb_t1(col1 varchar2,col2 varchar2,col3 varchar2); CREATE TABLE openGauss=# insert into tb_t1 select '1','1','2'; INSERT 0 1 openGauss=# create or replace procedure proc_test()as openGauss\$# v_a varchar2; openGauss\$# v_b varchar2; openGauss\$# v_cnt int; openGauss\$# v_sql varchar2; openGauss\$# begin openGauss\$# v_a = '1'; openGauss\$# v_b = '2'; openGauss\$# v_sql = 'select count(1) from tb_t1 where col1 = :va and col2 = :va and col3 = :vb;'; openGauss\$# execute immediate v_sql into v_cnt using v_a,v_a,v_b; openGauss\$# raise info 'v_cnt:%',v_cnt; openGauss\$# end; openGauss\$# / CREATE PROCEDURE openGauss=# call proc_test(); INFO: v_cnt:0 proc_test ----- (1 row) openGauss=# set behavior_compat_options='dynamic_sql_check'; SET openGauss=# call proc_test(); ERROR: argnum not match in Dynamic SQL, using args num : 3 , actual sql args num : 2 CONTEXT: SQL statement "select count(1) from tb_t1 where col1 = :va and col2 = :va and col3 = :vb;" PL/SQL function proc_test() line 10 at EXECUTE statemen</pre> <p><b>CAUTION</b><br/>If the <b>dynamic_sql_compat</b> option is enabled, the <b>dynamic_sql_check</b> option does not take effect.</p> |

## plsql\_compile\_check\_options

**Parameter description:** Specifies database compatibility behavior. Multiple items are separated by commas (,).

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** ""

### NOTE

- Currently, only items in [Table 17-8](#) are supported.
- Multiple items are separated by commas (,), for example, **set plsql\_compile\_check\_options='for\_loop,outparam';**

**Table 17-9** Compatibility configuration items

| Configuration Item | Behavior                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| for_loop           | Determines the behavior of the <b>FOR_LOOP</b> query statement in a stored procedure. When this parameter is set, if <b>rec</b> has been defined in the <b>FOR rec IN query LOOP</b> statement, the defined <b>rec</b> variable is not reused and a new variable is created. Otherwise, the defined <b>rec</b> variable is reused and no new variable is created. (It is the same as <b>proc_implicit_for_loop_variable</b> and will be incorporated later.) |
| outparam           | When the output parameter overloading condition is met, the output parameters are checked. The output parameters cannot be constant.                                                                                                                                                                                                                                                                                                                         |

## a\_format\_version

**Parameter description:** Specifies the database platform compatibility configuration item. The value of this parameter is an enumerated string.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** ""

### NOTE

- Currently, only items in [Table 17-8](#) are supported.
- Set a character string for the compatibility configuration item, for example, **set a\_format\_version='10c'**.

**Table 17-10** Compatibility configuration items

| Configuration Item | Compatibility Behavior Control   |
|--------------------|----------------------------------|
| 10c                | Compatible version of platform A |

## a\_format\_dev\_version

**Parameter description:** Specifies the database platform minor version compatibility configuration item. The value of this parameter is an enumerated string.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** ""

 NOTE

- Currently, only items in [Table 17-8](#) are supported.
- Set a character string for the compatibility configuration item, for example, **set a\_format\_dev\_version='s1'**.

**Table 17-11** Compatibility configuration items

| Configuration Item | Compatibility Behavior Control                                                |
|--------------------|-------------------------------------------------------------------------------|
| s1                 | After this parameter is enabled, the int casted from text can be rounded off. |

## plpgsql.variable\_conflict

**Parameter description:** Sets the priority of using stored procedure variables and table columns with the same name.

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** a string

- **error** indicates that a compilation error is reported when the name of a stored procedure variable is the same as that of a table column.
- **use\_variable** indicates that if the name of a stored procedure variable is the same as that of a table column, the variable is used preferentially.
- **use\_column** indicates that if the name of a stored procedure variable is the same as that of a table column, the column name is used preferentially.

**Default value:** error

## td\_compatible\_truncation

**Parameter description:** Specifies whether to enable features compatible with a Teradata database. You can set this parameter to **on** when connecting to a database compatible with the Teradata database, so that when you perform the **INSERT** operation, overlong strings are truncated based on the allowed maximum length before being inserted into char- and varchar-type columns in the target table. This ensures all data is inserted into the target table without errors reported.

 NOTE

The string truncation function cannot be used if the **INSERT** statement includes a foreign table.

If inserting multi-byte character data (such as Chinese characters) to database with the character set byte encoding (such as SQL\_ASCII or LATIN1), and the character data crosses the truncation position, the string is truncated based on its bytes instead of characters. Unexpected result will occur in tail after the truncation. If you want correct truncation result, you are advised to adopt encoding set such as UTF8, which has no character data crossing the truncation position.

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that overlong strings are truncated.
- **off** indicates that overlong strings are not truncated.

**Default value:** off

## uppercase\_attribute\_name

**Parameter description:** Specifies whether to return column names in uppercase to the client. This parameter is used only in the ORA-compatible mode and centralized environment.

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that column names are returned to the client in uppercase.
- **off** indicates that column names are not returned to the client in uppercase.

**Default value:** off

## lastval\_supported

**Parameter description:** Specifies whether the lastval function can be used.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the lastval function can be used and the nextval function cannot be pushed down.
- **off** indicates that the lastval function cannot be used and the nextval function can be pushed down.

**Default value:** off

## 17.3.16 Fault Tolerance

This section describes parameters used for controlling how the server processes an error occurring in the database system.

### exit\_on\_error

**Parameter description:** If this function is enabled, errors of the ERROR level will be upgraded to PANIC errors, and core stacks will be generated. It is mainly used to locate problems and test services.

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that errors of the ERROR level will be upgraded to PANIC errors.
- **off** indicates that errors of the ERROR level will not be upgraded.

**Default value:** off

## restart\_after\_crash

**Parameter description:** If this parameter is set to **on** and a backend process crashes, GaussDB automatically reinitializes the backend process.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** maximizes the availability of the database.  
In some circumstances (for example, when a management tool, such as xCAT, is used to manage GaussDB), setting this parameter to **on** maximizes the availability of the database.
- **off** indicates that a management tool is enabled to obtain control permission and take proper measures when a backend process crashes.

**Default value:** on

## omit\_encoding\_error

**Parameter description:** If this parameter is set to **on** and the client character set of the database is encoded in UTF-8 format, character encoding conversion errors will be recorded in logs. Additionally, converted characters that have conversion errors will be ignored and replaced with question marks (?).

This is a USERSET parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that characters that have conversion errors will be ignored and replaced with question marks (?), and error information will be recorded in logs.
- **off** indicates that characters that have conversion errors cannot be converted and error information will be directly displayed.

**Default value:** off

## cn\_send\_buffer\_size

**Parameter description:** Specifies the size of the data buffer used for data transmission on the primary database node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 8 to 128. The unit is KB.

**Default value:** 8 KB

## data\_sync\_retry

**Parameter description:** Specifies whether to keep running the database when updated data fails to be written into disks by using the **fsync** function. In some OSs, no error is reported even if **fsync** fails after the second attempt. As a result, data is lost.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).



**Value range:** Boolean

- **on** indicates that the database keeps running and **fsync** is executed again after **fsync** fails.
- **off** indicates that a PANIC-level error is reported and the database is stopped after **fsync** fails.

**Default value:** off

## remote\_read\_mode

**Parameter description:** Specifies whether to enable the remote read function.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **off** indicates that the remote read function is disabled.
- **non\_authentication** indicates that the remote read function is enabled but certificate authentication is not required.
- **authentication** indicates that the remote read function is enabled and certificate authentication is required.

**Default value:** authentication

## 17.3.17 Connection Pool Parameters

When a connection pool is used to access the database, database connections are established and then stored in the memory as objects during system running. When you need to access the database, no new connection is established. Instead, an existing idle connection is selected from the connection pool. After you finish accessing the database, the database does not disable the connection but puts it back into the connection pool. The connection can be used for the next access request.

## cache\_connection

**Parameter description:** Specifies whether to recycle the connections of a connection pool.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that the connections of a connection pool will be recycled.
- **off** indicates that the connections of a connection pool will not be recycled.

**Default value:** on

## 17.3.18 Transaction

This section describes the settings and value ranges of database transaction parameters.

## transaction\_isolation

**Parameter description:** specifies the isolation level of the current transaction.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string of case-sensitive characters. The values include:

- **serializable:** This value is equivalent to REPEATABLE READ in GaussDB.
- **read committed** indicates that only the data in committed transactions will be read.
- **repeatable read** indicates that only the data committed before transaction start is read. Uncommitted data or data committed in other concurrent transactions cannot be read.
- **default:** The value is the same as that of **default\_transaction\_isolation**.

**Default value:** read committed

## transaction\_read\_only

**Parameter description:** Specifies whether the current transaction is a read-only transaction.

This parameter has a fixed value **on** during database restoration or on the standby node. Otherwise, set this parameter to the value of **default\_transaction\_read\_only**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the current transaction is a read-only transaction.
- **off** indicates that the current transaction can be a read/write transaction.

**Default value:** off

## xc\_maintenance\_mode

**Parameter description:** Specifies whether the system is in maintenance mode.

This is a SUSERSET parameter. Set it based on method 3 in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the system is in maintenance mode.
- **off** indicates that the system is not in maintenance mode.

---

### NOTICE

Enable the maintenance mode with caution to avoid database data inconsistencies.

---

**Default value:** off

## allow\_concurrent\_tuple\_update

**Parameter description:** Specifies whether to allow concurrent update.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the concurrent update is allowed.
- **off** indicates that the system is not in maintenance mode.

**Default value:** on

## transaction\_deferrable

**Parameter description:** Specifies whether to delay the execution of a read-only serial transaction without incurring an execution failure. Assume this parameter is set to **on**. When the server detects that the tuples read by a read-only transaction are being modified by other transactions, it delays the execution of the read-only transaction until the other transactions finish modifying the tuples. This parameter is reserved and does not take effect in this version. Similar to this parameter, the [default\\_transaction\\_deferrable](#) parameter is used to specify whether to allow delayed execution of a transaction.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the execution of a read-only serial transaction can be delayed.
- **off** indicates that the execution of a read-only serial transaction cannot be delayed.

**Default value:** off

## enable\_show\_any\_tuples

**Parameter description:** This parameter is available only in a read-only transaction and is used for analysis. When this parameter is set to **on** or **true**, all versions of tuples in the table are displayed.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** or **true** indicates that all versions of tuples in the table are displayed.
- **off** or **false** indicates that no versions of tuples in the table are displayed.

**Default value:** off

## replication\_type

**Parameter description:** Specifies whether the current HA database is deployed in standalone or one-primary-multiple-standby mode.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

This parameter is an internal parameter. Do not set it.

**Value range:** 0 to 2

- **2** indicates the standalone mode. In this mode, no standby node can be added.
- **1** indicates the one-primary-multiple-standby mode, covering all scenarios. This mode is recommended.
- **0** indicates that the mode is not available currently.

**Default value:** 1

## pgxc\_node\_name

**Parameter description:** Specifies the name of a node.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

When a standby node requests to replicate logs on the primary node, if the **application\_name** parameter is not set, the **pgxc\_node\_name** parameter is used as the name of the streaming replication slot of the standby node on the primary node. The streaming replication slot is named in the following format: Value of this parameter\_IP address of the standby node\_Port number of the standby node. The IP address and port number of the standby node are obtained from the IP address and port number of the standby node specified by the **replconninfo** parameter. The maximum length of a streaming replication slot name is 61 characters. If the length of the concatenated string exceeds 61 characters, the truncated **pgxc\_node\_name** will be used for concatenation to ensure that the length of the streaming replication slot name is less than or equal to 61 characters.

---

### CAUTION

After this parameter is modified, the database instance will fail to be connected. You are advised not to modify this parameter.

---

**Value range:** a string

**Default value:** current node name

## enable\_defer\_calculate\_snapshot

**Parameter description:** Specifies the delay in calculating **xmin** and **oldestxmin**. Calculation is triggered only when 1000 transactions are executed or the interval is 1s. If this parameter is set to **on**, the overhead of calculating snapshots can be reduced in heavy-load scenarios, but the progress of **oldestxmin** is slow, affecting tuple recycling. If this parameter is set to **off**, **xmin** and **oldestxmin** can be calculated in real time, but the overhead for calculating snapshots increases.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that snapshots **xmin** and **oldestxmin** are calculated with a delay.
- **off** indicates that snapshot **xmin** and **oldestxmin** are calculated in real time.

**Default value:** on

## 17.3.19 Replication Parameters of Two Database Instances

### RepOriginId

**Parameter description:** This parameter is a session-level GUC parameter. In bidirectional logical replication, set it to a non-zero value to avoid infinite data replication.

This is a USERSET parameter. Set it based on method 3 in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647

**Default value:** 0

### stream\_cluster\_run\_mode

**Parameter description:** Specifies whether a DN is in the primary or standby instance in dual-instance streaming DR scenarios. For single-instance scenarios, the DN is in the primary instance by default.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **cluster\_primary** indicates that the node is in the primary instance.
- **cluster\_standby** indicates that the node is in the standby instance.

**Default value:** cluster\_primary

### enable\_roach\_standby\_cluster

**Parameter description:** Sets the standby database instances to read-only in dual-database instance mode. Only users with the sysadmin permission can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the read-only mode is enabled for the standby database instances.
- **off** indicates that the read-only mode is disabled for the standby database instances. In this case, the standby database instances can be read and written.

**Default value:** off

### hadr\_process\_type

**Parameter description:** Specifies the process type in dual-instance streaming DR scenarios.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** enumerated values

- **none** indicates that there is no process.
- **failover** indicates that the DR database instance is being upgraded.
- **switchover\_promote** indicates the process of upgrading the DR database instance to the primary database instance during the primary/standby database instance switchover.
- **switchover\_demote** indicates the process of demoting the primary database instance to the DR database instance during the primary/standby database instance switchover.

**Default value:** none

## 17.3.20 Developer Options

### allow\_system\_table\_mods

**Parameter description:** Specifies whether the structure of a system catalog or the name of a system schema can be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the structure of the system catalog or the name of the system schema can be modified.
- **off** indicates that the structure of the system catalog or the name of the system schema cannot be modified.

**Default value:** off

#### NOTE

You are not advised to change the default value of this parameter. If this parameter is set to **on**, system tables may be damaged and the database may fail to be started.

### allow\_create\_sysobject

**Parameter description:** Specifies whether objects such as functions, stored procedures, synonyms, aggregate functions, and operators can be created or modified in the system schema. The system schema refers to the schema provided by the database after initialization, excluding the public schema. The OID of the system schema is usually smaller than 16384.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the initial user and system administrator can create or modify objects such as functions, stored procedures, synonyms, and aggregate functions in the system schema, and the initial user can create operators in the system schema. For details about whether other users are allowed to create these objects, see the permission requirements of the corresponding schema.
- **off** indicates that all users are not allowed to create or modify objects such as functions, stored procedures, synonyms, aggregate functions, and operators in the system schema.

**Default value:** on

## debug\_assertions

**Parameter description:** Specifies whether to enable various assertion checks. This parameter assists in debugging. If you are experiencing strange problems or crashes, set this parameter to **on** to identify programming defects. To use this parameter, the macro USE\_ASSERT\_CHECKING must be defined (through the configure option **--enable-cassert**) during the GaussDB compilation.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that various assertion checks are enabled.
- **off** indicates that various assertion checks are disabled.

### NOTE

If you compile GaussDB with the assertion check enabled, **debug\_assertions** is set to **on** by default.

**Default value:** off

## ignore\_checksum\_failure

**Parameter description:** If this parameter is enabled, the system ignores the failure (but still generates an alarm). Continuing execution may result in breakdown, damaged data being transferred or hidden, failure of data recovery from remote nodes, or other serious problems. You are not advised to modify the settings.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that data check errors are ignored.
- **off** indicates that data check errors are reported.

**Default value:** off

## ignore\_system\_indexes

**Parameter description:** Specifies whether to ignore system indexes when reading system catalog (but still update the indexes when modifying the tables).

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).

---

### NOTICE

This parameter is useful for recovering data from tables whose system indexes are damaged.

---

**Value range:** Boolean

- **on** indicates that system indexes are ignored.
- **off** indicates that system indexes are not ignored.

**Default value:** off

## post\_auth\_delay

**Parameter description:** Specifies the delay in the connection to the server after a successful authentication. Developers can attach a debugger to the server startup process.

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147. The unit is s.

**Default value:** 0

### NOTE

This parameter is used only for commissioning and fault locating. To prevent impact on service running, ensure that the default value **0** is used in the production environment. If this parameter is set to a value other than 0, the cluster status may be abnormal due to a long authentication delay.

## pre\_auth\_delay

**Parameter description:** Specifies the period of delaying authentication after the connection to the server is started. Developers can attach a debugger to the authentication procedure.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 60. The unit is s.

**Default value:** 0

### NOTE

This parameter is used only for commissioning and fault locating. To prevent impact on service running, ensure that the default value **0** is used in the production environment. If this parameter is set to a value other than 0, the cluster status may be abnormal due to a long authentication delay.

## trace\_notify

**Parameter description:** Specifies whether to enable the function of generating debugging output for the **LISTEN** and **NOTIFY** commands. The level of [client\\_min\\_messages](#) or [log\\_min\\_messages](#) must be **debug1** or lower so that debugging output can be recorded in the client or server logs, respectively.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the function is enabled.
- **off** indicates that the function is disabled.

**Default value:** off



## trace\_recovery\_messages

**Parameter description:** Specifies whether to enable logging of recovery-related debugging output. This parameter allows users to overwrite the normal setting of [log\\_min\\_messages](#), but only for specific messages. This is intended for the use in debugging the standby node.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values. Valid values include **debug5**, **debug4**, **debug3**, **debug2**, **debug1**, and **log**. For details about the parameter values, see [log\\_min\\_messages](#).

**Default value:** log

### NOTE

- **log** indicates that recovery-related debugging information will not be logged.
- Except the default value **log**, each of the other values indicates that recovery-related debugging information at the specified level will also be logged. Common settings of **log\_min\_messages** enable logs to be unconditionally recorded into server logs.

## trace\_sort

**Parameter description:** Specifies whether to print information about resource usage during sorting operations. This parameter is available only when the macro **trace\_sort** is defined during the GaussDB compilation. However, **trace\_sort** is currently defined by default.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the function is enabled.
- **off** indicates that the function is disabled.

**Default value:** off

## zero\_damaged\_pages

**Parameter description:** Specifies whether to detect a damaged page header that causes GaussDB to report an error, aborting the current transaction.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

Setting this parameter to **on** causes the system to report a warning, zero out the damaged page, and continue processing. This behavior will destroy data, including all the rows on the damaged page. However, it allows you to bypass the error and retrieve rows from any undamaged pages that may be present in the table. Therefore, it is useful for restoring data if corruption has occurred due to a hardware or software error. In most cases, you are advised not to set this parameter to **on** if you want to restore data from damaged pages.

**Default value:** off

## remotetype

**Parameter description:** Specifies the remote connection type.

This is a BACKEND parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values. Valid values are **application**, **datanode**, and **internaltool**.

**Default value:** **application**

## max\_user\_defined\_exception

**Parameter description:** Specifies the maximum number of exceptions. The default value cannot be changed.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. Currently, only the fixed value **1000** is supported.

**Default value:** **1000**

## enable\_fast\_numeric

**Parameter description:** Specifies whether to enable optimization for numeric data calculation. Calculation of numeric data is time-consuming. Numeric data is converted into int64- or int128-type data to improve numeric data calculation performance.

This is a SUSER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** or **true** indicates that optimization for numeric data calculation is enabled.
- **off** or **false** indicates that optimization for numeric data calculation is disabled.

**Default value:** **on**

## enable\_compress\_spill

**Parameter description:** Specifies whether to enable the compression function of writing data to disk.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** or **true** indicates that optimization for writing data to disk is enabled.
- **off** or **false** indicates that optimization for writing data to a disk is disabled.

**Default value:** **on**

## support\_batch\_bind

**Parameter description:** Specifies whether to batch bind and execute PBE statements through interfaces such as JDBC, ODBC, and libpq.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that batch binding and execution are used.
- **off** indicates that batch binding and execution are not used.

**Default value:** on

## numa\_distribute\_mode

**Parameter description:** Specifies the distribution of some shared data and threads among NUMA nodes. This parameter is used to optimize the performance of large-scale Arm servers with multiple NUMA nodes. Generally, you do not need to set this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. The valid values are **none** and **all**.

- **none** indicates that this function is disabled.
- **all** indicates that some shared data and threads are distributed to different NUMA nodes to reduce the number of remote access times and improve performance. Currently, this function applies only to Arm servers with multiple NUMA nodes. All NUMA nodes must be available for database processes. You cannot select only some NUMA nodes.

### NOTE

In the current version, **numa\_distribute\_mode** cannot be set to **all** on the x86 architecture.

**Default value:** none

## log\_pagewriter

**Parameter description:** Specifies whether to display the page refresh information of a thread and details about an incremental check point after the incremental check point is enabled. You are not advised to set this parameter to **true** because a large amount of information will be generated.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** on

## advance\_xlog\_file\_num

**Parameter description:** Specifies the number of Xlog files that are periodically initialized in advance in the backend. This parameter is used to prevent the Xlog file initialization from affecting the performance during transaction commission. However, such a fault may occur only when the system is overloaded. Therefore, you do not need to set this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000. The value **0** indicates that initialization is not performed in advance. For example, the value **10** indicates that

the backend thread periodically initializes 10 Xlog files in advance based on the write location of the current Xlog.

**Default value:** 0

## enable\_beta\_opfusion

**Parameter description:** Specifies whether to accelerate the execution of SQL statements, such as aggregate functions and sorting in TPC-C when **enable\_opfusion** is set to **on**.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** enabled.
- **off:** disabled.

**Default value:** off

## enable\_csqual\_pushdown

**Parameter description:** Specifies whether to deliver filter criteria for a rough check during query.

This is a SUSERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that a rough check is performed with filter criteria delivered during query.
- **off** indicates that a rough check is performed without filter criteria delivered during query.

**Default value:** on

## string\_hash\_compatible

**Parameter description:** Specifies whether to use the same method to calculate char-type hash values and varchar- or text-type hash values. Based on the setting of this parameter, you can determine whether a redistribution is required when a distribution key is converted from a char-type data distribution into a varchar- or text-type data distribution.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the same calculation method is used and a redistribution is not required.
- **off** indicates that different calculation methods are used and a redistribution is required.

 NOTE

Calculation methods differ in the length of input strings used for calculating hash values. (For a char-type hash value, spaces following a string are not counted as the length. For a text- or varchar-type hash value, the spaces are counted.) The hash value affects the calculation result of queries. To avoid query errors, do not modify this parameter during database running once it is set.

**Default value:** off

## pldebugger\_timeout

**Parameter description:** Specifies the timeout interval for the pldebugger server to wait for a response from the debug end.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 86400. The unit is second.

**Default value:** 15min

## plsql\_show\_all\_error

**Parameter description:** Specifies whether to skip errors and continue compiling PL/SQL objects. For details about the impact, see the description of DBE\_PLDEVELOPER.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** off

## ustore\_attr

**Parameter description:** This parameter is used to control the information statistics of Ustore tables, rollback type, and data verification during the running of key modules (including data, indexes, rollback segments, and replay). This parameter helps R&D engineers locate faults.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string. This parameter is set in key-value mode. The mapping between keys and values is as follows: If multiple key-value pairs are used, use semicolons (;) to separate them. For example:

**ustore\_attr='ustore\_verify\_level=NORMAL;ustore\_verify\_module=UPAGE:UBTR  
EE:UNDO:REDO'.**

 NOTE

When setting **ustore\_attr**, do not leave spaces or other characters before and after the equal sign (=) between key and value, for example, **ustore\_attr='ustore\_verify\_level = NORMAL**. Otherwise, the parameter is invalid during kernel code verification and the parameter setting fails.

- **ustore\_verify\_level:** verification level.

**Value range:** a string. For details, see the following table.

**Table 17-12** Parameter value meaning of `ustore_verify_level`

| Parameter Value | Description                                                                                                                                                                                      |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FAST            | <b>FAST</b> indicates fast verification. A few content is verified and the impact on performance is minimized.                                                                                   |
| NORMAL          | <b>NORMAL</b> indicates that the verification is routinely performed. Compared with the fast verification, this verification involves more content, and the impact on the performance is medium. |
| SLOW            | <b>SLOW</b> indicates slow verification. The verification content is the largest, and the performance is greatly affected.                                                                       |

**Default value:** an empty string

- **ustore\_verify\_module:** module that controls verification.

**Value range:** a string of case-insensitive characters. The value can be one or more of **UPAGE**, **UBTREE**, **UNDO**, and **REDO**, or it can be **ALL** or **NULL**. When multiple values of **UPAGE**, **UBTREE**, **UNDO**, and **REDO** are used, separate them with colons (:).

For example, **ustore\_verify\_module=UPAGE:UBTREE:UNDO:REDO**.

**Table 17-13** Parameter value meaning of `ustore_verify_module`

| Parameter Value | Description                                                                                 |
|-----------------|---------------------------------------------------------------------------------------------|
| UPAGE           | Indicates that data page verification is enabled.                                           |
| UBTREE          | Indicates that UBtree index verification is enabled.                                        |
| UNDO            | Indicates that rollback segment data verification is enabled.                               |
| REDO            | Indicates that data page verification for the REDO process is enabled.                      |
| ALL             | Indicates that data verification is enabled for the UPAGE, UBTREE, UNDO, and REDO modules.  |
| NULL            | Indicates that data verification for the UPAGE, UBTREE, UNDO, and REDO modules is disabled. |

**Default value:** an empty string

- **index\_trace\_level:** specifies whether to enable index tracing and control the printing level. After this function is enabled, information about index tuples that meet the conditions is printed based on the printing level during index scanning.

**Value range:** a string. The values are described in the following table.

**Default value: NO**

**Table 17-14** Parameter value meaning of index\_trace\_level

| Parameter Value | Description                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NO              | No additional information is printed.                                                                                                                                                                                                                                                                                                                                                                                                          |
| NORMAL          | Information about visible index tuples are printed, including: <ul style="list-style-type: none"><li>• ID and offset of the index page where the current index tuple is located</li><li>• Current tuple status</li><li>• TID and partOid corresponding to the current tuple</li><li>• xmin and xmax information corresponding to the current tuple</li><li>• Current tuple content (if <b>enable_log_tuple</b> is set to <b>on</b>).</li></ul> |
| VISIBILITY      | On the basis of <b>NORMAL</b> , the information about the index tuples that do not pass the visibility check is printed and whether the index tuples are visible is marked.                                                                                                                                                                                                                                                                    |
| SHOWHIKEY       | On the basis of <b>VISIBILITY</b> , the system tries to print the information about the HIKEY tuple on the page.                                                                                                                                                                                                                                                                                                                               |
| ALL             | Information about all tuples on the scanned index page is printed.                                                                                                                                                                                                                                                                                                                                                                             |

- **enable\_log\_tuple**: specifies whether to print the contents of related tuples when printing log-level prompts for troubleshooting and locating.

**Value range:** on or off (case-insensitive)

**Default value:** off

Note: This parameter has been discarded.

- **enable\_ustore\_sync\_rollback**: indicates whether to enable synchronous rollback for Ustore tables.

**Value range:** Boolean

**Default value:** true

- **enable\_ustore\_async\_rollback**: indicates whether to enable asynchronous rollback for Ustore tables.

**Value range:** Boolean

**Default value:** true

- **enable\_ustore\_page\_rollback**: indicates whether to enable page rollback for Ustore tables.

**Value range:** Boolean

**Default value:** true

- **enable\_ustore\_partial\_seqscan**: indicates whether to enable partial scan for Ustore tables.

**Value range:** Boolean

**Default value:** false

- **enable\_candidate\_buf\_usage\_count:** indicates whether to enable buffer usage statistics.

**Value range:** Boolean

**Default value:** false

- **ustats\_tracker\_naptime:** specifies the interval for collecting statistics on Ustore tables.

**Value range:** [1,INT\_MAX/1000]

**Default value:** 20, in seconds

- **umax\_search\_length\_for\_prune:** specifies the maximum search depth of the prune operation on the Ustore table.

**Value range:** [1,INT\_MAX/1000]

**Default value:** 10 (times)

**Default value:** an empty string

 **NOTE**

This parameter applies only to Ustore tables in a centralized system.

## 17.3.21 Auditing

### 17.3.21.1 Audit Switch

#### **audit\_enabled**

**Parameter description:** Specifies whether to enable or disable the audit process. After the audit process is enabled, the auditing information written by the background process can be read from the pipe and written into audit files.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the auditing function is enabled.
- **off** indicates that the auditing function is disabled.

**Default value:** on

#### **audit\_directory**

**Parameter description:** Specifies the storage directory of audit files. A path relative to the **data** directory. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** **pg\_audit** If **om** is used for database deployment, audit logs are stored in *\$GAUSSLOG/pg\_audit/Instance name*.



---

**NOTICE**

- You need to set different audit file directories for different DNs. Otherwise, audit logs will be abnormal.
  - If the value of **audit\_directory** in the configuration file is an invalid path, the audit function cannot be used.
- 

 **NOTE**

- Valid path: Users must have read and write permissions on the path.
- Invalid path: Users do not have read or write permissions on an invalid path.

## audit\_data\_format

**Parameter description:** Audits the format of log files. Currently, only the binary format is supported. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** binary

## audit\_rotation\_interval

**Parameter description:** Specifies the interval of creating an audit log file. If the difference between the current time and the time when the previous audit log file is created is greater than the value of this parameter, a new audit log file will be generated.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to *INT\_MAX*/60. The unit is min.

**Default value:** 1d

---

**NOTICE**

Do not adjust this parameter unless necessary. Otherwise, **audit\_resource\_policy** may fail to take effect. To control the storage space and time of audit logs, set the [audit\\_resource\\_policy](#), [audit\\_space\\_limit](#), and [audit\\_file\\_remain\\_time](#) parameters.

---

## audit\_rotation\_size

**Parameter description:** Specifies the maximum capacity of an audit log file. If the total number of messages in an audit log exceeds the value of this parameter, the server will generate a new audit log file.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1024 to 1048576. The unit is KB.

**Default value:** 10 MB

---

**NOTICE**

- Do not adjust this parameter unless necessary. Otherwise, **audit\_resource\_policy** may fail to take effect. To control the storage space and time of audit logs, set the **audit\_resource\_policy**, **audit\_space\_limit**, and **audit\_file\_remain\_time** parameters.
  - If the space occupied by a single record in an audit log file exceeds the value of this parameter, the log file is regarded as an invalid log file.
- 

## audit\_resource\_policy

**Parameter description:** Specifies the policy for determining whether audit logs are preferentially stored by space or time.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that audit logs are preferentially stored by space. A maximum of [audit\\_space\\_limit](#) logs can be stored.
- **off** indicates that audit logs are preferentially stored by time. A minimum duration of [audit\\_file\\_remain\\_time](#) logs must be stored.

**Default value:** on

## audit\_file\_remain\_time

**Parameter description:** Specifies the minimum duration required for recording audit logs. This parameter is valid only when [audit\\_resource\\_policy](#) is set to **off**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 730. The unit is day. **0** indicates that the storage duration is not limited.

**Default value:** 90

## audit\_space\_limit

**Parameter description:** Specifies the total disk space occupied by audit files.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1024 KB to 1024 GB. The unit is KB.

**Default value:** 1GB

---

**NOTICE**

In the multi-audit thread scenario, the minimum disk space occupied by audit files is the product of values of **audit\_thread\_num** and **audit\_rotation\_size**. Ensure that the value of **audit\_space\_limit** is greater than the product of values of **audit\_thread\_num** and **audit\_rotation\_size**.

---

## audit\_file\_remain\_threshold

**Parameter description:** Specifies the maximum number of audit files in the audit directory.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 100 to 1048576

**Default value:** 1048576

---

### NOTICE

- Ensure that this parameter is set to **1048576**. Do not adjust this parameter unless necessary. Otherwise, **audit\_resource\_policy** may fail to take effect. To control the storage space and time of audit logs, set the **audit\_resource\_policy**, **audit\_space\_limit**, and **audit\_file\_remain\_time** parameters.
  - In the multi-audit thread scenario, do not adjust this parameter unless necessary. Ensure that the value of this parameter is equal to or greater than the value of **audit\_thread\_num**. Otherwise, the audit function cannot be used and the database is abnormal.
- 

## audit\_thread\_num

**Parameter description:** Specifies the number of audit threads.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 48

**Default value:** 1

---

### NOTICE

- When **audit\_dml\_state** is enabled and high performance is required, you are advised to increase the value of this parameter to ensure that audit messages can be processed and recorded in a timely manner.
  - Ensure that the value of this parameter is less than or equal to the maximum number of audit files in the audit directory (**audit\_file\_remain\_threshold**). Otherwise, the audit function cannot be used and the database is abnormal.
- 

## 17.3.21.2 User and Permission Audit

### audit\_login\_logout

**Parameter description:** Specifies whether to audit the GaussDB user's login (including login success and failure) and logout.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 7

- **0** indicates that the function of auditing users' logins and logouts is disabled.
- **1** indicates that only successful user logins are audited.
- **2** indicates that only failed user logins are audited.
- **3** indicates that successful and failed user logins are audited.
- **4** indicates that only user logouts are audited.
- **5** indicates that successful user logouts and logins are audited.
- **6** indicates that failed user logouts and logins are audited.
- **7** indicates that successful user logins, failed user logins, and logouts are audited.

**Default value:** 7

## audit\_database\_process

**Parameter description:** Specifies whether to audit GaussDB start, stop, recovery, and switchover operations.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing GaussDB start, stop, recovery, and switchover operations is disabled.
- **1** indicates that the function of auditing GaussDB start, stop, recovery, and switchover operations is enabled.

**Default value:** 1

---

### NOTICE

When the database is started, the standby DN is promoted to primary. Therefore, the DN type in the audit log is **system\_switch** when the DN is started.

---

## audit\_user\_locked

**Parameter description:** Specifies whether to audit the GaussDB user's locking and unlocking.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing user's locking and unlocking is disabled.
- **1** indicates that the function of auditing user's locking and unlocking is enabled.

**Default value:** 1

## audit\_user\_violation

**Parameter description:** Specifies whether to audit the access violation operations of a user.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing the access violation operations of a user is disabled.
- **1** indicates that the function of auditing the access violation operations of a user is enabled.

**Default value:** 0

## audit\_grant\_revoke

**Parameter description:** Specifies whether to audit the granting and reclaiming of the GaussDB user's permission.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing the granting and reclaiming of a user's permission is disabled.
- **1** indicates that the function of auditing the granting and reclaiming of a user's permission is enabled.

**Default value:** 1

### 17.3.21.3 Operation Audit

## audit\_system\_object

**Parameter description:** Specifies whether to audit the CREATE, DROP, and ALTER operations on the GaussDB database object. The GaussDB database objects include databases, users, schemas, and tables. The operations on the database object can be audited by changing the value of this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 134217727

- **0** indicates that the function of auditing the CREATE, DROP, and ALTER operations on the GaussDB database object can be disabled.
- Other values indicate that the CREATE, DROP, and ALTER operations on a certain or some GaussDB database objects are audited.

**Value description:**

The value of this parameter is calculated by 27 binary bits. The 27 binary bits represent 27 types of GaussDB objects. If the corresponding binary bit is set to **0**, the CREATE, DROP, and ALTER operations on corresponding database objects are not audited. If it is set to **1**, the CREATE, DROP, and ALTER operations are audited.

For details about the audit contents represented by these 27 binary bits, see [Table 17-15](#).

When SQL patches are audited and **audit\_dml\_state\_select** are enabled, an SQL patch operation will be audited twice and recorded as DML and DDL operations in the audit log, respectively.

**Default value: 67121159**

**Table 17-15** Meaning of each value for the **audit\_system\_object** parameter

| Binary Bit | Description                                                                  | Value Description                                                                                                                                                                                                                                                             |
|------------|------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit 0      | Whether to audit the CREATE, DROP, and ALTER operations on databases.        | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li></ul>                     |
| Bit 1      | Whether to audit the CREATE, DROP, and ALTER operations on schemas.          | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li></ul>                     |
| Bit 2      | Whether to audit the CREATE, DROP, and ALTER operations on users.            | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li></ul>                     |
| Bit 3      | Whether to audit the CREATE, DROP, ALTER, and TRUNCATE operations on tables. | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, ALTER, and TRUNCATE operations on these objects are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, ALTER, and TRUNCATE operations on these objects are audited.</li></ul> |
| Bit 4      | Whether to audit the CREATE, DROP, and ALTER operations on indexes.          | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li></ul>                     |

| Binary Bit | Description                                                                                                                                                                 | Value Description                                                                                                                                                                                                                                            |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit 5      | Whether to audit the CREATE and DROP operations on VIEW and MATVIEW objects.                                                                                                | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE and DROP operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE and DROP operations on these objects are audited.</li> </ul>                 |
| Bit 6      | Whether to audit the CREATE, DROP, and ALTER operations on triggers.                                                                                                        | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul> |
| Bit 7      | Whether to audit the CREATE, DROP, and ALTER operations on procedures/ functions.                                                                                           | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul> |
| Bit 8      | Whether to audit the CREATE, DROP, and ALTER operations on tablespaces.                                                                                                     | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul> |
| Bit 9      | Whether to audit the CREATE, DROP, and ALTER operations on resource pools. For details, see section "Resource Load Management" in <i>Lab Feature Description</i> .          | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul> |
| Bit 10     | Whether to audit the CREATE, DROP, and ALTER operations on the WORKLOAD object. For details, see section "Resource Load Management" in the <i>Lab Feature Description</i> . | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul> |
| Bit 11     | Reserved.                                                                                                                                                                   | -                                                                                                                                                                                                                                                            |

| Binary Bit | Description                                                                                                    | Value Description                                                                                                                                                                                                                                                        |
|------------|----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit 12     | Whether to audit the CREATE, DROP, and ALTER operations on data sources. You are advised not to use it.        | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul>             |
| Bit 13     | Reserved.                                                                                                      | -                                                                                                                                                                                                                                                                        |
| Bit 14     | Whether to audit the CREATE, DROP, and ALTER operations on ROW LEVEL SECURITY objects.                         | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on these objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on these objects are audited.</li> </ul>             |
| Bit 15     | Whether to audit the CREATE, DROP, and ALTER operations on types.                                              | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on types are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on types are audited.</li> </ul>                             |
| Bit 16     | Whether to audit the CREATE, DROP, and ALTER operations on text search objects (CONFIGURATION and DICTIONARY). | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on text search objects are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on text search objects are audited.</li> </ul> |
| Bit 17     | Whether to audit the CREATE, DROP, and ALTER operations on directories.                                        | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations on directories are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations on directories are audited.</li> </ul>                 |
| Bit 18     | Whether to audit the CREATE, DROP, and ALTER operations on synonyms.                                           | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations are not audited.</li> <li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations are audited.</li> </ul>                                               |
| Bit 19     | Whether to audit the CREATE, DROP, and ALTER operations on sequences.                                          | <ul style="list-style-type: none"> <li>• <b>0</b> indicates that the operations are not audited.</li> <li>• <b>1</b> indicates that the operations are audited.</li> </ul>                                                                                               |
| Bit 20     | Reserved.                                                                                                      | -                                                                                                                                                                                                                                                                        |



| Binary Bit | Description                                                                                          | Value Description                                                                                                                                                                                                                                                         |
|------------|------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit 21     | Specifies whether to audit the CREATE, DROP, and ALTER operations on packages.                       | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the operations are not audited.</li><li>• <b>1</b> indicates that the operations are audited.</li></ul>                                                                                                   |
| Bit 22     | Specifies whether to audit the CREATE and DROP operations on models.                                 | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE and ALTER operations are not audited.</li><li>• <b>1</b> indicates that the CREATE and DROP operations are audited.</li></ul>                                                                  |
| Bit 23     | Specifies whether to audit the CREATE, DROP, and ALTER operations on publications and subscriptions. | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations are audited.</li></ul>                                                   |
| Bit 24     | Whether to audit the ALTER and DROP operations on the <b>gs_global_config</b> objects.               | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the ALTER and DROP operations are not audited.</li><li>• <b>1</b> indicates that the ALTER and DROP operations are audited.</li></ul>                                                                     |
| Bit 25     | Specifies whether to audit the CREATE, DROP, and ALTER operations on foreign data wrappers.          | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, DROP, and ALTER operations are not audited.</li><li>• <b>1</b> indicates that the CREATE, DROP, and ALTER operations are audited.</li></ul>                                                   |
| Bit 26     | Whether to audit the CREATE, ENABLE, DISABLE, and DROP operations on SQL patches.                    | <ul style="list-style-type: none"><li>• <b>0</b> indicates that the CREATE, ENABLE, DISABLE, and DROP operations on SQL patches are not audited.</li><li>• <b>1</b> indicates that the CREATE, ENABLE, DISABLE, and DROP operations on SQL patches are audited.</li></ul> |

## audit\_dml\_state

**Parameter description:** Specifies whether to audit the INSERT, UPDATE, and DELETE operations on a specific table.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing the DML operations (except SELECT) is disabled.

- **1** indicates that the function of auditing the DML operations (except SELECT) is enabled.

**Default value:** 0

## audit\_dml\_state\_select

**Parameter description:** Specifies whether to audit the SELECT operation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the SELECT auditing function is disabled.
- **1** indicates that the SELECT auditing function is enabled.

**Default value:** 0

## audit\_function\_exec

**Parameter description:** Specifies whether to record the audit information during the execution of the stored procedures, anonymous blocks, or user-defined functions (excluding system functions).

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of auditing the procedure or function execution is disabled.
- **1** indicates that the function of auditing the procedure or function execution is enabled.

**Default value:** 0

## audit\_copy\_exec

**Parameter description:** Specifies whether to audit the COPY operation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the COPY auditing function is disabled.
- **1** indicates that the COPY auditing function is enabled.

**Default value:** 1

## audit\_set\_parameter

**Parameter description:** Specifies whether to audit the SET operation.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the SET auditing function is disabled.

- **1** indicates that the SET auditing function is enabled.

**Default value:** 0

## audit\_xid\_info

**Parameter description:** Specifies whether to record the transaction ID of the SQL statement in the **detail\_info** column of the audit log.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1

- **0** indicates that the function of recording transaction IDs in audit logs is disabled.
- **1** indicates that the function of recording transaction IDs in audit logs is enabled.

**Default value:** 0

### NOTICE

If this function is enabled, the **detail\_info** information in audit logs starts with *xid*. For example:

```
detail_info: xid=14619 , create table t1(id int);
```

If transaction IDs do not exist, *xid* is recorded as **NA** in audit logs.

## enableSeparationOfDuty

**Parameter description:** Specifies whether the separation of three duties is enabled.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the separation of three duties is enabled.
- **off** indicates that the separation of three duties is disabled.

**Default value:** off

## enable\_nonsysadmin\_execute\_direct

**Parameter description:** Specifies whether non-system administrators and non-monitor administrator are allowed to execute the EXECUTE DIRECT ON statement.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that any user is allowed to execute the EXECUTE DIRECT ON statement.
- **off** indicates that only the system administrators and monitor administrators are allowed to execute the EXECUTE DIRECT ON statement.

**Default value:** off

## enable\_access\_server\_directory

**Parameter description:** Specifies whether to allow non-initial users to create, modify, and delete directories.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that non-initial users have the permission to create, modify, and delete directories.
- **off** indicates that non-initial users do not have the permission to create, modify, and delete directories.

**Default value:** off

---

### NOTICE

- For security purposes, only the initial user can create, modify, and delete DIRECTORY objects by default.
  - If **enable\_access\_server\_directory** is enabled, users with the SYSADMIN permission and users who inherit the **gs\_role\_directory\_create** permission of the built-in role can create directory objects. A user with the SYSADMIN permission, the owner of a directory, a user who is granted with the DROP permission for the directory, or a user who inherits the **gs\_role\_directory\_drop** permission of the built-in role can delete a directory. A user with the SYSADMIN permission and the owner of a directory object can change the owner of the directory object, and the user must be a member of the new owner.
- 

## 17.3.22 CM Parameters

Modifying CM parameters affects the running mechanism of GaussDB. You are advised to ask GaussDB engineers to do it for you. For details about how to modify the CM parameters, see method 1 in [Table 17-2](#).

### 17.3.22.1 CM Agent Parameters

#### log\_dir

**Parameter description:** Specifies the directory where CM Agent logs are stored. The value can be an absolute path, or relative to the CM Agent data directory.

**Value range:** a string The modification takes effect after CM Agent is restarted. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** "log", indicating that CM Agent logs are generated in the CM Agent data directory.

## log\_file\_size

**Parameter description:** Specifies the size of a log file. If a log file exceeds the specified size, a new one is created to record log information.

**Value range:** an integer ranging from 0 to 2047. The unit is MB. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 16 MB

## log\_min\_messages

**Parameter description:** Specifies the message levels to be written to the CM Agent logs. Each level covers all the levels following it. The lower the level is, the fewer messages will be written into the log.

**Value range:** enumerated type. Valid values are **debug5**, **debug1**, **log**, **warning**, **error**, and **fatal**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** warning

## incremental\_build

**Parameter description:** Specifies whether a standby DN is incrementally built. If this parameter is enabled, a standby DN is incrementally built.

**Value range:** Boolean. The value can be **on** or **off**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** on

## alarm\_component

**Parameter description:** Specifies the location of the alarm component that processes alarms.

**Value range:** a string For details about how to modify this parameter, see [Table 17-2](#).

- If **--alarm-type** in the **gs\_preinstall** script is set to **5**, no third-party component is connected and alarms are written into the **system\_alarm** log. In this case, the value of **alarm\_component** is **/opt/huawei/snas/bin/snas\_cm\_cmd**.
- If **--alarm-type** in the **gs\_preinstall** script is set to **1**, a third-party component is connected. In this case, the value of **alarm\_component** is the absolute path of the executable program of the third-party component.

**Default value:** /opt/huawei/snas/bin/snas\_cm\_cmd

## alarm\_report\_interval

**Parameter description:** Specifies the interval at which an alarm is reported. For details about how to modify this parameter, see [Table 17-2](#).

**Value range:** a non-negative integer (unit: s)

**Default value:** 1

## alarm\_report\_max\_count

**Parameter description:** Specifies the maximum number of times an alarm is reported. For details about how to modify this parameter, see [Table 17-2](#).

**Value range:** a non-negative integer

**Default value:** 1

## agent\_report\_interval

**Parameter description:** Specifies the interval at which CM Agent reports the instance status.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1

## agent\_phony\_dead\_check\_interval

**Parameter description:** Specifies the interval at which CM Agent checks whether the DN process is suspended.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10

## agent\_check\_interval

**Parameter description:** Specifies the interval at which CM Agent queries for the status of instances, such as the DNs.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 2

## agent\_heartbeat\_timeout

**Parameter description:** Specifies the heartbeat timeout interval for CM Agent to connect to CM Server.

**Value range:** an integer ranging from 2 to  $2^{31} - 1$ . The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 8

## agent\_connect\_timeout

**Parameter description:** Specifies the time to wait before the attempt of CM Agent to connect to CM Server times out.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1

## agent\_connect\_retries

**Parameter description:** Specifies the number of times CM Agent tries to connect to the CM Server.

**Value range:** an integer. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 15

## agent\_kill\_instance\_timeout

**Parameter description:** Specifies the interval from the time when CM Agent fails to connect to the primary CM Server to the time when CM Agent terminates all instances on the node.

**Value range:** an integer. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0, indicating that the operation of terminating all instances on the node is not initiated.

## security\_mode

**Parameter description:** Specifies whether DNs are started in secure mode. If this parameter is set to **on**, DNs are started in secure mode. Otherwise, DNs are started in non-secure mode.

**Value range:** Boolean. The value can be **on** or **off**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** off

## upgrade\_from

**Parameter description:** Specifies the internal version number of the database before an in-place upgrade. Do not modify the value of this parameter.

**Value range:** a non-negative integer For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0

## process\_cpu\_affinity

**Parameter description:** Specifies whether to bind a primary DN process to a CPU core before starting the process. If this parameter is set to **0**, core binding will not be performed. If it is set to another value, core binding will be performed, and the number of physical CPU cores is  $2^n$ . Restart the database and CM Agent for any modification of this parameter to take effect. Only Arm is supported. For details about how to modify this parameter, see [Table 17-2](#).

**Value range:** an integer ranging from 0 to 2

**Default value:** 0

## log\_threshold\_check\_interval

**Parameter description:** Specifies the interval for compressing and clearing logs. Logs are compressed and cleared every 1800 seconds.

**Value range:** an integer ranging from 0 to 2147483647. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1800

## dilatation\_shard\_count\_for\_disk\_capacity\_alarm

**Parameter description:** Specifies the number of shards to be added in the scale-out scenario. This parameter is used to calculate the threshold for reporting a disk capacity alarm.

### NOTE

The parameter value must be the same as the actual number of shards to be added.

**Value range:** an integer ranging from 0 to  $2^{32} - 1$ . If this parameter is set to 0, the disk scale-out alarm is not reported. If this parameter is set to a value greater than 0, the disk scale-out alarm is reported and the threshold is calculated based on the number of shards specified by this parameter. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1

## log\_max\_size

**Parameter description:** Specifies the maximum size of a log file.

**Value range:** an integer ranging from 0 to 2147483647. The unit is MB. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10240

## log\_max\_count

**Parameter description:** Specifies the maximum number of logs that can be stored on hard disks.

**Value range:** an integer ranging from 0 to 10000. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10000

## log\_saved\_days

**Parameter description:** Specifies the number of days for storing logs.

**Value range:** an integer ranging from 0 to 1000. The unit is day. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 90



## enable\_log\_compress

**Parameter description:** Specifies whether to enable log compression.

**Value range:** Boolean For details about how to modify this parameter, see [Table 17-2](#).

- **on** indicates that log compression is enabled.
- **off** indicates that log compression is disabled.

**Default value:** on

## agent\_backup\_open

**Parameter description:** Specifies the DR database instance settings. After this parameter is enabled, CM runs in DR database instance mode.

**Value range:** 0 or 1. The modification takes effect after CM Agent is restarted. For details about how to modify this parameter, see [Table 17-1](#).

- 0: disabled.
- 1: enabled.

**Default value:** 0

## enable\_xc\_maintenance\_mode

**Parameter description:** Specifies whether the **pgxc\_node** system catalog can be modified when the database instance is in read-only mode.

**Value range:** Boolean The modification takes effect after CM Agent is restarted. For details about how to modify this parameter, see [Table 17-1](#).

- **on** indicates that the **pgxc\_node** system catalog can be modified.
- **off** indicates that the **pgxc\_node** system catalog cannot be modified.

**Default value:** on

## unix\_socket\_directory

**Parameter description:** Specifies the directory location of the Unix socket.

**Value range:** a string For details about how to modify this parameter, see [Table 17-1](#).

**Default value:** "

## enable\_dcf

**Parameter description:** Specifies the status of the DCF mode.

**Value range:** Boolean The modification takes effect after CM Agent is restarted. For details about how to modify this parameter, see [Table 17-1](#).

- 0: disabled.
- 1: enabled.

**Default value:** off

## disaster\_recovery\_type

**Parameter description:** Specifies the type of the DR relationship between the primary and standby database instances.

**Value range:** an integer ranging from 0 to 2 For details about how to modify this parameter, see [Table 17-1](#).

- 0 indicates that no DR relationship is established.
- 1 indicates that the OBS DR relationship is established.
- 2 indicates that the streaming DR relationship is established.

**Default value:** 0

## 17.3.22.2 Parameters Related to CM Server

### log\_dir

**Parameter description:** Specifies the directory where CM Server logs are stored. The value can be an absolute path, or relative to the CM Server data directory.

**Value range:** a string You need to restart CM Server for the modification to take effect. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** "log", indicating that CM Server logs are generated in the CM Server data directory.

### log\_file\_size

**Parameter description:** Specifies the size of a log file. If a log file exceeds the specified size, a new one is created to record log information.

**Value range:** an integer ranging from 0 to 2047. The unit is MB. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 16MB

### log\_min\_messages

**Parameter description:** Specifies which message levels are written to the cm\_server log. Each level covers all the levels following it. The lower the level is, the fewer messages will be written into the log.

**Value range:** enumerated type. Valid values are **debug5**, **debug1**, **log**, **warning**, **error**, and **fatal**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** warning

### thread\_count

**Parameter description:** Specifies the number of threads in the CM Server thread pool.

**Value range:** an integer ranging from 2 to 1000. You need to restart CM Server for the modification to take effect. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1000

## alarm\_component

**Parameter description:** Specifies the location of the alarm component that processes alarms.

**Value range:** a string For details about how to modify this parameter, see [Table 17-2](#).

- If **--alarm-type** in the **gs\_preinstall** script is set to **5**, no third-party component is connected and alarms are written into the **system\_alarm** log. In this case, the value of **alarm\_component** is **/opt/huawei/snas/bin/snas\_cm\_cmd**.
- If **--alarm-type** in the **gs\_preinstall** script is set to **1**, a third-party component is connected. In this case, the value of **alarm\_component** is the absolute path of the executable program of the third-party component.

**Default value:** **/opt/huawei/snas/bin/snas\_cm\_cmd**

## instance\_failover\_delay\_timeout

**Parameter description:** Specifies the delay in the CM Server failover when the primary CM Server breakdown is detected.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0

## instance\_heartbeat\_timeout

**Parameter description:** Specifies the time to wait before the instance heartbeat times out.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 6

## cmserver\_ha\_connect\_timeout

**Parameter description:** Specifies the time to wait before the connection between the primary and standby cm\_servers times out.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 2

## cmserver\_ha\_heartbeat\_timeout

**Parameter description:** Specifies the time to wait before the heartbeat between the primary and standby cm\_servers times out.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 6

## phony\_dead\_effective\_time

**Parameter description:** Specifies the maximum number of times DN processes are detected as zombie. If the number of times a process is detected as zombie is greater than the specified value, the process is considered to be a zombie process and will be restarted.

**Value range:** an integer. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 5

## enable\_transaction\_read\_only

**Parameter description:** Specifies whether the database is in read-only mode.

**Value range:** Boolean values **on**, **off**, **true**, **false**, **yes**, **no**, **1**, and **0**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** on

## datastorage\_threshold\_check\_interval

**Parameter description:** Specifies the interval for checking the disk usage. The system checks the disk usage at the interval specified by the user.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10

## datastorage\_threshold\_value\_check

**Parameter description:** Specifies the usage threshold of a read-only disk in a database. When the disk usage of the data directory exceeds the specified value, the database is automatically set to read-only mode. When adjusting this parameter, you are advised to adjust the **max\_size\_for\_xlog\_retention** parameter of the DN to prevent the cluster read-only threshold from being triggered by backup operations.

**Value range:** an integer ranging from 1 to 99, in percentage. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 85

## max\_datastorage\_threshold\_check

**Parameter description:** Specifies the maximum interval for checking the disk usage. After you modify the **enable\_transaction\_read\_only** parameter, the system automatically checks whether the disk usage reaches the threshold at the specified interval.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

Default value: **43200**

## cmserver\_ha\_status\_interval

**Parameter description:** Specifies the interval between synchronizations of primary and standby CM Server status.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

Default value: **1**

## cmserver\_self\_vote\_timeout

**Parameter description:** Specifies the time to wait before the CM Server self-vote times out.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

Default value: **6**

## alarm\_report\_interval

**Parameter description:** Specifies the interval at which an alarm is reported.

**Value range:** a non-negative integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

Default value: **3**

## alarm\_report\_max\_count

**Parameter description:** Specifies the maximum number of times an alarm is reported.

**Value range:** a non-negative integer. For details about how to modify this parameter, see [Table 17-2](#).

Default value: **1**

## enable\_az\_auto\_switchover

**Parameter description:** Specifies whether to enable automatic AZ switchover. If it is set to **1**, cm\_server automatically switches over services among AZs. Otherwise, when a DN is faulty, services will not be automatically switched to another AZ.

even if the current AZ is unavailable. You can run the switchover command to manually switch services to another AZ.

**Value range:** a non-negative integer. The value **0** indicates that automatic AZ switchover is disabled, and the value **1** indicates that automatic AZ switchover is enabled. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1

## **instance\_keep\_heartbeat\_timeout**

**Parameter description:** The CM Agent periodically checks the instance status and reports the status to the CM Server. If the instance status cannot be detected for a long time and the accumulated number of times exceeds the value of this parameter, the CM Server delivers a command to the CM Agent to restart the instance.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 40

## **az\_switchover\_threshold**

**Parameter description:** If the failure rate of a DN shard in an AZ (Number of faulty DN shards/Total number of DN shards x 100%) exceeds the specified value, automatic AZ switchover is triggered.

**Value range:** an integer ranging from 0 to 100 For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 100

## **az\_check\_and\_arbitrate\_interval**

**Parameter description:** Specifies the interval for checking the AZ status. If the status of an AZ is abnormal, automatic AZ switchover is triggered.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 2

## **az\_connect\_check\_interval**

**Parameter description:** Specifies the interval at which the network connection between AZs is checked.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 60

## **az\_connect\_check\_delay\_time**

**Parameter description:** Specifies the delay between two retries to check the network connection between AZs.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 150

## cmserver\_demote\_delay\_on\_etcd\_fault

**Parameter description:** Specifies the interval at which CM Server switches from the primary state to the standby state due to unhealthy etcd.

**Value range:** an integer. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 8

## instance\_phony\_dead\_restart\_interval

**Parameter description:** Specifies the interval at which the CM Agent process restarts and kills a zombie DN. The interval between two consecutive kill operations cannot be less than the value of this parameter. Otherwise, the CM Agent process does not deliver commands.

**Value range:** an integer. The unit is s. The minimum value that takes effect is **1800**. If this parameter is set to a value less than **1800**, the value **1800** takes effect. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 21600

## cm\_auth\_method

**Parameter description:** Specifies the port authentication mode of the CM. **trust** indicates that port authentication is not configured. **gss** indicates that Kerberos port authentication is used. Note that you can change the value to **gss** only after the Kerberos server and client are successfully installed. Otherwise, the CM cannot communicate properly, affecting the database status.

**Value range:** **gss** or **trust**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** trust

## cm\_krb\_server\_keyfile

**Parameter description:** Specifies the location of the key file on the Kerberos server. The value must be an absolute path. The file is usually stored in the `${GAUSSHOMER}/kerberos` directory and ends with keytab. The file name is the same as the name of the user who runs the database. This parameter is used together with **cm\_auth\_method**. If the **cm\_auth\_method** parameter is changed to **gss**, **cm\_krb\_server\_keyfile** must also be configured as the correct path. Otherwise, the database status will be affected.

**Value range:** a string. For details about how to modify the parameter, see [Table 17-2](#).

**Default value:** `${GAUSSHOMER}/kerberos/${UserName}.keytab`. The default value cannot take effect and is used only as a prompt.

## cm\_server\_arbitrate\_delay\_base\_time\_out

**Parameter description:** Specifies the basic delay duration for cm\_server quorum. If the primary CM Server is disconnected, the arbitration starts to be timed. If the disconnection duration exceeds the arbitration delay duration, a new primary CM Server will be selected. The quorum delay duration is determined by the basic delay duration, the node index (server ID), and the incremental delay duration. The formula is as follows: Quorum delay duration = Basic delay duration + Node index x Incremental delay duration

**Value range:** an integer. The unit is s. The index should be larger than 0. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10

## cm\_server\_arbitrate\_delay\_incremental\_time\_out

**Parameter description:** Specifies the incremental delay duration for CM Server arbitration. If the primary CM Server is disconnected, the arbitration starts to be timed. If the disconnection duration exceeds the arbitration delay duration, a new primary CM Server will be selected. The arbitration delay duration is determined by the basic delay duration, the node index (server ID), and the incremental delay duration. The formula is as follows: Arbitration delay duration = Basic delay duration + Node index x Incremental delay duration

**Value range:** an integer. The unit is s. The index should be larger than 0. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 3

## force\_promote

**Parameter description:** Specifies whether CM Server enables forcible switchover (that is, when the cluster status is unknown, the basic functions of the cluster are available at the cost of data loss). The value **0** indicates that the function is disabled, and the value **1** indicates that the function is enabled. This parameter applies to DNs.

**Value range:** **0** or **1**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0

## switch\_rto

**Parameter description:** Specifies the delay for the forcible startup of CM Server. When **force\_promote** is set to **1** and a shard in the database does not have primary CM Server, the system starts timing. After the delay, forcible switchover is executed.

**Value range:** an integer ranging from 60 to 2147483647. The unit is s. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 600



## backup\_open

**Parameter description:** Specifies the DR database instance settings. After this parameter is enabled, CM runs in DR database instance mode.

**Value range:** 0 or 1. Any modification of this parameter takes effect only after CM Server is restarted. This parameter cannot be enabled for a non-DR database instance. For details about how to modify this parameter, see [Table 17-2](#).

- 0: disabled.
- 1: enabled.

**Default value:** 0

## enable\_dcf

**Parameter description:** Specifies the status of the DCF mode.

**Value range:** Boolean Any modification of this parameter takes effect only after CM Server is restarted. For details about how to modify this parameter, see [Table 17-2](#).

- 0: disabled.
- 1: enabled.

**Default value:** off

## install\_type

**Parameter description:** Specifies the settings related to the DR database instance. It is used to distinguish whether the database instance is based on Dorado.

**Value range:** an integer ranging from 0 to 2 Any modification of this parameter takes effect only after CM Server is restarted. This parameter cannot be enabled for a non-DR database instance. For details about how to modify this parameter, see [Table 17-2](#).

- 0 indicates the database instance for which the DR relationship is not established.
- 1 indicates a Dorado-based database instance.
- 2 indicates a streaming-based database instance.

**Default value:** 0

## enable\_ssl

**Parameter description:** Specifies whether to enable SSL.

**Value range:** Boolean After this function is enabled, the SSL certificate is used to encrypt communication. Any modification of this parameter takes effect only after a restart. For details about how to modify this parameter, see [Table 17-2](#).

- on indicates that SSL is enabled.
- off indicates that SSL is disabled.
- **Default value:** off

---

**NOTICE**

To ensure security, you are advised not to disable it. After this function is disabled, the CM **does not** use encrypted communication and all information is transmitted in plaintext, which may bring security risks such as eavesdropping, tampering, and spoofing.

---

## ssl\_cert\_expire\_alert\_threshold

**Parameter description:** Specifies the SSL certificate expiration alarm time.

**Value range:** an integer. The unit is day. If the certificate expiration time is less than the value of this parameter, an alarm indicating that the certificate is about to expire is reported. Any modification of this parameter takes effect only after a restart. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 90

## ssl\_cert\_expire\_check\_interval

**Parameter description:** Specifies the period for checking whether the SSL certificate expires.

**Value range:** an integer. The unit is s. Any modification of this parameter takes effect only after a restart. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 86400

## delay\_arbitrate\_timeout

**Parameter description:** Specifies the waiting time for a node in the same AZ as the primary DN to be promoted to primary after redo replay.

**Value range:** an integer, in the range [0,21474836] (unit: second). For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0

## ddb\_type

**Parameter description:** Specifies whether to switch between ETCD and DCC modes.

**Value range:** an integer. **0:** ETCD; **1:** DCC. You need to restart cm\_server for the modification to take effect. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 0

## ddb\_log\_level

**Parameter description:** Sets the DDB log level.

To disable the log function, set this parameter to **NONE**, which cannot be used together with the following log levels:

To enable the log function, set this parameter to one or a combination of the following log levels: **RUN\_ERR|RUN\_WAR|RUN\_INF|DEBUG\_ERR|DEBUG\_WAR|DEBUG\_INF|TRACE|PROFILE|OPER**. If two or more log levels are used together, separate them with vertical bars (|). The log level cannot be set to an empty string.

**Value range:** a string containing one or a combination of the following log levels: **RUN\_ERR|RUN\_WAR|RUN\_INF|DEBUG\_ERR|DEBUG\_WAR|DEBUG\_INF|TRACE|PROFILE|OPER**. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** **RUN\_ERR|RUN\_WAR|DEBUG\_ERR|OPER|RUN\_INF|PROFILE**

## **ddb\_log\_backup\_file\_count**

**Parameter description:** Specifies the maximum number of log files that can be saved.

**Value range:** an integer ranging from 1 to 100. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10

## **ddb\_max\_log\_file\_size**

**Parameter description:** Specifies the maximum number of bytes in a log.

**Value range:** a string, in the range from 1 MB to 1000 MB. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 10M

## **ddb\_log\_suppress\_enable**

**Parameter description:** Specifies whether to enable the log suppression function.

**Value range:** an integer. **0**: disabled; **1**: enabled. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 1

## **ddb\_election\_timeout**

**Parameter description:** Specifies the DCC election timeout period.

**Value range:** an integer, in the range [1,600], in seconds. For details about how to modify this parameter, see [Table 17-2](#).

**Default value:** 3

# **17.3.23 Upgrade Parameters**

## **IsInplaceUpgrade**

**Parameter description:** Specifies whether an upgrade is ongoing. This parameter cannot be modified by users. Only the sysadmin user can access this parameter.

This is a SUSE parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates an upgrade is ongoing.
- **off** indicates no upgrade is ongoing.

**Default value:** off

### inplace\_upgrade\_next\_system\_object\_oids

**Parameter description:** Specifies the OID of a new system object during the in-place upgrade. The value of this parameter cannot be changed.

This is a SUSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** empty

### upgrade\_mode

**Parameter description:** Specifies the upgrade mode.

This is a fixed INTERNAL parameter. It can be viewed but cannot be modified.

**Value range:** an integer ranging from 0 to *INT\_MAX*

- **0** indicates that no upgrade is ongoing.
- **1** indicates that a local upgrade is ongoing.
- **2** indicates that a grayscale upgrade is ongoing.

**Default value:** 0

#### NOTE

Special case: When the gray upgrade is used, if the major version upgrade policy is selected, that is, the upgrade script needs to be executed and the binary package needs to be replaced, the value of **upgrade\_mode** is set to **2**; if the minor version upgrade policy is selected, that is, only the binary package needs to be replaced, the value of **upgrade\_mode** is not set to **2**.

## 17.3.24 Miscellaneous Parameters

### enable\_default\_ustore\_table

**Parameter description:** Specifies whether to enable the Ustore storage engine by default. If this parameter is set to **on**, all created tables are Ustore tables.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#). Note that the **track\_counts** and **track\_activities** parameters must be enabled when the Ustore table is used. Otherwise, space expansion may occur.

**Value range:** [off,on]

**Default value:** off

### enable\_ustore

**Parameter description:** Specifies whether to enable the Ustore storage engine. If this parameter is set to **on**, Ustore tables can be created. Note that the

**track\_counts** and **track\_activities** parameters must be enabled when the Ustore table is used. Otherwise, space expansion may occur.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** [off,on]

**Default value:** on

## reserve\_space\_for\_nullable\_atts

**Parameter description:** Specifies whether to reserve space for the nullable attribute of an Ustore table. If this parameter is set to **on**, space is reserved for the nullable attribute of the Ustore table by default.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** [off,on]

**Default value:** on

## ustore\_attr

**Parameter description:** Specifies the Ustore test parameters.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

You can set **enable\_ustore\_partial\_seqscan** (copy selective columns only during the sequential scan in the Ustore table), **enable\_candidate\_buf\_usage\_count** (whether dirty pages are evicted and added to the use count weight), and **ustats\_tracker\_naptime** (time for reloading the statistics file).

**umax\_search\_length\_for\_prune** (number of blocks to be pruned before table extension) and **ustore\_unit\_test** (starting the Ustore white box test). The setting method is **ustore\_attr='Parameter to be set'**. For example, if you want to set **enable\_ustore\_partial\_seqscan**, set **ustore\_attr='enable\_ustore\_partial\_seqscan=on'**.

**Value range:** a string

## server\_version

**Parameter description:** Specifies the server version number.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified. This parameter is inherited from the PostgreSQL kernel, indicating that the current database kernel is compatible with the *server\_version* version corresponding to PostgreSQL. This parameter is reserved to ensure the ecosystem compatibility of the northbound external tool API (query when the tool is connected). This parameter is not recommended. To query the server version, use the `opengauss_version()` function.

**Value range:** a string

**Default value:** 9.2.4

## server\_version\_num

**Parameter description:** Specifies the server version number.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified. This parameter is inherited from the PostgreSQL kernel, indicating that the current database kernel is compatible with the *server\_version\_num* version corresponding to PostgreSQL. This parameter is reserved to ensure the ecosystem compatibility of the northbound external tool API (query when the tool is connected).

**Value range:** an integer

**Default value:** 90204

## block\_size

**Parameter description:** Specifies the block size of the current database.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Value:** 8192

**Default value:** 8192

## segment\_size

**Parameter description:** Specifies the segment file size of the current database.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Unit:** 8 KB

**Default value:** 131072, that is, 1 GB

## max\_index\_keys

**Parameter description:** Specifies the maximum number of index keys supported by the current database.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Default value:** 32

## integer\_datetimes

**Parameter description:** Specifies whether the date and time are in the 64-bit integer format.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Value range:** Boolean

- **on:** yes.
- **off:** no.

**Default value:** on

## lc\_collate

**Parameter description:** Specifies the locale in which sorting of textual data is done.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Default value:** Determined by the configuration set during the database installation and deployment.

## lc\_ctype

**Parameter description:** Specifies the locale that determines character classifications. For example, it specifies what a letter and its upper-case equivalent are.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Default value:** Determined by the configuration set during the database installation and deployment.

## max\_identifier\_length

**Parameter description:** Specifies the maximum identifier length.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Value range:** an integer

**Default value:** 63

## server\_encoding

**Parameter description:** Specifies the database encoding (character set).

By default, `gs_initdb` will initialize the setting of this parameter based on the current system environment. You can also run the **locale** command to check the current configuration environment.

This is a fixed parameter of the INTERNAL type. It can be viewed but cannot be modified.

**Default value:** determined by the current system environment when the database is created.

## enable\_upgrade\_merge\_lock\_mode

**Parameter description:** If this parameter is set to **on**, the delta merge operation internally increases the lock level, and errors can be prevented when update and delete operations are performed at the same time.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- If this parameter is set to **on**, the delta merge operation internally increases the lock level. In this way, when the **DELTAMERGE** operation is concurrently performed with the **UPDATE** or **DELETE** operation, one operation can be performed only after the previous one is complete.
- If this parameter is set to **off** and the **DELTAMERGE** operation is concurrently performed with the **UPDATE** or **DELETE** operation to the data in a row in the delta table of the table, errors will be reported during the later operation, and the operation will stop.

**Default value:** off

## transparent\_encrypt\_kms\_region

**Parameter description:** Specifies the deployment region of the entire database. It must contain only the characters specified in RFC3986, and the maximum length is 2047 bytes. This parameter applies only to the DWS scenario in the current version.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** empty

## basebackup\_timeout

**Parameter description:** Specifies the timeout interval for a connection that has no read or write operations after a backup transfer is complete.

When the gs\_basebackup tool is used for transmission and a high compression rate is specified, the transmission of the tablespace may time out (the client needs to compress the transmitted data).

**Value range:** an integer ranging from 0 to *INT\_MAX*. The unit is s. **0** indicates that archiving timeout is disabled.

**Default value:** 600s

## datanode\_heartbeat\_interval

**Parameter description:** Specifies the interval at which heartbeat messages are sent between heartbeat threads. You are advised to set this parameter to a value no more than wal\_receiver\_timeout/2.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1000 to 60000. The unit is ms.

**Default value:** 1s

## max\_concurrent\_autonomous\_transactions

**Parameter description:** Specifies the maximum number of autonomous transaction connections, that is, the maximum number of concurrent autonomous transactions executed at the same time. If this parameter is set to **0**, autonomous transactions cannot be executed.



This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0–1024

**Default value:** 10

## cluster\_run\_mode

**Parameter description:** Specifies whether a DN belongs to the primary or standby database instance in the dual-database instance DR scenario. For a single database instance, use the default primary database instance.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

- **cluster\_primary** indicates the primary database instance.
- **cluster\_standby** indicates the standby database instance.

**Default value:** cluster\_primary

## dfs\_partition\_directory\_length

**Parameter description:** Specifies the maximum directory name length for the partition directory of a table partitioned by VALUE in the HDFS.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 92 to 7999

**Default value:** 512

## enable\_gpi\_auto\_update

**Parameter description:** Determines whether global indexes are updated by default in partition DDL commands.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- The value **on** indicates that the global index is updated regardless of whether the partition DDL commands contain the UPDATE GLOBAL INDEX clause.
- The value **off** indicates that the global index is not updated unless the partition DDL commands contain the UPDATE GLOBAL INDEX clause.

**Default value:** off

## 17.3.25 Wait Events

### enable\_instr\_track\_wait

**Parameter description:** Specifies whether to enable real-time collection of wait event information.

In the x86-based centralized deployment scenario, the hardware configuration specifications are 32-core CPU and 256 GB memory. When the Benchmark SQL 5.0

tool is used to test performance, the performance fluctuates by about 1.4% by enabling or disabling this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the function of collecting wait event information is enabled.
- **off** indicates that the function of collecting wait event information is disabled.

**Default value:** on

## 17.3.26 Query

### instr\_unique\_sql\_count

**Parameter description:** Specifies the maximum number of Unique SQL records to be collected. The value **0** indicates that the function of collecting Unique SQL information is disabled.

If the value is changed from a larger one to a smaller one, the original data in the system will be cleared and re-collected (the standby node does not support this function). There is no impact if the value is changed from a smaller one to a larger one.

When the number of unique SQL records generated in the system is greater than the value of **instr\_unique\_sql\_count**, the extra unique SQL records are not collected.

In the x86-based centralized deployment scenario, the hardware configuration specifications are 32-core CPU and 256 GB memory. When the Benchmark SQL 5.0 tool is used to test performance, the performance fluctuates by about 3% by enabling or disabling this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647

**Default value:** 200000

### instr\_unique\_sql\_track\_type

**Parameter description:** Specifies which SQL statements are recorded in Unique SQL.

This is an INTERNAL parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated values

**top:** Only top-level SQL statements are recorded.

**Default value:** top

### enable\_instr\_rt\_percentile

**Parameter description:** Specifies whether to enable the function of calculating the response time of 80% and 95% SQL statements in the system.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the function is enabled.
- **off** indicates that the function is disabled.

**Default value:** on

## percentile

**Parameter description:** Specifies the percentage of SQL statements whose response time is to be calculated by the background calculation thread.

This is an INTERNAL parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** 80, 95

## instr\_rt\_percentile\_interval

**Parameter description:** Specifies the interval at which the background calculation thread calculates the SQL response time.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 3600. The unit is s.

**Default value:** 10s

## enable\_instr\_cpu\_timer

**Parameter description:** Specifies whether to capture the CPU time consumed during SQL statement execution.

In the x86-based centralized deployment scenario, the hardware configuration specifications are 32-core CPU and 256 GB memory. When the Benchmark SQL 5.0 tool is used to test performance, the performance fluctuates by about 3.5% by enabling or disabling this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates the CPU time consumed during SQL statement execution is captured.
- **off** indicates the CPU time consumed during SQL statement execution is not captured.

**Default value:** on

## enable\_stmt\_track

**Parameter description:** Specifies whether to enable the full/slow SQL statement feature.

In the x86-based centralized deployment scenario, the hardware configuration specifications are 32-core CPU and 256 GB memory. When the Benchmark SQL 5.0

tool is used to test performance, the performance fluctuates by about 1.2% by enabling or disabling this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** Full/Slow SQL capture is enabled.
- **off:** Full /Slow SQL capture is disabled.

**Default value:** on

## track\_stmt\_parameter

**Parameter description:** After **track\_stmt\_parameter** is enabled, the executed statements recorded in **statement\_history** are not normalized. The complete SQL statement information can be displayed to help the database administrator locate problems. For a simple query, the complete statement information is displayed. For a PBE statement, the complete statement information and information about each variable value are displayed. The format is query string; parameters: \$1=value1,\$2=value2, .... This parameter is used to display complete SQL information and is not controlled by the **track\_activity\_query\_size** parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** The function of displaying complete SQL statement information is enabled.
- **off:** The function of displaying complete SQL statement information is disabled.

**Default value:** off

## track\_stmt\_session\_slot

**Parameter description:** Specifies the maximum number of full/slow SQL statements that can be cached in a session. If the number of full/slow SQL statements exceeds this value, new statements will not be traced until the flush thread flushes the cached statements to the disk to reserve idle space.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647

**Default value:** 1000

## track\_stmt\_details\_size

**Parameter description:** Specifies the maximum size (in bytes) of execution events that can be collected by a single statement.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000000

**Default value:** 4096

## track\_stmt\_retention\_time

**Parameter description:** Specifies the retention period of full/slow SQL statement records. This parameter is a combination of parameters. This parameter is read every 60 seconds and records exceeding the retention period are deleted. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

This parameter consists of two parts in the format of 'full sql retention time, slow sql retention time'.

**full sql retention time** indicates the retention time of full SQL statements. The value ranges from 0 to 86400.

**slow sql retention time** indicates the retention time of slow SQL statements. The value ranges from 0 to 604800.

**Default value:** 3600,604800

## track\_stmt\_stat\_level

**Parameter description:** Determines the level of statement execution tracing.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#). The value is case-insensitive.

**Value range:** a string

This parameter consists of two parts in the format of 'full sql stat level, slow sql stat level'.

The first part indicates the tracing level of full SQL statements. The value can be **OFF**, **L0**, **L1**, or **L2**.

The second part indicates the tracing level of slow SQL statements. The value can be **OFF**, **L0**, **L1**, or **L2**.

### NOTE

If the tracing level of full SQL statements is not **OFF**, the current SQL statement tracing level is a higher level ( $L2 > L1 > L0$ ) of full and slow SQL statements. For details about the levels, see "System Catalogs and System Views > System Catalogs > STATEMENT\_HISTORY > STATEMENT\_HISTORY columns" in *Developer Guide*.

**Default value:** OFF,L0

## enable\_auto\_clean\_unique\_sql

**Parameter description:** Specifies whether to enable the automatic elimination function of unique SQL statements when the number of unique SQL statements generated in the system is equal to or greater than the value of **instr\_unique\_sql\_count**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** off



**CAUTION**

Some snapshot information comes from unique SQL statements. Therefore, when automatic elimination is enabled, if the selected start snapshot and end snapshot exceed the elimination time, the WDR report cannot be generated.

---

## asp\_log\_directory

**Parameter description:** Specifies the directory for storing ASP log files on the server when **asp\_flush\_mode** is set to **all** or **file**. The value can be an absolute path, or relative to the **data** directory. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

---

**NOTICE**

If the value of **asp\_log\_directory** in the configuration file is an invalid path, the database instance cannot be restarted.

---



**NOTE**

- Valid path: Users must have read and write permissions on the path.
- Invalid path: Users do not have read or write permissions on an invalid path.

**Value range:** a string

**Default value:** specified during installation

## perf\_directory

**Parameter description:** Specifies the directory of the output file of the performance view dotting task. Only the sysadmin user can access this parameter. The value can be an absolute path, or relative to the data directory.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).



**NOTE**

- Valid path: Users must have read and write permissions on the path.
- Invalid path: Users do not have read or write permissions on an invalid path.

**Value range:** a string

**Default value:** specified during installation

## unique\_sql\_retention\_time

**Parameter description:** Specifies the memory cleanup interval for the unique SQL hash table. The default value is 30 minutes.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 3650

**Default value:** 30 min

## 17.3.27 System Performance Snapshot

### enable\_wdr\_snapshot

**Parameter description:** Specifies whether to enable the database monitoring snapshot function.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the database monitoring snapshot function is enabled.
- **off** indicates that the database monitoring snapshot function is disabled.

**Default value:** on

### wdr\_snapshot\_retention\_days

**Parameter description:** Specifies the number of days for storing database monitoring snapshot data in the system. When the number of snapshots generated during database running exceeds the maximum number of snapshots that can be generated within the retention period, the system clears the snapshot data with the smallest **snapshot\_id** at the interval specified by **wdr\_snapshot\_interval**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 1 to 8

**Default value:** 8

### wdr\_snapshot\_query\_timeout

**Parameter description:** Specifies the execution timeout for the SQL statements associated with database monitoring snapshot operations. If the SQL statement execution is not complete and a result is not returned within the specified time, the snapshot operation fails.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 100 to *INT\_MAX*. The unit is s.

**Default value:** 100s

### wdr\_snapshot\_interval

**Parameter description:** Specifies the interval at which the backend thread Snapshot automatically performs snapshot operations on the database monitoring data.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 10 to 60. The unit is min.

**Default value:** 1h

## asp\_flush\_mode

**Parameter description:** Specifies the mode for the ASP to update data to the disk. The value can be **file** (default value), **table** (system catalog), or **all** (system catalog and file). Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** a string, which can be **table**, **file**, or **all**

**Default value:** table

## asp\_flush\_rate

**Parameter description:** When the number of samples reaches the value of **asp\_sample\_num**, the samples in the memory are updated to the disk based on a certain proportion. **asp\_flush\_rate** indicates the update proportion. If this parameter is set to **10**, it indicates that the update ratio is 10:1.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** an integer ranging from 1 to 10

**Default value:** 10

## asp\_log\_filename

**Parameter description:** Specifies the file name format when writing files using ASP. Only the sysadmin user can access this parameter.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** a string

**Default value:** asp-%Y-%m-%d\_%H%M%S.log

## asp\_retention\_days

**Parameter description:** Specifies the maximum number of days for reserving ASP samples when they are written to the system catalog.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 7

**Default value:** 2

## asp\_sample\_interval

**Parameter description:** Specifies the sampling interval.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 10. The unit is s.



**Default value:** 1s

## asp\_sample\_num

**Parameter description:** Specifies the maximum number of samples allowed in the LOCAL\_ACTIVE\_SESSION view. Only the sysadmin user can access this parameter.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 100000.

**Default value:** 100000

## enable\_asp

**Parameter description:** Specifies whether to enable the active session profile function.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** The function is enabled.
- **off:** The function is disabled.

**Default value:** on

## 17.3.28 Security Configuration

### elastic\_search\_ip\_addr

**Parameter description:** Specifies the IP address of the Elasticsearch system.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

**Default value:** 'https:127.0.0.1'

### enable\_security\_policy

**Parameter description:** Specifies whether the unified audit and dynamic data masking policies take effect.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**on:** The security policy is enabled.

**off:** The security policy is disabled.

**Default value:** off

### use\_elastic\_search

**Parameter description:** Specifies whether to send unified audit logs to Elasticsearch. If **enable\_security\_policy** and this parameter are enabled, unified

audit logs are sent to Elasticsearch through HTTP or HTTPS (used by default). After this parameter is enabled, ensure that the Elasticsearch service corresponding to **elastic\_search\_ip\_addr** can be properly connected. Otherwise, the process fails to be started.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**on:** Unified audit logs are sent to Elasticsearch.

**off:** Unified audit logs are not sent to Elasticsearch.

**Default value:** off

## is\_sysadmin

**Parameter description:** Specifies whether the current user is an initial user.

This is a fixed INTERNAL parameter and cannot be modified.

**Value range:** Boolean

**on** indicates that the user is an initial user.

**off** indicates that the user is not an initial user.

**Default value:** off

## 17.3.29 Global Temporary Table

### max\_active\_global\_temporary\_table

**Parameter description:** Specifies whether global temporary tables can be created. Currently, the Ustore engine does not support global temporary tables. The value of this parameter determines the memory reserved in the shared cache for hash tables required by global temporary tables. The total number of active global temporary tables in all sessions is not forcibly limited.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000

- **0:** The global temporary table function is disabled.
- **> 0:** The global temporary table function is enabled.

**Default value:** 1000

### vacuum\_gtt\_defer\_check\_age

**Parameter description:** Checks the differences between the global temporary table relfrozenxid and the ordinary table after VACUUM is executed. WARNING is generated if the difference value exceeds the specified parameter value. Use the default value for this parameter.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000

**Default value:** 10000

## **enable\_gtt\_concurrent\_truncate**

**Parameter description:** Specifies whether to support concurrent execution of TRUNCATE TABLE and DML operations on global temporary tables and concurrent execution of TRUNCATE TABLE on global temporary tables.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on/true** indicates that the preceding operations can be performed concurrently.
- **off/false** indicates that the preceding operations cannot be performed concurrently.

**Default value:** on

## **17.3.30 HyperLogLog**

### **hll\_default\_log2m**

**Parameter description:** Specifies the number of buckets for HLL data. The number of buckets affects the precision of distinct values calculated by HLL. The more buckets there are, the smaller the deviation is. The deviation range is as follows:  $[-1.04/2^{\log_2 m^{*1/2}}, +1.04/2^{\log_2 m^{*1/2}}]$

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 16

**Default value:** 14

### **hll\_default\_log2explicit**

**Parameter description:** Specifies the default threshold for switching from the explicit mode to the sparse mode.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 12 The value 0 indicates that the explicit mode is skipped. The value 1 to 12 indicates that the mode is switched when the number of distinct values reaches  $2^{\text{hll\_default\_log2explicit}}$ .

**Default value:** 10

### **hll\_default\_log2sparse**

**Parameter description:** Specifies the default threshold for switching from the sparse mode to the full mode.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 14 The value 0 indicates that the explicit mode is skipped. The value 1 to 14 indicates that the mode is switched when the number of distinct values reaches  $2^{\text{hll\_default\_log2sparse}}$ .

**Default value:** 12

## hll\_duplicate\_check

**Parameter description:** Specifies whether duplicatecheck is enabled by default.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1 0: disabled; 1: enabled

**Default value:** 0

## 17.3.31 User-defined Functions

### udf\_memory\_limit

**Parameter description:** Controls the maximum physical memory that can be used when each database node executes UDFs. This parameter does not take effect in the current version. Use **FencedUDFMemoryLimit** and **UDFWorkerMemHardLimit** to control virtual memory used by fenced udf worker.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer. The value range is from 200 x 1024 to *max\_process\_memory* and the unit is KB.

**Default value:** 200 MB

### FencedUDFMemoryLimit

**Parameter description:** Specifies the virtual memory used by each fenced udf worker process.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 KB to 2147483647 KB. The unit can also be MB or GB. 0 indicates that the memory is not limited.

**Default value:** 0

### UDFWorkerMemHardLimit

**Parameter description:** Specifies the maximum value of **fencedUDFMemoryLimit**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 KB to 2147483647 KB. The unit can also be MB or GB.

**Default value:** 1 GB

## 17.3.32 Scheduled Task

### job\_queue\_processes

**Parameter description:** Specifies the number of jobs that can be concurrently executed. This parameter is a POSTMASTER parameter. You can set it using **gs\_guc**, and you need to restart **gaussdb** to make the setting take effect.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 to 1000

Function:

- Setting **job\_queue\_processes** to **0** indicates that the scheduled job function is disabled and that no job will be executed. (Enabling scheduled jobs may affect the system performance. At sites where this function is not required, you are advised to disable it.)
- Setting **job\_queue\_processes** to a value that is greater than **0** indicates that the scheduled job function is enabled and this value is the maximum number of jobs that can be concurrently processed.

After the scheduled job function is enabled, the job\_scheduler thread polls the **pg\_job** system catalog at a scheduled interval. The scheduled job check is performed every second by default.

Too many concurrent jobs consume many system resources, so you need to set the number of concurrent jobs to be processed. If the current number of concurrent jobs reaches the value of **job\_queue\_processes** and some of them expire, these jobs will be postponed to the next polling period. Therefore, you are advised to set the polling interval (the **Interval** parameter of the submit API) based on the execution duration of each job to avoid the problem that jobs in the next polling period cannot be properly processed because of overlong job execution time.

Note: If the number of concurrent jobs is large and the value is too small, these jobs will wait in queues. However, a large parameter value leads to large resource consumption. You are advised to set this parameter to **100** and change it based on the system resource condition.

**Default value:** 10

### enable\_prevent\_job\_task\_startup

**Parameter description:** Specifies whether to start the job thread.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the job thread is not started.
- **off** indicates that the job thread is started.

**Default value:** off

## 17.3.33 Thread Pool

### enable\_thread\_pool

**Parameter description:** Specifies whether to enable the thread pool function. This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the thread pool function is enabled.
- **off** indicates that the thread pool function is disabled.

**Default value:** on

### thread\_pool\_attr

**Parameter description:** Specifies the detailed attributes of the thread pool function. This parameter is valid only when **enable\_thread\_pool** is set to **on**. Only the sysadmin user can access this parameter. This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

This parameter consists of three parts: thread\_num, group\_num, and cpubind\_info. The meanings of the three parts are as follows:

- **thread\_num** is the total number of initial threads in the thread pool, which can be dynamically expanded. The value ranges from 0 to 4096. The value **0** indicates that the database automatically configures the number of threads in the thread pool based on the number of CPU cores. If the value is greater than **0**, the number of threads in the thread pool is the same as the value of **thread\_num**. You are advised to set the thread pool size based on the hardware configuration. The formula is as follows: Value of **thread\_num** = Number of CPU cores x 3–5. The maximum value of **thread\_num** is **4096**.
- **group\_num** indicates the number of thread groups in the thread pool. The value ranges from 0 to 64. The value **0** indicates that the database automatically configures the number of thread groups in the thread pool based on the number of NUMA groups. If the value is greater than **0**, the number of thread groups in the thread pool is the same as the value of **group\_num**.
- **cpubind\_info** indicates whether the thread pool is bound to a core. The available configuration modes are as follows: 1. '**(nobind)**': The thread is not bound to a core. 2. '**(allbind)**': Use all CPU cores that can be queried in the current system to bind threads. 3. '**(nodebind: 1, 2)**': Use the CPU cores in NUMA groups 1 and 2 to bind threads. 4. '**(cpubind: 0-30)**': Use CPU cores 0 to 30 to bind threads. 5. '**(numabind: 0-30)**': Use CPU cores 0 to 30 in the NUMA group to bind threads. This parameter is case-insensitive.

**Default value:**

'**4096,2,(nobind)**' (128-core CPU/1024-GB memory, 104-core CPU/1024-GB memory, and 96-core CPU/1024-GB memory); '**2048,2,(nobind)**' (96-core CPU/768-GB memory); '**1024,2,(nobind)**' (64-core CPU/512-GB memory, 60-core CPU/480-GB memory, and 32-core CPU/256-GB memory); '**512,2,(nobind)**' (16-core CPU/128-GB memory); '**256,2,(nobind)**' (8-core CPU/64-GB memory);

## thread\_pool\_stream\_attr

**Parameter description:** Specifies the detailed attributes of the stream thread pool function. This parameter is valid only when **enable\_thread\_pool** is set to **on** and only takes effect on DNs. Only the sysadmin user can access this parameter. This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

This parameter consists of four parts: 'stream\_thread\_num, stream\_proc\_ratio ,group\_num ,cpubind\_info'. The meanings of the four parts are as follows:

- **stream\_thread\_num** indicates the total number of threads in the stream thread pool. The value ranges from 0 to 4096. The value **0** indicates that the database automatically configures the number of threads in the thread pool based on the number of CPU cores. If the value is greater than **0**, the number of threads in the thread pool is the same as the value of **stream\_thread\_num**. You are advised to set the thread pool size based on the hardware configuration. The formula is as follows: Value of **stream\_thread\_num** = Number of CPU cores x 3-5. The maximum value of **stream\_thread\_num** is **4096**.
- **stream\_proc\_ratio** indicates the ratio of proc resources reserved for stream threads. The value is a floating point number. The default value is **0.2**. The reserved proc resources are calculated as follows: **stream\_proc\_ratio** x **stream\_thread\_num**.
- **group\_num** indicates the number of thread groups in the thread pool. The value ranges from 0 to 64. The value **0** indicates that the database automatically configures the number of thread groups in the thread pool based on the number of NUMA groups. If the value is greater than **0**, the number of thread groups in the thread pool is the same as the value of **group\_num**. The value of **group\_num** in **thread\_pool\_stream\_attr** must be the same as that in **thread\_pool\_attr**. If they are set to different values, the value of **group\_num** in **thread\_pool\_attr** is used.
- **cpubind\_info** indicates whether the thread pool is bound to a core. The available configuration modes are as follows: 1. '**(nobind)**': The thread is not bound to a core. 2. '**(allbind)**': Use all CPU cores that can be queried in the current system to bind threads. 3. '**(nodebind: 1, 2)**': Use the CPU cores in NUMA groups 1 and 2 to bind threads. 4. '**(cpubind: 0-30)**': Use CPU cores 0 to 30 to bind threads. 5. '**(numabind: 0-30)**': Use CPU cores 0 to 30 in the NUMA group to bind threads. This parameter is case-insensitive. The value of **cpubind\_info** in **thread\_pool\_stream\_attr** must be the same as that in **thread\_pool\_attr**. If they are set to different values, the value of **cpubind\_info** in **thread\_pool\_attr** is used.

**Default value:**

**stream\_thread\_num:** 16

**stream\_proc\_ratio:** 0.2

**group\_num** and **cpubind\_info:** For details, see [thread\\_pool\\_attr](#).

## resilience\_threadpool\_reject\_cond

**Parameter description:** Specifies the percentage of accumulated sessions in the thread pool for escape from overload. This parameter takes effect only when the GUC parameters **use\_workload\_manager** and **enable\_thread\_pool** are enabled. This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string, consisting of one or more characters

This parameter consists of **recover\_threadpool\_percent** and **overload\_threadpool\_percent**. The meanings of the two parts are as follows:

- **recover\_threadpool\_percent:** Percentage of the number of sessions that are recovered to the normal state in the initial number of threads in the thread pool. When the number of accessed sessions is less than the initial number of threads in the thread pool multiplied by the value of this parameter, the escape from overload function is disabled and new connections are allowed. The value ranges from 0 to *INT\_MAX*. The value indicates a percentage.
- **overload\_threadpool\_percent:** Percentage of the number of accessed sessions to the initial number of threads in the thread pool when the thread pool is overloaded. If the number of accessed sessions is greater than the initial number of threads in the thread pool multiplied by the value of this parameter, the current thread pool is overloaded. In this case, the escape from overload function is enabled to kill sessions and forbid new connections to access the thread pool. The value ranges from 0 to *INT\_MAX*. The value indicates a percentage.

**Default value:** '0,0', indicating that the thread pool escape function is disabled.

### Example:

```
resilience_threadpool_reject_cond = '100,200'
```

When the number of stacked sessions exceeds 200% of the initial number of threads in the thread pool, new connections are forbidden and stacked sessions are killed. When the number of stacked sessions is less than 100% of the initial number of threads in the thread pool, new connections are allowed.

---

### NOTICE

- The number of stacked sessions can be obtained by querying the number of data records in the **pg\_stat\_activity** view. A few background threads need to be filtered out. The initial number of threads in the thread pool can be obtained by querying the **thread\_pool\_attr** parameter.
  - If this parameter is set to a small value, the thread pool escape from overload process is frequently triggered. As a result, ongoing sessions are forcibly logged out, and new connections fail to be connected for a short period of time. Therefore, exercise caution when setting this parameter based on the actual thread pool usage.
  - The values of **recover\_threadpool\_percent** and **overload\_threadpool\_percent** can be 0 at the same time. In addition, the value of **recover\_threadpool\_percent** must be smaller than that of **overload\_threadpool\_percent**. Otherwise, the setting does not take effect.
-



## 17.3.34 Backup and Restoration

### operation\_mode

**Parameter description:** Specifies whether the system enters the backup and restoration mode.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that the system is in the backup and restoration mode.
- **off** indicates that the system is not in the backup and restoration mode.

**Default value:** off

### enable\_cbm\_tracking

**Parameter description:** This parameter must be enabled when Roach is used to perform full and incremental backups. If this parameter is disabled, the backup will fail.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on:** The cbm tracking is enabled.
- **off:** The cbm tracking is disabled.

**Default value:** off

### max\_size\_for\_xlog\_retention

**Parameter description:** Specifies when to forcibly push the backup replication slot to prevent the disk from being full and the cluster from being read-only because logs cannot be recycled during backup. It is recommended that the value of this parameter be a little smaller than the value of **datastorage\_threshold\_value\_check** of the CM Server component to prevent the cluster from entering the read-only state.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** -100 to 2147483647

- The value **0** indicates that this function is disabled.
- A negative value indicates that the backup replication slot is forcibly pushed when the disk usage exceeds 80% and log recycling is blocked due to backup operations. For example, **-80** indicates that the backup replication slot is forcibly pushed when the disk usage exceeds 80%.
- A positive value indicates that the backup replication slot is triggered based on the size of stacked logs. For example, **32** indicates that the backup replication slot is forcibly pushed when the redo position of the current checkpoint exceeds 32 log segments (the size of each log segment is 16 MB) and logs are recycled because the backup operation is blocked.

**Default value:** 0

## 17.3.35 Undo

### undo\_space\_limit\_size

**Parameter description:** Specifies the threshold for forcibly recycling undo space. When the undo space usage reaches 80% of the threshold, forcible recycling starts. You can set this parameter to a large value based on service requirements and then set it to a proper value based on the actual undo space usage. The value of **undo\_space\_limit\_size** must be equal to or greater than that of **undo\_limit\_size\_per\_transaction**.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 102400 to 2147483647. The unit is 8 KB.

**Default value:** 256GB

### undo\_limit\_size\_per\_transaction

**Parameter description:** Specifies the undo space threshold of a single transaction. If the threshold is reached, the transaction is rolled back due to an error. You are advised to set **undo\_limit\_size\_per\_transaction** to a value less than or equal to that of **undo\_space\_limit\_size**. If the value of **undo\_limit\_size\_per\_transaction** is greater than that of **undo\_space\_limit\_size**, the displayed value is the same as the configured value when you run the **show undo\_limit\_size\_per\_transaction** command to query the parameter. The only difference is that the smaller value between **undo\_space\_limit\_size** and **undo\_limit\_size\_per\_transaction** is used as the actual undo space threshold of a single transaction. If the **undo\_limit\_size\_per\_transaction** is set to a value greater than 1 TB, the system performance and stability may be affected.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 256 to 2147483647. The unit is 8 KB.

**Default value:** 32 GB

## 17.3.36 DCF Parameters Settings

### enable\_dcf

**Parameter description:** Specifies whether to enable the DCF mode. This parameter cannot be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean. The value can be **on** or **off**. The value **on** indicates that the current installation mode is DCF, and the value **off** indicates that the current installation mode is not DCF.

**Default value:** off

### dcf\_ssl

**Parameter description:** Specifies whether to enable SSL. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean. The value can be **on** or **off**. The value **on** indicates that SSL is used, and the value **off** indicates that SSL is not used.

**Default value:** on

## dcf\_config

**Parameter description:** Specifies the customized configuration information during installation. This parameter cannot be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Default value:** a string, which is specified by users during installation

## dcf\_data\_path

**Parameter description:** Specifies the DCF data path. This parameter cannot be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Default value:** a string, which is the **dcf\_data** directory under the data directory of the DN

## dcf\_log\_path

**Parameter description:** Specifies the DCF log path. This parameter cannot be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Default value:** a string, which is the **dcf\_log** directory under the data directory of the DN.

## dcf\_node\_id

**Parameter description:** Specifies the ID of the DN where the DCF is located. This parameter is defined by users during installation and cannot be modified.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Default value:** an integer, which is specified by users during installation

## dcf\_max\_workers

**Parameter description:** Specifies the number of DCF callback function threads. If the number of nodes exceeds 7, increase the value of this parameter (for example, to 40). Otherwise, the primary node may remain in the promoting state and the log replication between the primary and standby nodes has no progress.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 262143

**Default value:** 10

## dcf\_truncate\_threshold

**Parameter description:** Specifies the threshold for a DN to truncate DCF logs.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647

**Default value:** 100000

## dcf\_election\_timeout

**Parameter description:** Specifies the timeout interval for selecting the DCF leader and follower. The election timeout interval depends on the status of the network between DNs. If the timeout interval is short and the network quality is poor, timeout occurs. After the network recovers, the election becomes normal. You are advised to set a proper timeout interval based on the current network status. Restriction on the DCF node clock: The maximum clock difference between DCF nodes is less than half of the election timeout period.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 600, in seconds

**Default value:** 3

## dcf\_enable\_auto\_election\_priority

**Parameter description:** Specifies whether the DCF priority can be automatically adjusted. The value **0** indicates that automatic adjustment is not allowed, and the value **1** indicates that automatic adjustment is allowed.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** 0 or 1.

**Default value:** 1

## dcf\_election\_switch\_threshold

**Parameter description:** Specifies the DCF threshold for preventing frequent switchover to primary. It is recommended that this parameter be set based on the maximum fault duration acceptable for user services.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647, in seconds.

**Default value:** 0

## dcf\_run\_mode

**Parameter description:** Specifies the DCF election mode. The value **0** indicates that the automatic DCF election mode is enabled, and the value **2** indicates that the DCF election mode is disabled. Currently, the election mode can be disabled only in minority restoration scenarios. If the election mode is disabled, the database instance will become unavailable.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated type. The value can be **0** or **2**.

**Default value:** **0**

## dcf\_log\_level

**Parameter description:** Specifies the DCF log level.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** a string

- Disable the log function: **NONE**, indicating that the log function is disabled and cannot be used for the following log levels:
- Enable the log function: **RUN\_ERR|RUN\_WAR|RUN\_INF|DEBUG\_ERR|DEBUG\_WAR|DEBUG\_INF|TRACE|PROFILE|OPER**

You can select a string from the preceding strings and use vertical bars (|) to combine the strings. The log level cannot be left blank.

**Default value:** **RUN\_ERR|RUN\_WAR|DEBUG\_ERR|OPER|RUN\_INF|PROFILE**

## dcf\_log\_backup\_file\_count

**Parameter description:** Specifies the number of DCF run log backups.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 100

**Default value:** **10**

## dcf\_max\_log\_file\_size

**Parameter description:** Specifies the maximum size of a DCF run log file.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 1000, in MB

**Default value:** **10**

## dcf\_socket\_timeout

**Parameter description:** Specifies the timeout interval for the DCF communication module to connect to the socket. This parameter takes effect upon the system restart. In an environment where the network quality is poor, if the timeout interval is set to a small value, the connection may fail to be set up. In this case, you need to increase the value.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 600000, in ms

**Default value:** **5000**

## dcf\_connect\_timeout

**Parameter description:** Specifies the timeout interval for the DCF communication module to set up a connection. This parameter takes effect upon the system restart. In an environment where the network quality is poor, if the timeout interval is set to a small value, the connection may fail to be set up. In this case, you need to increase the value.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 10 to 600000, in ms

**Default value:** 60000

## dcf\_mec\_fragment\_size

**Parameter description:** Specifies the fragment size of the DCF communication module. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 10240, in KB

**Default value:** 64

## dcf\_stg\_pool\_max\_size

**Parameter description:** Specifies the maximum size of the memory pool of the DCF storage module. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 2147483647, in MB

**Default value:** 2048

## dcf\_stg\_pool\_init\_size

**Parameter description:** Specifies the minimum size of the memory pool of the DCF storage module. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 2147483647, in MB

**Default value:** 32

## dcf\_mec\_pool\_max\_size

**Parameter description:** Specifies the maximum size of the memory pool of the DCF communication module. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 2147483647, in MB

**Default value:** 200

## dcf\_flow\_control\_disk\_rawait\_threshold

**Parameter description:** Specifies the disk waiting threshold for DCF flow control.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647, in microseconds

**Default value:** 100000

## dcf\_flow\_control\_net\_queue\_message\_num\_threshold

**Parameter description:** Specifies the threshold for the number of messages in a network queue for DCF flow control.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647

**Default value:** 1024

## dcf\_flow\_control\_cpu\_threshold

**Parameter description:** Specifies the threshold for DCF CPU flow control.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 2147483647, in percentage (%)

**Default value:** 100

## dcf\_mec\_batch\_size

**Parameter description:** Specifies the number of batch messages for DCF communication. When the value is 0, the DCF automatically adjusts the value based on the network and the amount of data to be written. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1024

**Default value:** 0

## dcf\_mem\_pool\_max\_size

**Parameter description:** Specifies the maximum DCF memory. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 2147483647, in MB

**Default value:** 2048

## dcf\_mem\_pool\_init\_size

**Parameter description:** Specifies the initial size of the DCF memory. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 32 to 2147483647, in MB

**Default value:** 32

## dcf\_compress\_algorithm

**Parameter description:** Specifies the compression algorithm for DCF run log transmission. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer

- 0 indicates no compression.
- 1 indicates the ZSTD compression algorithm.
- 2 indicates the LZ4 compression algorithm.

**Default value:** 0

## dcf\_compress\_level

**Parameter description:** Specifies the compression level for DCF log transmission. This parameter takes effect upon the system restart. Before this parameter takes effect, a valid compression algorithm must be configured, that is, the **dcf\_compress\_algorithm** parameter is set.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 22

If compression is disabled, the configured compression level does not take effect.

**Default value:** 1

## dcf\_mec\_channel\_num

**Parameter description:** Specifies the number of DCF communication channels. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 64

**Default value:** 1

## dcf\_rep\_append\_thread\_num

**Parameter description:** Specifies the number of DCF log replication threads. This parameter takes effect upon the system restart.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 1000

**Default value:** 2



## dcf\_mec\_agent\_thread\_num

**Parameter description:** Specifies the number of DCF communication working threads. This parameter takes effect upon the system restart. It is recommended that the value of **dcf\_mec\_agent\_thread\_num** be equal to or greater than 2 x Number of nodes x Value of **dcf\_mec\_channel\_num**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 1000

**Default value:** 10

## dcf\_mec\_reactor\_thread\_num

**Parameter description:** Specifies the number of reactor threads used by the DCF. This parameter takes effect upon the system restart. It is recommended that the ratio of the value of **dcf\_mec\_reactor\_thread\_num** to the value of **dcf\_mec\_agent\_thread\_num** be 1:40.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 100

**Default value:** 1

## dcf\_log\_file\_permission

**Parameter description:** Specifies the attribute of the DCF run log file. The parameter setting takes effect after the system is restarted. This parameter is configured during installation and cannot be modified. To allow other users in the same group to access logs, ensure that all parent directories can be accessed by other users in the same group. That is, if **dcf\_log\_path\_permission** is set to **750**, **dcf\_log\_file\_permission** can only be set to **600** or **640**. If **dcf\_log\_path\_permission** is set to **700**, **dcf\_log\_file\_permission** can only be set to **600**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated type. The value can be **600** or **640**.

**Default value:** 600

## dcf\_log\_path\_permission

**Parameter description:** Specifies the attribute of the DCF run log directory. The parameter setting takes effect after the system is restarted. This parameter is configured during installation and cannot be modified. To allow other users in the same group to access the log path, set this parameter to **750**. Otherwise, set this parameter to **700**.

This is a POSTMASTER parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** enumerated type. The value can be **700** or **750**.

**Default value:** 700

## 17.3.37 Flashback

This section describes parameters related to the flashback function. In this version, only the Ustore engine supports flashback, while the Astore engine does not support flashback.

### enable\_recyclebin

**Parameter description:** Specifies whether the recycle bin is enabled or disabled in real time.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

**Default value:** off

### recyclebin\_retention\_time

**Parameter description:** Specifies the retention period of objects in the recycle bin. The objects will be automatically deleted after the retention period expires.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 2147483647. The unit is s.

**Default value:** 15 min (900s)

### version\_retention\_age

**Parameter description:** Specifies the number of transactions retained in the earlier version. If the number of transactions exceeds the value of this parameter, the earlier version will be recycled and cleared.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 576460752303423487. 0 means no delay.

**Default value:** 0



This parameter has been deprecated.

---

### vacuum\_defer\_cleanup\_age

**Parameter description:** Specifies the number of transactions by which **VACUUM** will defer the cleanup of invalid row-store table records, so that **VACUUM** and **VACUUM FULL** do not clean up deleted tuples immediately. You can also set this parameter to configure the retention period of the flashback function in the earlier version.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 1000000. **0** means no delay. The value range needs to be extended to 100 million.

**Default value:** 0

---

 **CAUTION**

This parameter can be ignored when you use the Ustore engine to flash back. It serves the Astore flashback function of the earlier version and has other functions. The flashback function is not used in this version.

---

## undo\_retention\_time

**Parameter description:** Specifies the period for retaining undo logs of earlier versions.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 0 to 259200. The unit is second.

**Default value:** 0

---

 **CAUTION**

1. If this parameter is set to **0** during the Ustore flashback query, the snapshot information at the flashback point will be cleared. In earlier versions, no flashback query can be performed. When a flashback query is performed, the error message "cannot find the restore point" is displayed.
  2. If the time within which the undo logs of earlier versions need to be retained is **time1** and the statement execution time for the flashback query is **time2**, you need to set **undo\_retention\_time** to a value greater than the sum of **time1** and **time2**. That is, set **undo\_retention\_time** to a value greater than  $\text{time1} + \text{time2} + 3\text{s}$ . You are advised to set **undo\_retention\_time** to a value equal to  $\text{time1} + 1.5 \times \text{time2}$ . For example, if you want to retain the logs of earlier versions within the latest 3 hours, and the SQL statement execution time for the flashback query is 1 hour, set **undo\_retention\_time** to a value equal to 3 hours + 1.5 x 1 hour, that is, 4.5 hours.
- 

## 17.3.38 Rollback Parameters

### max\_undo\_workers

**Parameter description:** Specifies the number of undo worker threads invoked during asynchronous rollback. The parameter setting takes effect after the system is restarted.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** an integer ranging from 1 to 100

**Default value:** 5

## 17.3.39 AI Features

### enable\_hypo\_index

**Parameter description:** Specifies whether the database optimizer considers the created virtual index when executing the **EXPLAIN** statement. By executing **EXPLAIN** on a specific query statement, you can evaluate whether the index can improve the execution efficiency of the query statement based on the execution plan provided by the optimizer.

This is a USERSET parameter. Set it based on instructions in [Table 17-1](#).

**Value range:** Boolean

- **on** indicates that a virtual index is created during **EXPLAIN** execution.
- **off** indicates that no virtual index is created during **EXPLAIN** execution.

**Default value:** off

## 17.3.40 Global SysCache Parameters

### enable\_global\_syscache

**Parameter description:** Specifies whether to enable the global system cache function. This is a POSTMASTER parameter. Set it based on instructions in [Table 17-2](#).

**Value range:** Boolean

- **on** indicates that the global system cache function is enabled.
- **off** indicates that the global system cache function is disabled.

**Default value:** on

You are advised to use this parameter together with the thread pool parameter. After this parameter is enabled, you are advised to set **wal\_level** of the standby node to **hot\_standby** or higher if you need to access the standby node.

### global\_syscache\_threshold

**Parameter description:** Specifies the maximum memory usage of the global system cache.

This is a SIGHUP parameter. Set it based on instructions in [Table 17-2](#).

The **enable\_global\_syscache** parameter must be enabled.

**Value range:** an integer ranging from 16384 to 1073741824. The unit is KB.

**Default value:** 163840

Recommended calculation formula: The smaller value of the number of hot databases and the number of threads x Memory size allocated to each database.

That is, **global\_syscache\_threshold** = **min(count(hot dbs), count(threads))** x **memofdb**.

The number of hot databases refers to the number of frequently accessed databases. In thread pool mode, the number of threads is the sum of the number of threads in the thread pool and the number of background threads. In non-thread pool mode, the number of hot databases is used.

**memofdb** indicates the average memory allocated to each database. The background noise memory of each database is 2 MB. Each time a table or index is added, 11 KB memory is added.

If this parameter is set to a small value, memory is frequently evicted, and a large number of memory fragments cannot be recycled. As a result, memory control fails.

## 17.3.41 Reserved Parameters

### NOTE

The following parameters are reserved and do not take effect in this version.

acce\_min\_datasize\_per\_thread  
cstore\_insert\_mode  
enable\_constraint\_optimization  
enable\_hadoop\_env  
enable\_hdfs\_predicate\_pushdown  
enable\_orc\_cache  
schedule\_splits\_threshold  
backend\_version  
undo\_zone\_count  
version\_retention\_age  
enable\_incremental\_catchup  
wait\_dummy\_time

## Deprecated Functions

enable\_adio\_debug  
enable\_adio\_function  
enable\_fast\_allocate  
prefetch\_quantity  
backwrite\_quantity  
cstore\_prefetch\_quantity  
cstore\_backwrite\_quantity  
cstore\_backwrite\_max\_threshold  
fast\_extend\_file\_size

effective\_io\_concurrency  
acceleration\_with\_compute\_pool  
enable\_slow\_query\_log  
query\_log\_directory  
query\_log\_file  
hll\_default\_regwidth  
hll\_default\_expthresh  
hll\_default\_sparseon  
hll\_max\_sparse  
enable\_compress\_hll