

GaussDB

Tool Guide(Distributed_V2.0-3.x)

Issue	01
Date	2025-08-19



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 Overview..... 1

2 How to Use gsql..... 10

3 Obtaining Help Information..... 13

4 Command Reference..... 15

5 Meta-Command Reference..... 22

6 Troubleshooting..... 43

1 Overview

Basic Features

- **Connect to the database:** By default, only the local server can be connected. To connect to a remote database, you must configure the server. For details, see "Database Quick Start > Connecting to a Database > Using gsql to Connect to a Database > Remotely Connecting to a Database" in *Developer Guide*.

NOTE

If gsql is used to connect to a database, the connection timeout period will be 5 minutes. If the database has not correctly set up a connection and authenticated the identity of the client within this period, gsql will time out and exit.

To resolve this problem, see [Troubleshooting](#).

- **Run SQL statements:** Interactively entered SQL statements and specified SQL statements in a file can be run.
- **Run meta-commands:** Meta-commands help the administrator view database object information, query cache information, format SQL output, and connect to a new database. For details about meta-commands, see [Meta-Command Reference](#).

Advanced Features

[Table 1-1](#) lists the advanced features of gsql.

Table 1-1 Advanced features of gsql

Feature Name	Description
Variable	gsql provides a variable feature that is similar to the shell command of Linux. The following \set meta-command of gsql can be used to set a variable: <pre>\set varname value</pre> To delete the variables set by the \set command, run the following command: <pre>\unset varname</pre> NOTE • A variable is a simple name-value pair. The value can be any characters in any length. • Variable names must consist of case-sensitive letters (including non-Latin letters), digits, and underscores (_). • If the \set varname meta-command (without the second parameter) is used, the variable is set without a value specified. • If the \set meta-command without parameters is used, values of all variables are displayed. For details about variable examples and descriptions, see Variables.
SQL substitution	Common SQL statements can be set to variables using the variable feature of gsql to simplify operations. For details about examples and descriptions about SQL substitution, see SQL substitution.
Customized prompt	Prompts of gsql can be customized. Prompts can be modified by changing the reserved three variables of gsql: <i>PROMPT1</i>, <i>PROMPT2</i>, and <i>PROMPT3</i>. These variables can be set to customized values or the values predefined by gsql. For details, see Prompt.
Historical client operation records	gsql can record historical client operations. This function is enabled by specifying the -r parameter when a client is connected. The number of historical records can be set using the \set command. For example, \set HISTSIZE 50 indicates that the number of historical records is set to 50. \set HISTSIZE 0 indicates that the operation history is not recorded. NOTE • The default number of historical records is 32. The maximum number of historical records is 500. If interactively entered commands contain Chinese characters, only the UTF-8 encoding environment is supported. • For security reasons, the records containing sensitive words (such as PASSWORD, IDENTIFIED, GS_ENCRYPT_AES128, GS_DECRYPT_AES128, GS_ENCRYPT, GS_DECRYPT, PG_CREATE_PHYSICAL_REPLICATION_SLOT_EXTERN, SECRET_ACCESS_KEY, SECRETKEY, CREATE_CREDENTIAL, ACCESSKEY, and SECRET_KEY) are regarded sensitive and not recorded in historical information. This indicates that you cannot view these records in command output histories.

- Variable

To set a variable, run the `\set` meta-command of `gsql`. For example, to set variable `foo` to `bar`, run the following command:

```
gaussdb=# \set foo bar
```

To reference the value of a variable, add a colon (:) before the variable. For example, to view the value of variable `foo`, run the following command:

```
gaussdb=# \echo :foo  
bar
```

This variable reference method is applicable to regular SQL statements and meta-commands except `\copy`, `\ef`, `\help`, `\sf`, and `\!`.

`gsql` pre-defines some special variables and plans the values of these variables. To ensure compatibility with later versions, do not use these variables for other purposes. For details about special variables, see [Table 1-2](#).

 NOTE

- All the special variables consist of upper-case letters, digits, and underscores (_).
- To view the default value of a special variable, run the `\echo :varname` meta-command, for example, `\echo :DBNAME`.

Table 1-2 Settings of special variables

Variable	Setting Method	Description
DBNAME	<code>\set DBNAME dbname</code>	Name of the connected database. This variable is set again when a database is connected.
ECHO	<code>\set ECHO all queries</code>	- If this variable is set to all, only the query information is displayed. This has the same effect as specifying the -a parameter when <code>gsql</code> is used to connect to a database. - If this variable is set to queries, the command line and query information are displayed. This has the same effect as specifying the -e parameter when <code>gsql</code> is used to connect to a database.

Variable	Setting Method	Description
ECHO_HIDDEN	\set ECHO_HIDDEN on off noexec	When a meta-command (such as \dg) is used to query database information, the value of this variable determines the query behavior. - If this variable is set to on, the query statements that are called by the meta-command are displayed, and then the query result is displayed. This has the same effect as specifying the -E parameter when gsql is used to connect to a database. - If this variable is set to off, only the query result is displayed. - If this variable is set to noexec, only the query information is displayed, and the query is not run.
ENCODING	\set ENCODING <i>encoding</i>	Character set encoding of the current client.
FETCH_COUNT	\set FETCH_COUNT <i>variable</i>	- If the value is an integer greater than 0, for example, <i>n</i>, <i>n</i> lines will be selected from the result set to the cache and displayed on the screen when the SELECT statement is run. - If this variable is not set or set to a value less than or equal to 0, all results are selected at a time to the cache when the SELECT statement is run. NOTE A proper variable value helps reduce the memory usage. The recommended value range is from 100 to 1000.
HISTCONTROL	\set HISTCONTROL ignorespace ignoredups ignoreboth none	ignorespace: A line started with a space is not written to the historical record. ignoredups: A line that exists in the historical record is not written to the historical record. ignoreboth, none, or other values: All the lines read in interaction mode are saved in the historical record. NOTE none indicates that HISTCONTROL is not set.
HISTFILE	\set HISTFILE <i>filename</i>	Specifies the file for storing historical records. The default value is ~/.bash_history .

Variable	Setting Method	Description
HISTSIZE	<code>\set HISTSIZE <i>size</i></code>	Specifies the number of commands to store in the command history. The default value is 500 .
HOST	<code>\set HOST <i>hostname</i></code>	Specifies the name of a connected host.
IGNOREEOF	<code>\set IGNOREEOF <i>variable</i></code>	• If this variable is set to a number, for example, 10, the first nine EOF characters (generally Ctrl+C) entered in gsql are neglected and the gsql program exits when the tenth Ctrl+C is entered. • If this variable is set to a non-numeric value, the default value is 10. • If this variable is deleted, gsql exits when an EOF is entered.
LASTOID	<code>\set LASTOID <i>oid</i></code>	Specifies the last OID, which is the value returned by an INSERT or lo_import command. This variable is valid only before the output of the next SQL statement is displayed.
ON_ERROR_ROLLBACK	<code>\set ON_ERROR_ROLLBACK on interactive off</code>	• If the value is on, an error that may occur in a statement in a transaction block is ignored and the transaction continues. • If the value is interactive, the error is ignored only in an interactive session. • If the value is off (default value), the error triggers the rollback of the transaction block. In on_error_rollback-on mode, a savepoint is set before each statement of a transaction block, and an error triggers the rollback of the transaction block.
ON_ERROR_STOP	<code>\set ON_ERROR_STOP on off</code>	• on: specifies that the execution stops if an error occurs. In interactive mode, gsql returns the output of executed commands immediately. • off (default value): specifies that an error, if occurring during the execution, is ignored, and the execution continues.
PORT	<code>\set PORT <i>port</i></code>	Specifies the port number of a connected database.

Variable	Setting Method	Description
USER	<code>\set USER <i>username</i></code>	Specifies the database user you are currently connected as.
VERBOSITY	<code>\set VERBOSITY terse default verbose</code>	This variable can be set to terse, default, or verbose to control redundant lines in error reports. • terse: Only critical and major error texts and text locations are returned (which is generally suitable for single-line error information). • default: Critical and major error texts and text locations, error details, and error messages (possibly involving multiple lines) are all returned. • verbose: All error information is returned.

- SQL substitution

gsql, like a parameter of a meta-command, provides a key feature that enables you to substitute a standard SQL statement for a gsql variable. gsql also provides a new alias or identifier for the variable. To replace the value of a variable using the SQL substitution method, add a colon (:) before the variable. For example:

```
gaussdb=# \set foo 'HR.areaS'
gaussdb=# select * from :foo;
 area_id |    area_name
-----+-----
      4 | Middle East and Africa
      3 | Asia
      1 | Europe
      2 | Americas
(4 rows)
```

The above command queries the HR.areaS table.

NOTICE

The value of the variable is copied literally, so it can even contain unbalanced quotation marks or backslash commands. Therefore, the input content must be meaningful.

- Prompt

The gsql prompt can be set using the three variables in [Table 1-3](#). These variables consist of characters and special escape characters.

Table 1-3 Prompt variables

Variable	Description	Examples
PROMPT1	Specifies the normal prompt used when gsql requests a new command. The default value of <i>PROMPT1</i> is: %/R%#	<i>PROMPT1</i> can be used to change the prompt. - Change the prompt to [local]: gaussdb=> \set PROMPT1 %M [local:/tmp/gaussdba_mppdb] - Change the prompt to name: gaussdb=> \set PROMPT1 name name - Change the prompt to =: gaussdb=> \set PROMPT1 %R =
PROMPT2	Specifies the prompt displayed when more input is expected because the command that is not terminated with a semicolon (;) or a quote (") is not closed.	<i>PROMPT2</i> can be used to display the prompt. gaussdb=# \set PROMPT2 TEST gaussdb=# select * from HR.areaS TEST; area_id area_name -----+----- 1 Europe 2 Americas 4 Middle East and Africa 3 Asia (4 rows))
PROMPT3	Specifies the prompt displayed when the COPY statement (such as COPY FROM STDIN) is run and data input is expected.	<i>PROMPT3</i> can be used to display the COPY prompt. gaussdb=# \set PROMPT3 '>>>>' gaussdb=# copy HR.areaS from STDIN; Enter data to be copied followed by a newline. End with a backslash and a period on a line by itself. >>>>1 aa >>>>2 bb >>>>\.

The value of the selected prompt variable is printed literally. However, a value containing a percent sign (%) is replaced by the predefined contents depending on the character following the percent sign (%). For details about the defined substitutions, see [Table 1-4](#).

Table 1-4 Defined substitutions

Symbol	Description
%M	Replaced with the full host name (with domain name). The full name is [local] if the connection is over UDS, or [local:/dir/name] if the UDS is not at the compiled default location.
%m	Replaced with the host name truncated at the first dot. It is [local] if the connection is over UDS.

Symbol	Description
%>	Replaced with the number of the port that the host is listening on.
%n	Replaced with the database session username.
%/	Replaced with the name of the current database.
%~	Similar to %/. However, the output is tilde (~) if the database is your default database.
%#	Uses # if the session user is the database administrator. Otherwise, uses >.
%R	• In <i>PROMPT1</i> normally =, but ^ if in single-line mode, or ! if the session is disconnected from the database (which can happen if \connect fails). • In <i>PROMPT2</i> %R is replaced with a hyphen (-), an asterisk (*), a single or double quotation mark, or a dollar sign (\$), depending on whether gsql expects more input or the query is in the /* ... */ the comment, quotation mark, or dollar sign extension.
%x	Replaced with the transaction status. • An empty string when it is not in a transaction block • An asterisk (*) when it is in a transaction block • An exclamation mark (!) when it is in a failed transaction block • A question mark (?) when the transaction status is indefinite (for example, because there is no connection).
%digits	Replaced with the character with the specified byte.
%:name	Replaced with the value of the <i>name</i> variable of gsql.
%command	Replaced with the command output, similar to substitution with the "^" symbol.
%[...%]	Prompts may contain terminal control characters which, for example, change the color, background, or style of the prompt text, or change the title of the terminal window. For example: <pre>gaussdb=> \set PROMPT1 '%[%033[1;33;40m%]%n@%/%R%[%033[0m%]%'</pre> The output is a boldfaced (1;) yellow-on-black (33;40) prompt on VT100-compatible, color-capable terminals.

Environment Variables

Table 1-5 Environment variables related to `gsql`

Name	Description
COLUMNS	If <code>\set columns</code> is set to 0 , this parameter controls the width of the wrapped format. This width determines whether to change the wide output mode into the vertical output mode if automatic expansion is enabled.
PAGER	If the query results do not fit on the screen, they are redirected through this command. You can use the <code>\pset</code> command to disable the pager. Typically, the more or less command is used for viewing the query result page by page. The default is platform-dependent. NOTE Display of the less command is affected by the <code>LC_CTYPE</code> environment variable.
PSQL_EDITOR	The <code>\e</code> and <code>\ef</code> commands use the editor specified by the environment variables. The variables are examined in the order listed. The default editor on Unix is <code>vi</code> .
EDITOR	
VISUAL	
PSQL_EDITOR_LINENUMBER_ARG	When the <code>\e</code> or <code>\ef</code> command is used with a line number parameter, this variable specifies the command-line parameter used to pass the starting line number to the editor. For editors, such as Emacs or <code>vi</code> , this is a plus sign. Include a space in the value of the variable if space is needed between the option name and the line number. For example: <pre>PSQL_EDITOR_LINENUMBER_ARG = '+' PSQL_EDITOR_LINENUMBER_ARG='--line '</pre> A plus sign (+) is used by default on Unix.
PSQLRC	Specifies the location of the user's <code>.gsqlrc</code> file.
SHELL	Has the same effect as the <code>\!</code> command.
TMPDIR	Specifies the directory for storing temporary files. The default value is <code>/tmp</code> .

2 How to Use gsql

Prerequisites

- The user using gsql must have the permission to access the database.
- The gsql version must match the database version.

Background

You can use the **gsql** command to connect to the local database or remote database. When the remote database service is connected, remote connection must be allowed on the remote server. For details, see "Database Quick Start > Connecting to a Database > Using gsql to Connect to a Database > Remotely Connecting to a Database" in *Developer Guide*.

Procedure

Step 1 Connect to the GaussDB server using the **gsql** tool.

The **gsql** tool uses the **-d** parameter to specify the target database name, the **-U** parameter to specify the database username, the **-h** parameter to specify the host name, and the **-p** parameter to specify the port number.

NOTE

If the database name is not specified, the default database name generated during initialization will be used. If the database username is not specified, the current OS username will be used by default. If a variable does not belong to any parameter (such as **-d** and **-U**), and **-d** is not specified, the variable will be used as the database name. If **-d** is specified but **-U** is not specified, the variable will be used as the database username.

Example 1: Connect to the 8000 port of the local **gaussdb** database as user **omm**.

```
gsql -d gaussdb -p 8000
```

Example 2: Connect to the 8000 port of the remote **gaussdb** database as user **jack**.

```
gsql -h 10.180.123.163 -d gaussdb -U jack -p 8000
```

Example 3: If **gaussdb** and **omm** are not parameter values, they are interpreted as the database name and the username, respectively.

```
gsql gaussdb omm -p 8000
```

Equals

```
gsql -d gaussdb -U omm -p 8000
```

For details about the gsql parameters, see [Command Reference](#).

Step 2 Run an SQL statement.

The following takes creating database **human_staff** as an example:

```
CREATE DATABASE human_staff;  
CREATE DATABASE
```

Ordinarily, input lines end when a command-terminating semicolon is reached. If the command is sent and executed without any error, the command output is displayed on the screen.

Step 3 Execute gsql meta-commands.

The following takes all GaussDB databases and description information as an example:

```
gaussdb=# \l
```

Name	Owner	Encoding	Collate	Ctype	Access privileges
human_resource	omm	SQL_ASCII	C	C	
postgres	omm	SQL_ASCII	C	C	
template0	omm	SQL_ASCII	C	C	=c/omm +
					omm=CTc/omm
template1	omm	SQL_ASCII	C	C	=c/omm +
					omm=CTc/omm
human_staff	omm	SQL_ASCII	C	C	

(5 rows)

For details about gsql meta-commands, see [Meta-Command Reference](#).

----End

Example

The example shows how to spread a command over several lines of input. Note the prompt change:

```
gaussdb=# CREATE TABLE HR.areaS(  
gaussdb(# area_ID NUMBER,  
gaussdb(# area_NAME VARCHAR2(25)  
gaussdb(# )tablespace EXAMPLE;  
CREATE TABLE
```

Query the table definition:

```
gaussdb=# \d HR.areaS
```

Column	Type	Modifiers
area_id	numeric	not null
area_name	character varying(25)	

Insert four lines of data into **HR.areaS**.

```
gaussdb=# INSERT INTO HR.areaS (area_ID, area_NAME) VALUES (1, 'Europe');  
INSERT 0 1  
gaussdb=# INSERT INTO HR.areaS (area_ID, area_NAME) VALUES (2, 'Americas');  
INSERT 0 1  
gaussdb=# INSERT INTO HR.areaS (area_ID, area_NAME) VALUES (3, 'Asia');  
INSERT 0 1
```

```
gaussdb=# INSERT INTO HR.areaS (area_ID, area_NAME) VALUES (4, 'Middle East and Africa');
INSERT 0 1
```

Change the prompt.

```
gaussdb=# \set PROMPT1 '%n@%m %~%R%#'
omm@[local] gaussdb=#
```

Query the table:

```
omm@[local] gaussdb=# SELECT * FROM HR.areaS;
 area_id |   area_name
-----+-----
      1 | Europe
      4 | Middle East and Africa
      2 | Americas
      3 | Asia
(4 rows)
```

Use the **\pset** command to display the table in different ways:

```
omm@[local] gaussdb=# \pset border 2
Border style is 2.
omm@[local] gaussdb=# SELECT * FROM HR.areaS;
+-----+-----+
| area_id |   area_name   |
+-----+-----+
|      1 | Europe        |
|      2 | Americas      |
|      3 | Asia          |
|      4 | Middle East and Africa |
+-----+-----+
(4 rows)
omm@[local] gaussdb=# \pset border 0
Border style is 0.
omm@[local] gaussdb=# SELECT * FROM HR.areaS;
 area_id   area_name
-----
      1 Europe
      2 Americas
      3 Asia
      4 Middle East and Africa
(4 rows)
```

Use the meta-command:

```
omm@[local] gaussdb=# \a \t \x
Output format is unaligned.
Showing only tuples.
Expanded display is on.
omm@[local] gaussdb=# SELECT * FROM HR.areaS;
 area_id|2
 area_name|Americas

 area_id|1
 area_name|Europe

 area_id|4
 area_name|Middle East and Africa

 area_id|3
 area_name|Asia
omm@[local] gaussdb=#
```

3 Obtaining Help Information

Procedure

- When connecting to the database, run the following command to obtain the help information:

```
gsql --help
```

The following help information is displayed:

```
.....
Usage:
gsql [OPTION]... [DBNAME [USERNAME]]

General options:
-c, --command=COMMAND  run only single command (SQL or internal) and exit
-d, --dbname=DBNAME    database name to connect to (default: "omm")
-f, --file=<FILE_NAME> execute commands from file, then exit
.....
```

- When connecting to the database, run the following command to obtain the help information:

```
help
```

The following help information is displayed:

```
You are using gsql, the command-line interface to gaussdb.
Type: \copyright for distribution terms
      \h for help with SQL commands
      \? for help with gsql commands
      \g or terminate with semicolon to execute query
      \q to quit
```

Examples

Step 1 Connect to the database.

```
gsql -d gaussdb -p 8000
```

gaussdb is the name of the database, and 8000 is the port number of the CN.

If information similar to the following is displayed, the connection succeeds:

```
gsql ((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.
```

Step 2 View the **gsql** help information. For details, see [Table 3-1](#).

Table 3-1 gsql online help

Description	Examples
Query the copyright.	\copyright
View help information about SQL statements supported by GaussDB.	View help information about SQL statements supported by GaussDB. For example, view all SQL statements supported by GaussDB. <pre>gaussdb=# \h Available help: ABORT ALTER AGGREGATE</pre> For example, view parameters of the CREATE DATABASE command: <pre>gaussdb=# \help CREATE DATABASE Command: CREATE DATABASE Description: create a new database Syntax: CREATE DATABASE database_name [[WITH] {[OWNER [=] user_name]} [TEMPLATE [=] template]} [ENCODING [=] encoding]} [LC_COLLATE [=] lc_collate]} [LC_CTYPE [=] lc_ctype]} [DBCOMPATIBILITY [=] compatibility_type]} [TABLESPACE [=] tablespace_name]} [CONNECTION LIMIT [=] connlimit]}{...} ;</pre>
View the help information about gsql commands.	For example, view commands supported by gsql. <pre>gaussdb=# \? General \copyright show GaussDB Kernel usage and distribution terms \g [FILE] or ; execute query (and send results to file or pipe) \h(\help) [NAME] help on syntax of SQL commands, * for all commands \q quit gsql</pre>

----End

4 Command Reference

For details about gsql parameters, see [Table 4-1](#), [Table 4-2](#), [Table 4-3](#), and [Table 4-4](#).

Table 4-1 Common parameters

Parameter	Description	Value Range
-c, -- command=CO MMAND	Specifies that gsql is to run a string command and then exit.	-
-d, -- dbname=DBNA ME	Specifies the name of the database to connect to. In addition, gsql allows you to use extended DBNAMES, that is, connection strings in the format of 'postgres[ql]://[user[:password]@][netloc][:port][, ...] [/dbname] [?param1=value1&...]' or '[key=value] [...]' as DBNAMES. gsql parses connection information from the connection strings and preferentially uses the information. CAUTION When gsql uses the extended DBNAME to create a connection, the replication parameter cannot be specified.	String
-f, -- file=<FILE_NAM E>	Specifies that files are used as the command source instead of interactively-entered commands. After the files are processed, gsql exits. If <FILE_NAME> is - (hyphen), then standard input is read.	An absolute path or relative path that meets the OS path naming convention
-l, --list	Lists all available databases and then exits.	-

Parameter	Description	Value Range
-v, --set, --variable=NAME =VALUE	Sets gsql variable <i>NAME</i> to <i>VALUE</i> . For details about variable examples and descriptions, see Variables .	-
-X, --no-gsqlrc	Does not read the startup file (neither the system-wide gsqlrc file nor the user's ~/.gsqlrc file). NOTE The startup file is ~/.gsqlrc by default or it can be specified by the environment variable PSQLRC .	-
-1 ("one"), --single-transaction	When gsql uses the -f parameter to execute a script, START TRANSACTION and COMMIT are added to the start and end of the script, respectively, so that the script is executed as one transaction. This ensures that the script is executed successfully. If the script cannot be executed, the script is invalid. NOTE If the script has used START TRANSACTION , COMMIT , or ROLLBACK , this parameter is invalid.	-
-, --help	Displays help information about gsql command parameters, and exits.	-
-V, --version	Prints the gsql version information and exits.	-

Table 4-2 Input and output parameters

Parameter	Description	Value Range
-a, --echo-all	Prints all input lines to standard output as they are read. CAUTION When this parameter is used in some SQL statements, the sensitive information, such as user password, may be disclosed. Use this parameter with caution.	-
-e, --echo-queries	Displays all queries sent to the server to the standard output as well. CAUTION When this parameter is used in some SQL statements, the sensitive information, such as user password, may be disclosed. Use this parameter with caution.	-

Parameter	Description	Value Range
-E, --echo-hidden	Echoes the actual queries generated by \d and other backslash commands.	-
-k, --with-key=KEY	Uses gsql to decrypt imported encrypted files. NOTICE - For key characters, such as the single quotation mark (') or double quotation mark (") in shell commands, Linux shell checks whether the input single quotation mark (') or double quotation mark (") matches. If no match is found, Linux shell does not enter the gsql program until input is complete. - Stored procedures and functions cannot be decrypted and imported.	-
-L, --log-file=<FILE_NAME>	Writes normal output source and all query output into the <FILE_NAME> file. CAUTION - When this parameter is used in some SQL statements, the sensitive information, such as user password, may be disclosed. Use this parameter with caution. - This parameter retains only the query result in the corresponding file, so that the result can be easily found and parsed by other callers (for example, automatic O&M scripts). Logs about gsql operations are not retained.	An absolute path or relative path that meets the OS path naming convention
-m, --maintenance	Allows a cluster to be connected when a two-phase transaction is being restored. NOTE The parameter is for engineers only. When this parameter is used, gsql can be connected to the standby node to check data consistency between the primary and standby nodes.	-
-n, --no-libedit	Closes command line editing.	-
-o, --output=<FILE_NAME>	Puts all query output into the <FILE_NAME> file.	An absolute path or relative path that meets the OS path naming convention
-q, --quiet	Specifies the quiet mode and no additional information will be printed.	By default, gsql displays various information.

Parameter	Description	Value Range
-s, --single-step	Runs in single-step mode. It indicates that the user is prompted before each command is sent to the server. This option can be also used for canceling execution. Use this option to debug scripts. CAUTION When this parameter is used in some SQL statements, the sensitive information, such as user password, may be disclosed. Use this parameter with caution.	-
-S, --single-line	Runs in single-line mode where a line break terminates a command, as a semicolon does.	-

Table 4-3 Output format parameters

Parameter	Description	Value Range
-A, --no-align	Switches to unaligned output mode.	The default output mode is aligned.
-F, --field-separator=STRING	Specifies the field separator. The default is the vertical bar ().	-
-H, --html	Turns on the HTML tabular output.	-
-P, --pset=VAR[=ARG]	Specifies the print option in the \pset format in the command line. NOTE The equal sign (=), instead of the space, is used here to separate the name and value. For example, enter -P format=latex to set the output format to LaTeX .	-
-R, --record-separator=STRING	Sets the record separator.	-
-r	Enables the insert mode on the client.	This function is disabled by default.
-t, --tuples-only	Prints only tuples.	-

Parameter	Description	Value Range
-T, --table-attr=TEXT	Specifies options to be placed within the HTML table tag. Use this parameter with the -H,--html parameter to specify the output to the HTML format.	-
-x, --expanded	Turns on the expanded table formatting mode.	-
-z, --field-separator-zero	Sets the field separator in the unaligned output mode to be blank. Use this parameter with the -A, --no-align parameter to switch to unaligned output mode.	-
-0, --record-separator-zero	Sets the record separator in the unaligned output mode to be blank. Use this parameter with the -A, --no-align parameter to switch to unaligned output mode.	-
-2, --pipeline	Uses a pipe to transmit the password. This parameter cannot be used on devices and must be used together with the -c or -f parameter.	-
-g,	Prints all SQL statements from a file.	-

Table 4-4 Connection parameters

Parameter	Description	Value Range
-h, --host=HOSTNAME	Specifies the host name of the running server, the UDS path, or the domain name. You can specify multiple host addresses by using character strings separated by commas (,). If multiple host addresses are specified, the faulty node is skipped by default and the first available primary node is automatically selected for connection. You can set the <code>PGTARGETSESSIONATTRS</code> environment variable to connect to different types of nodes. The mapping between variables and node types is as follows: A value of the <code>PGTARGETSESSIONATTRS</code> environment variable indicates a type of node to be connected. read-write: readable and writable node. read-only: read-only node. primary or not set: primary node. standby: standby node. prefer-standby: preferred standby node. If there is no standby node, any is used. any: The role is not checked. NOTE Primary CNs and DN are primary nodes. Generally, DNs are read-only nodes. You can set <code>PGTARGETSESSIONATTRS</code> to read-write to connect to CNs.	If the host name is omitted, <code>gsq</code> connects to the server of the localhost over UDS or over TCP/IP to connect to localhost without UDS.
-p, --port=PORT	Specifies the port number of the database server. You can configure one or more port numbers. When one port number is configured, all host addresses use the same port for connection. When multiple port numbers are configured, the sequence is the same as the host address sequence, and the number of port numbers must be the same as the number of host addresses. If they are different, an error is reported. You can modify the default port number using the -p, --port=PORT parameter.	The default port number can be specified by using compilation parameters. If the port number is not specified, the default port number 5432 is used.

Parameter	Description	Value Range
-U, -- username=USE RNAME	Specifies the user who connects to the database. NOTE - If this parameter is specified, you also need to enter your password for identity authentication when connecting to the database. You can enter the password interactively or use the -W parameter to specify a password. - To connect to a database, add an escape character before any dollar sign (\$) in the username.	String. The default user is the current user who operates the system.
-W, -- password=PASS WORD	Specifies a password when the -U parameter is used to connect to the local database or a remote database. NOTE - When you attempt to connect to the CN after you have logged in to the server where the CN resides, the trust connection is used by default, and this parameter does not need to be set. - To connect to a database, add an escape character before any backslash (\) or back quote (') in the password. - If this parameter is not specified but database connection requires your password, you will be prompted to enter your password in interactive mode. The maximum length of the password is 999 bytes, which is restricted by the maximum value of the GUC parameter password_max_length.	String

5 Meta-Command Reference

This section describes meta-commands provided by gsql after the GaussDB database CLI tool is used to connect to a database. A gsql meta-command can be anything that you enter in gsql and begins with an unquoted backslash.

Precautions

- The format of the gsql meta-command is a backslash (\) followed by a command verb, and then a parameter. The parameters are separated from the command verb and from each other by any number of whitespace characters.
- To include whitespace characters into an argument, you must quote them with a single straight quotation mark. To include a single straight quotation mark into such an argument, precede it by a backslash. Anything contained in single quotation marks is furthermore subject to C-like substitutions for \n (new line), \t (tab), \b (backspace), \r (carriage return), \f (form feed), \digits (octal), and \xdigits (hexadecimal).
- Within a parameter, text enclosed in double quotation marks (") is taken as a command line input to the shell. The output of the command (with any trailing newline removed) is taken as the argument value.
- If an unquoted argument begins with a colon (:), the argument is taken as a gsql variable and the value of the variable is used as the argument value instead.
- Some commands take an SQL identifier (such as a table name) as a parameter. These parameters follow the SQL syntax rules: Unquoted letters are forced to lowercase, while double quotation marks (") protect letters from case conversion and allow incorporation of whitespace into the identifier. Within double quotation marks, paired double quotation marks reduce to a single double quotation mark in the result name. For example, **FOO"BAR"BAZ** is interpreted as **fooBARbaz**, and **"Aweird""name"** becomes **A weird"name**.
- Parsing for arguments stops when another unquoted backslash is found. This is taken as the beginning of a new meta-command. The special sequence \\ (two backslashes) marks the end of parameters and continues parsing SQL statements if any. In this way, SQL statements and gsql commands can be freely mixed in a row. But in any case, the arguments of a meta-command cannot continue beyond the end of the line.

Meta-commands

For details about meta-commands, see [Table 5-1](#), [Table 5-2](#), [Table 5-3](#), [Table 5-4](#), [Table 5-6](#), [Table 5-8](#), [Table 5-9](#), [Table 5-10](#), and [Table 5-12](#).

NOTICE

FILE mentioned in the following commands indicates a file path. This path can be an absolute path such as **/home/gauss/file.txt** or a relative path, such as **file.txt**. By default, a **file.txt** is created in the path where the user runs gsql commands.

Table 5-1 Common meta-commands

Parameter	Description	Value Range
\copyright	Displays GaussDB version and copyright information.	-
\g [FILE] or ;	Performs a query operation and sends the result to a file or pipe.	-
\h(\help) [NAME]	Provides syntax help on the specified SQL statement.	If the name is not specified, then gsql will list all the commands for which syntax help is available. If the name is an asterisk (*), syntax help on all SQL statements is displayed.

Parameter	Description	Value Range
<code>\parallel [on [num]] off]</code>	Controls the parallel execution function. • on: The switch is enabled and the maximum number of concurrently executed tasks is num. • off: This switch is disabled. NOTE • Parallel execution is not allowed in a running transaction and a transaction is not allowed to be started during parallel execution. • Parallel execution of \d meta-commands is not allowed. • If SELECT statements are run concurrently, customers can accept the problem that the return results are displayed randomly but they cannot accept it if a core dump or process response failure occurs. • It is not recommended that you run SET statements in concurrent tasks because they may cause unexpected results. • Temporary tables cannot be created in parallel. If temporary tables are required, create them before parallel execution is enabled, and use them only in the parallel execution. Temporary tables cannot be created in parallel execution. • When \parallel is executed, <i>num</i> independent gsql processes can be connected to the database server. • The total duration of all \parallel tasks cannot exceed session_timeout. Otherwise, the connection may fail during concurrent execution.	The default value of <i>num</i> is 1024. NOTICE • The maximum number of connections allowed by the server is determined based on max_connection and the number of current connections. • Set the value of <i>num</i> based on the allowed number of connections.
<code>\q</code>	Exits the gsql program. This command is executed only when a script terminates in a script file.	-

Table 5-2 Query buffer meta-commands

Parameter	Description
<code>\e [FILE] [LINE]</code>	Uses an external editor to edit the query buffer or file.
<code>\ef [FUNCNAME [LINE]]</code>	Edits the function definition using an external editor. If LINE is specified, the cursor will point to the specified line of the function body.

Parameter	Description
\p	Prints the current query buffer to the standard output.
\r	Resets or clears the query buffer.
\w FILE	Outputs the current query buffer to a file.

Table 5-3 Input/Output commands

Parameter	Description
<code>\copy { table [(column_list)] (query) } { from to } { filename stdin stdout pstdin pstdout }[LOAD] [LOAD_DISCARD 'string'] [with] [binary] [oids] [delimiter [as] 'character'] [null [as] 'string'] [useeof] [csv [header] [quote [as] 'character'] [escape [as] 'character'] [force quote column_list *] [force not null column_list]] [parallel integer]</code>	After logging in to the database on any gsql client, you can import and export data. This is an operation of running the SQL COPY command, but not the server that reads or writes data to a specified file. Instead, data is transferred between the server and the local file system. In this way, local user permissions instead of server permissions are required for file access, and the user does not need to be the initial user. NOTE • \COPY only applies to small-batch data import with uniform formats. GDS or COPY is preferred for data import. • \COPY specifies the number of clients to import data to implement parallel import of data files. Currently, the value ranges from 1 to 8. • The parallel import using \COPY has the following constraints: Parallel import of temporary tables is not supported. Parallel import within transactions is not supported. Parallel import of binary files is not supported. Parallel import of data encrypted using AES-128 is not supported. The COPY option contains EOL. In these cases, even if the parallel parameter is specified, a non-parallel process is performed. • Both the TEXT and CSV formats of \COPY support the header function. • The LOAD function is used by gs_loader to call COPY after syntax conversion. It is not an active calling function. • The LOAD_DISCARD function is used by gs_loader to discard file path after parsing. It is not an active calling function.
<code>\echo [STRING]</code>	Writes character strings to the standard output.

Parameter	Description
\i FILE	Reads content from <i>FILE</i> and uses them as the input for a query.
\i+ FILE KEY	Runs commands in an encrypted file.
\ir FILE	Similar to \i, but resolves relative path names differently.
\ir+ FILE KEY	Similar to \i+, but resolves relative path names differently.
\o [FILE]	Saves all query results to a file.
\qecho [STRING]	Writes character strings to the query output flow.

 NOTE

In [Table 5-4](#), **S** indicates displaying the system object and **+** indicates displaying the additional description information of the object. **PATTERN** specifies the name of an object to be displayed.

Table 5-4 Information display meta-commands

Parameter	Description	Value Range	Examples
\d[S+]	Lists all tables, views, and sequences of all schemas in the search_path. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	-	Lists all tables, views, and sequences of all schemas in search_path . gaussdb=# \d
\d[S+] NAME	Lists the structure of specified tables, views, and indexes. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	-	Lists the structure of table a . gaussdb=# \dtable+ a

Parameter	Description	Value Range	Examples
\d+[PATTERN]	Lists all tables, views, and indexes. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only tables, views, and indexes whose names match PATTERN are displayed.	Lists all tables, views, and indexes whose names start with f . gaussdb=# \d+ f*
\da[S][PATTERN]	Lists all available aggregate functions, together with the data type they perform operations on and the return value types. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only aggregate functions whose names match PATTERN are displayed.	Lists all available aggregate functions whose names start with f , together with their return value types and the data types. gaussdb=# \da f*
\db[+][PATTERN]	Lists all available tablespaces.	If PATTERN is specified, only tablespaces whose names match PATTERN are displayed.	Lists all available tablespaces whose names start with p . gaussdb=# \db p*
\dc[S+][PATTERN]	Lists all available conversions between character-set encodings. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only conversions whose names match PATTERN are displayed.	Lists all available conversions between character-set encodings. gaussdb=# \dc *
\dC[+][PATTERN]	Lists all available type conversions. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed. PATTERN must be the actual type name and cannot be an alias.	If PATTERN is specified, only conversions whose names match PATTERN are displayed.	Lists all type conversions whose pattern names start with c . gaussdb=# \dC c*

Parameter	Description	Value Range	Examples
\dd[S] [PATTERN]	Lists descriptions about objects matching PATTERN . When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If no parameter is specified, all visible objects are displayed. The objects include aggregations, functions, operators, types, relations (tables, views, indexes, sequences, and large objects), and rules.	Lists all visible objects. gaussdb=# \dd
\ddp [PATTERN]	Lists all default permissions.	If PATTERN is specified, only permissions whose names match PATTERN are displayed.	Lists all default permissions. gaussdb=# \ddp
\dD[S+] [PATTERN]	Lists all available domains. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only domains whose names match PATTERN are displayed.	Lists all available domains. gaussdb=# \dD
\det[+] [PATTERN]	Lists all external tables. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only tables whose names match PATTERN are displayed.	Lists all external tables. gaussdb=# \det
\des[+] [PATTERN]	Lists all external servers.	If PATTERN is specified, only servers whose names match PATTERN are displayed.	Lists all external servers. gaussdb=# \des

Parameter	Description	Value Range	Examples
\deu[+] [PATTERN]	Lists all user mappings.	If PATTERN is specified, only information whose name matches PATTERN is displayed.	Lists all user mappings. gaussdb=# \deu
\dew[+] [PATTERN]	Lists all encapsulated external data.	If PATTERN is specified, only data whose name matches PATTERN is displayed.	Lists all encapsulated external data. gaussdb=# \dew
\df[antw][S+] [PATTERN]	Lists all available functions, together with their parameters and return types. a indicates an aggregate function, n indicates a common function, t indicates a trigger, and w indicates a window function. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only functions whose names match PATTERN are displayed.	Lists all available functions, together with their parameters and return types. gaussdb=# \df
\dF[+] [PATTERN]	Lists all text search configuration information. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only configurations whose names match PATTERN are displayed.	Lists all text search configuration information. gaussdb=# \dF+
\dFd[+] [PATTERN]	Lists all text search dictionaries. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only dictionaries whose names match PATTERN are displayed.	Lists all text search dictionaries. gaussdb=# \dFd

Parameter	Description	Value Range	Examples
<code>\dFp[+]</code> <code>[PATTERN]</code>	Lists all text search analyzers. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only analyzers whose names match PATTERN are displayed.	Lists all text search analyzers. <code>gaussdb=# \dFp</code>
<code>\dFt[+]</code> <code>[PATTERN]</code>	Lists all text search templates. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only templates whose names match PATTERN are displayed.	Lists all text search templates. <code>gaussdb=# \dFt</code>
<code>\dg[+]</code> <code>[PATTERN]</code>	Lists all database roles. NOTE Since the concepts of "users" and "groups" have been unified into "roles", this command is now equivalent to <code>\du</code> . Both commands are retained to ensure compatibility with earlier versions.	If PATTERN is specified, only roles whose names match PATTERN are displayed.	Lists all database roles named <code>j?e</code> (the question mark (?) indicates any character). <code>gaussdb=# \dg j?e</code>
<code>\dl</code>	This is an alias for <code>\lo_list</code> , which shows a list of large objects.	-	Lists all large objects. <code>gaussdb=# \dl</code>
<code>\dL[S+]</code> <code>[PATTERN]</code>	Lists all available program languages.	If PATTERN is specified, only languages whose names match PATTERN are displayed.	Lists all available program languages. <code>gaussdb=# \dL</code>
<code>\dm[S+]</code> <code>[PATTERN]</code>	Lists materialized views. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only materialized views whose names match PATTERN are displayed.	Lists materialized views. <code>gaussdb=# \dm</code>

Parameter	Description	Value Range	Examples
<code>\dn[S+] [PATTERN]</code>	Lists all schemas (namespace). If + is added to the command, the permission and description of each schema are listed.	If PATTERN is specified, only schemas whose names match the pattern are shown. By default, only schemas you created are displayed.	Lists information about all schemas whose names start with d . gaussdb=# \dn+ d*
<code>\do[S] [PATTERN]</code>	Lists available operators with their operand and return types. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only operators whose names match PATTERN are displayed. By default, only the operators created by the user are listed.	Lists available operators with their operand and return types. gaussdb=# \do
<code>\dO[S+] [PATTERN]</code>	Lists collation rules. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only rules whose names match PATTERN are displayed. By default, only user-created rules are shown.	Lists collation rules. gaussdb=# \dO
<code>\dp [PATTERN]</code>	Lists tables, views, and related permissions. When objects with the same name exist in different schemas in search_path, only the object in the schema that ranks first in search_path is displayed. The following result about \dp is displayed: <pre>rolename=xxxx/yyyy -- Assigns permissions to a role. =xxxx/yyyy -- Assigns permissions to public.</pre> <i>xxxx</i> indicates the assigned permissions, and <i>yyyy</i> indicates the roles with the assigned permissions. For details about permission descriptions, see Table 5-5.	If PATTERN is specified, only tables and views whose names match the pattern are shown.	Lists tables, views, and related permissions. gaussdb=# \dp

Parameter	Description	Value Range	Examples
<code>\drds</code> [PATTERN1 [PATTERN2]]	Lists all parameters that have been modified. These settings can be for roles, for databases, or for both. PATTERN1 and PATTERN2 indicate a role pattern and a database pattern, respectively.	If PATTERN is specified, only collations rules whose names match PATTERN are displayed. If the default value is used or * is specified, all settings are listed.	List all the modified configuration parameters of the database. gaussdb=# \drds * dbname
<code>\dT[S+]</code> [PATTERN N]	Lists all data types. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed.	If PATTERN is specified, only types whose names match PATTERN are displayed.	Lists all data types. gaussdb=# \dT
<code>\du[+]</code> [PATTERN N]	Lists all database roles. NOTE Since the concepts of "users" and "groups" have been unified into "roles", this command is now equivalent to <code>\dg</code> . Both commands are retained to ensure compatibility with earlier versions.	If PATTERN is specified, only roles whose names match PATTERN are displayed.	Lists all database roles. gaussdb=# \du
<code>\dE[S+]</code> [PATTERN N] <code>\di[S+]</code> [PATTERN N] <code>\ds[S+]</code> [PATTERN N] <code>\dt[S+]</code> [PATTERN N] <code>\dv[S+]</code> [PATTERN N]	In this group of commands, the letters E, i, s, t, and v stand for foreign table, index, sequence, table, and view, respectively. You can specify any or a combination of these letters sequenced in any order to obtain an object list. When objects with the same name exist in different schemas in search_path , only the object in the schema that ranks first in search_path is displayed. For example, <code>\dit</code> lists all indexes and tables. If + is added to the end of a command name, the physical size and related description of each object are also listed.	If PATTERN is specified, only objects whose names match PATTERN are displayed. By default, only objects you created are displayed. You can specify PATTERN or S to view other system objects.	Lists all indexes and views. gaussdb=# \div

Parameter	Description	Value Range	Examples
\dx[+] [PATTERN]	Lists installed extensions.	If PATTERN is specified, only extensions whose names match PATTERN are displayed.	Lists installed extensions. gaussdb=# \dx
\l[+]	Lists the names, owners, character set encodings, and permissions of all the databases in the server.	-	List the names, owners, character set encodings, and permissions of all the databases in the server. gaussdb=# \l
\sf[+] FUNCTION NAME	Displays the definition of a function. NOTE If the function name contains parentheses, enclose the function name with double quotation marks and add the parameter type list following the double quotation marks. Also enclose the list with parentheses.	-	Assume a function function_a and a function func()name . This parameter will be as follows: gaussdb=# \sf function_a gaussdb=# \sf "func()name"(argtype1, argtype2)
\z [PATTERN]	Lists all tables, views, and sequences in the database and their access permissions.	If a pattern is given, it is a regular expression, and only matched tables, views, and sequences are shown.	Lists all tables, views, and sequences in the database and their access permissions. gaussdb=# \z

Table 5-5 Description of permissions

Parameter	Description
r	SELECT: allows users to read data from specified tables and views.
w	UPDATE: allows users to update columns for specified tables.
a	INSERT: allows users to insert data to specified tables.
d	DELETE: allows users to delete data from specified tables.

Parameter	Description
D	TRUNCATE: allows users to delete all data from specified tables.
x	REFERENCES: allows users to create a foreign key constraint (not supported in distributed mode).
t	TRIGGER: allows users to create a trigger on specified tables.
X	EXECUTE: allows users to use specified functions and the operators that are realized by the functions.
U	USAGE: • For procedural languages, allows users to specify a procedural language when creating a function. • For schemas, allows users to access objects included in specified schemas. • For sequences, allows users to use the nextval function.
C	CREATE: • For databases, allows new schemas to be created within the database. • For schemas, allows users to create objects in a schema. • For tablespaces, allows users to create tables in a tablespace and set the tablespace to default one when creating databases and schemas.
c	CONNECT: allows users to connect to specified databases.
T	TEMPORARY: allows users to create temporary tables.
A	ALTER: allows users to modify the attributes of a specified object.
P	DROP: allows users to delete specified objects.
m	COMMENT: allows users to define or modify comments of a specified object.
i	INDEX: allows users to create indexes on specified tables.
v	VACUUM: allows users to perform ANALYZE and VACUUM operations on specified tables.
*	Authorization options for preceding permissions.

Table 5-6 Formatting meta-commands

Parameter	Description
\a	Switches between aligned and unaligned table output formats.

Parameter	Description
\C [STRING]	Sets the title of any table being printed as the result of a query or unsets any such title.
\f [STRING]	Sets the field separator for unaligned query outputs.
\H	- If the text format schema is used, switches to the HTML format. - If the HTML format schema is used, switches to the text format.
\pset NAME [VALUE]	Sets options affecting the output of query result tables. For details about the value of NAME , see Table 5-7 .
\t [on off]	Switches the display of output name information and row count footer.
\T [STRING]	Specifies attributes to be placed within the table tag in HTML output format. If this parameter is empty, no attribute is specified.
\x [on off auto]	Switches expanded table formatting mode.

Table 5-7 Adjustable printing options

Option	Description	Value Range
border	The value must be a number. In general, the larger the number, the more borders and lines the tables will have, but this depends on the particular format.	- The value is an integer greater than 0 in HTML format. - The value range in other formats is as follows: 0: no border 1: internal dividing line 2: table frame

Option	Description	Value Range
expanded (or x)	Switches between regular and expanded formats.	- When the expanded format is enabled, query results are displayed in two columns, with the column name on the left and the data on the right. This mode is useful if the data does not fit on the screen in the normal "horizontal" mode. - Use the expanded format when the query output format is wider than the screen in regular format. The regular format is effective only in the aligned and wrapped formats.
fieldsep	Specifies the field separator to be used in unaligned output mode. In this way, you can create tab- or comma-separated output required by other programs. To set a tab as field separator, type \pset fieldsep '\t' . The default field separator is a vertical bar (' ').	-
fieldsep_zero	Sets the field separator to use in unaligned output format to a zero byte.	-
footer	Switches the display of the default footer.	-

Option	Description	Value Range
format	Selects the output format. Unique abbreviations are allowed (this indicates that one letter is enough).	Value range: • unaligned: Write all columns of a row on one line, separated by the currently active column separator. • aligned: This format is standard and human-readable. • wrapped: This format is similar to aligned, but includes the packaging cross-line width data value to suit the width of the target field output. • html: This format outputs tables to the markup language for a document. The output is not a complete document. • latex: This format outputs tables to the markup language for a document. The output is not a complete document. • troff-ms: This format outputs tables to the markup language for a document. The output is not a complete document.
null	Sets a character string to be printed in place of a null value.	The default is to print nothing, which can be easily mistaken for an empty string.
numericlocale	Switches the display of a locale-aware character to separate groups of digits to the left of the decimal marker.	• on: The specified separator is displayed. • off: The specified separator is not displayed If this parameter is ignored, the default separator is displayed.

Option	Description	Value Range
pager	Controls the use of a pager for query and gsql help outputs. If the PAGER environment variable is set, the output is redirected to the specified program. Otherwise, the platform-dependent default value is used.	• on: The pager is used for terminal output that does not fit the screen. • off: The pager is not used. • always: The pager is used for all terminal output regardless of whether it fits the screen.
recordsep	Specifies the record separator to use in unaligned output mode.	-
recordsep_zero	Sets the record separator to use in unaligned output format to a zero byte.	-
tableattr (or T)	Specifies attributes to be placed inside the HTML table tag in HTML output format (such as cellpadding or bgcolor). Note that you do not need to specify border here because it has been used by \pset border . If no value is given, the table attributes do not need to be set.	-
title	Sets the table title for any subsequently printed tables. This can be used to give your output descriptive tags. If no value is given, the title does not need to be set.	-
tuples_only (or t)	Enables or disables the tuples-only mode. Full display may show extra information, such as column headers, titles, and various footers. In tuples-only mode, only the table data is shown.	-
feedback	Specifies whether to output the number of result lines.	-

Table 5-8 Connection meta-commands

Parameter	Description	Value Range
\c[onnect] [DBNAME - USER - HOST - PORT -]	Connects to a new database. If a database name contains more than 63 bytes, only the first 63 bytes are valid and are used for connection. However, the database name displayed in the command line of <code>gsq</code> is still the name before the truncation. NOTE If the database login user is changed during reconnection, you need to enter the password of the new user. The maximum length of the password is 999 bytes, which is restricted by the maximum value of the GUC parameter <code>password_max_length</code> .	-
\encoding [ENCODING]	Sets the client character set encoding.	Without an argument, this command shows the current encoding.
\conninfo	Prints information about the current connected database.	-

Table 5-9 OS meta-commands

Parameter	Description	Value Range
\cd [DIR]	Changes the current working directory.	An absolute path or relative path that meets the OS path naming convention
\setenv NAME [VALUE]	Sets the NAME environment variable to VALUE . If VALUE is not provided, do not set the environment variable.	-
\timing [on off]	Toggles a display of how long each SQL statement takes, in milliseconds (exclude the displayed printing time).	• on: specifies that the display is enabled. • off: indicates that the display is disabled.
\! [COMMAND]	Escapes to a separate Unix shell or runs a Unix command.	-

Table 5-10 Variable meta-commands

Parameter	Description
\prompt [TEXT] NAME	Prompts the user to use texts to specify a variable name.
\set [NAME [VALUE]]	Sets the <i>NAME</i> internal variable to VALUE. If more than one value is provided, <i>NAME</i> is set to the concatenation of all of them. If no second argument is given, the variable is just set with no value. Some common variables are processed differently in gsql and they are combinations of uppercase letters, numbers and underscores. Table 5-11 describes a list of variables that are processed in a way different from other variables.
\unset NAME	Deletes the variable name of gsql.

Table 5-11 Common \set commands

Command	Command Description	Value Range
\set VERBOSITY value	This variable can be set to default , verbose , or terse to control redundant lines in error reports.	Value range: default , verbose , terse
\set ON_ERROR_STOP value	If this variable is set, the script execution stops immediately. If this script is called from another script, that script will be stopped immediately as well. If the primary script is called using the -f option rather than from one gsql session, gsql will return error code 3 , indicating the difference between the current error and critical errors. (The error code for critical errors is 1 .)	Value range: on/off , true/false , yes/no , and 1/0

Command	Command Description	Value Range
<code>\set AUTOCOMMIT [on off]</code>	Sets the auto commit behavior of the current gsql connection. on indicates that auto commit is enabled, and off indicates that auto commit is disabled. By default, the gsql connection is automatically committed, and each individual statement is implicitly committed. If auto commit is disabled for performance or other purposes, you need to explicitly run the COMMIT command to ensure that transactions are committed. For example, execute the COMMIT statement to explicitly commit transactions after a specified service SQL statement is executed. Particularly, ensure that all transactions are committed before the gsql client exits. NOTE The auto commit is enabled in gsql by default. If you disable it, all the statements executed later will be packaged in implicit transactions, and you cannot execute statements that cannot be executed within transactions.	• on: The auto commit is enabled. • off: The auto commit is disabled.

Table 5-12 Large object meta-commands

Parameter	Description
<code>\lo_list</code>	Displays a list of all GaussDB large objects stored in the database, along with the comments provided for them.

PATTERN

The various `\d` commands accept a **PATTERN** parameter to specify the object name to be displayed. In the simplest case, PATTERN is the exact name of the object. Characters in **PATTERN** are usually converted to lowercase (as in SQL names), for example, `\dt FOO` will display a table named **foo**. As in SQL names, placing double quotation marks (") around a pattern prevents them being folded to lower case. If you need to include a double quotation mark (") in a pattern, write it as a pair of double quotation marks (") within a double-quote sequence, which is in accordance with the rules for SQL quoted identifiers. For example, `\dt "FOO""BAR"` will be displayed as a table named **FOO"BAR** instead of **foo"bar**. You cannot put double quotation marks around just part of a pattern, which is different from the normal rules for SQL names. For example, `\dt FOO"FOO"BAR` will be displayed as a table named **fooFOObar** if just part of a pattern is quoted.

Whenever the **PATTERN** parameter is omitted completely, the `\d` commands display all objects that are visible in the current schema search path, which is equivalent to using an asterisk (*) as the pattern. An object is regarded to be visible if it can be referenced by name without explicit schema qualification. To see

all objects in the database regardless of their visibility, use a dot within double quotation marks (".") as the pattern.

Within a pattern, the asterisk (*) matches any sequence of characters (including no characters) and a question mark (?) matches any single character. This notation is comparable to Unix shell file name patterns. For example, `\dt int*` displays tables whose names start with **int**. But within double quotation marks, the asterisk (*) and the question mark (?) lose these special meanings and are just matched literally.

A pattern that contains a dot (.) is interpreted as a schema name pattern followed by an object name pattern. For example, `\dt foo*.bar*` displays all tables (whose names include **bar**) in schemas starting with **foo**. If no dot appears, then the pattern matches only visible objects in the current schema search path. Likewise, the dot within double quotation marks loses its special meaning and becomes an ordinary character.

Senior users can use regular-expression notations, such as character classes. For example [0-9] can be used to match any digit. All regular-expression special characters work as specified in POSIX. The following characters are excluded:

- A dot (.) is used as a separator.
- An asterisk (*) is translated into an asterisk prefixed with a dot (.*), which is a regular-expression marking.
- A question mark (?) is translated into a dot (.).
- A dollar sign (\$) is matched literally.

You can write `?`, `(R+|)`, `(R|)`, and `R` to the following pattern characters: `.`, `R*`, and `R?`. The dollar sign (\$) does not need to be used as a regular expression character because **PATTERN** must match the entire name instead of being interpreted as a regular expression (in other words, \$ is automatically appended to **PATTERN**). If you do not expect a pattern to be anchored, write an asterisk (*) at its beginning or end. All regular-expression special characters within double quotation marks lose their special meanings and are matched literally. Regular-expression special characters in operator name patterns (such as the `\do` parameter) are also matched literally.

6 Troubleshooting

Low Connection Performance

- **log_hostname** is enabled, but DNS is incorrect.

Connect to the database, and run **show log_hostname** to check whether **log_hostname** is enabled in the database.

If it is enabled, the database kernel will use DNS to check the name of the host where the client is deployed. If the host where the CN resides database is configured with an incorrect or unreachable DNS server, the database connection will take a long time to set up. For more details about **log_hostname**, see the section "GUC parameters".

- The database kernel slowly runs the initialization statement.

Problems are difficult to locate in this scenario. Try using the **strace** Linux trace command.

```
strace gsql -U MyUserName -d gaussdb -h 127.0.0.1 -p 23508 -r -c 'q'  
Password for MyUserName:
```

The database connection process will be printed on the screen. If the following statement takes a long time to run:

```
sendto(3, "Q\0\0\0\25SELECT VERSION()\0", 22, MSG_NOSIGNAL, NULL, 0) = 22  
poll([fd=3, events=POLLIN|POLLERR], 1, -1) = 1 ([fd=3, revents=POLLIN])
```

It can be determined that the database executes the **SELECT VERSION()** statement slowly.

After the database is connected, you can run the **explain performance select version()** statement to find the reason why the initialization statement was run slowly. For more information, see "SQL Optimization > Introduction to the SQL Execution Plan" in *Developer Guide*.

An uncommon scenario is that the disk of the machine where the database CN resides is full or faulty, affecting queries and leading to user authentication failures. As a result, the connection process is suspended. To solve this problem, simply clear the data disk space of the database CN.

- TCP connection is set up slowly.

Adapt the steps of troubleshooting slow initialization statement execution. Use **strace**. If the following statement is run slowly:

```
connect(3, {sa_family=AF_FILE, path="/home/test/tmp/ gaussdb_llt1/s.PGSQL.61052"}, 110) = 0
```

Or,

```
connect(3, {sa_family=AF_INET, sin_port=htons(61052), sin_addr=inet_addr("127.0.0.1")}, 16) = -1  
EINPROGRESS (Operation now in progress)
```

It indicates that the physical connection between the client and the database is set up slowly. In this case, check whether the network is unstable or has high throughput.

- The connection is slow because the resource load is full.

Cause analysis: When the CPU, memory, or I/O load is close to 100%, the gsql connection is slow.

Solution:

- a. Run the **top** command to check the CPU usage. Run the **free** command to check the memory usage. Run the **iostat** command to check the I/O load. You can also check the monitor logs in the CM Agent and the monitoring records on the database O&M platform.
- b. 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 [*Port number of the database server* + 1] to query the pg_stat_activity view. For slow queries, you can use the system function pg_terminate_backend to kill sessions.
- c. 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.

Problems in Setting Up Connections

- gsql: could not connect to server: No route to host
This problem occurs generally because an unreachable IP address or port number was specified. Check whether the values of **-h** and **-p** parameters are correct.
- gsql: FATAL: Invalid username/password,login denied.
This problem occurs generally because an incorrect username or password was entered. Contact the database administrator to check whether the username and password are correct.
- gsql: FATAL: Forbid remote connection with trust method!
For security purposes, remote login in trust mode is forbidden. In this case, you need to modify the connection authentication information in the **pg_hba.conf** file. Contact an administrator

NOTE

Do not modify the configurations of database cluster hosts in the **pg_hba.conf** file. Otherwise, the database may become faulty. It is recommended that service applications be deployed outside the database cluster.

- The the host where the CN resides can connect to the database if **-h 127.0.0.1** is specified, and the connection will fail if **-h 127.0.0.1** is removed.
Run the SQL statement **show unix_socket_directory** to check whether the **unix socket directory** used by the database CN is the same as that specified by the environment variable **\$PGHOST** in the **shell** directory.
If they are different, set **\$PGHOST** to the directory specified by **unix_socket_directory**.

- The "libpq.so" loaded mismatch the version of gsql, please check it.
This problem occurs because the version of **libpq.so** used in the environment does not match that of **gsql**. Run the **ldd gsql** command to check the version of the loaded **libpq.so**, and then load correct **libpq.so** by modifying the environment variable **LD_LIBRARY_PATH**.

 NOTE

Modify the **LD_LIBRARY_PATH** environment variable by referring to the example command provided below. In this example, **\${path_to_correct_libpq_dir}** indicates the directory where **libpq.so** is located.

```
export LD_LIBRARY_PATH=${path_to_correct_libpq_dir}:$LD_LIBRARY_PATH
```

- gsql: symbol lookup error: xxx/gsql: undefined symbol: libpqVersionString
This problem occurs because the version of **libpq.so** used in the environment does not match that of **gsql** (or the PostgreSQL **libpq.so** exists in the environment). Run the **ldd gsql** command to check the version of the loaded **libpq.so**, and then load correct **libpq.so** by modifying the environment variable **LD_LIBRARY_PATH**.

- gsql: connect to server failed: Connection timed out
Is the server running on host "xx.xxx.xxx.xxx" and accepting TCP/IP connections on port xxxx?
This problem is caused by network connection faults. Check the network connection between the client and the database server. If you cannot ping from the client to the database server, the network connection is abnormal. Contact network management personnel for troubleshooting.

```
ping -c 4 10.10.10.1
PING 10.10.10.1 (10.10.10.1) 56(84) bytes of data.
From 10.10.10.1: icmp_seq=2 Destination Host Unreachable
From 10.10.10.1 icmp_seq=2 Destination Host Unreachable
From 10.10.10.1 icmp_seq=3 Destination Host Unreachable
From 10.10.10.1 icmp_seq=4 Destination Host Unreachable
--- 10.10.10.1 ping statistics ---
4 packets transmitted, 0 received, +4 errors, 100% packet loss, time 2999ms
```

- gsql: FATAL: permission denied for database "gaussdb"
DETAIL: User does not have CONNECT privilege.
This problem occurs because the user does not have the permission to access the database. To solve this problem, perform the following steps:

- a. Connect to the database as the system administrator **dbadmin**.
gsql -d gaussdb -U dbadmin -p 8000
- b. Grant the user with the permission to access the database.
GRANT CONNECT ON DATABASE gaussdb TO user1;

 NOTE

Actually, some common misoperations may also cause a database connection failure, for example, entering an incorrect database name, username, or password. In this case, the client tool will display the corresponding error messages.

```
gsql -d gaussdb -p 8000
gsql: FATAL: database "gaussdb" does not exist
```

```
gsql -d gaussdb -U user1 -p 8000
Password for user user1:
gsql: FATAL: Invalid username/password,login denied.
```

- gsql: FATAL: sorry, too many clients already, active/non-active: 197/3.

This problem occurs because the number of system connections exceeds the allowed maximum. Contact the DBA database administrator to release unnecessary sessions.

You can check the number of connections as described in [Table 6-1](#).

You can view the session status in the **PG_STAT_ACTIVITY** view. To release unnecessary sessions, use the **pg_terminate_backend** function.

```
select datid,pid,state from pg_stat_activity;  
datid |      pid      | state  
-----+-----  
13205 | 139834762094352 | active  
13205 | 139834759993104 | idle  
(2 rows)
```

The value of **pid** is the thread ID of the session. Terminate the session using its thread ID.

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

If a command output similar to the following is displayed, the session is successfully terminated:

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

Table 6-1 Viewing the number of session connections

Description	Command
View the maximum number of sessions connected to a specific user.	Run the following command to view the upper limit of the number of USER1 's session connections, where -1 indicates that no upper limit is set for the number of USER1 's session connections: <pre>SELECT ROLNAME,ROLCONNLIMIT FROM PG_ROLES WHERE ROLNAME='user1'; rolname rolconnlimit -----+----- user1 -1 (1 row)</pre>
View the number of session connections that have been used by a specified user.	Run the following command to view the number of session connections that have been used by USER1 , where 1 indicates the number of session connections that have been used by USER1 : <pre>SELECT COUNT(*) FROM dv_sessions WHERE USERNAME='user1'; count ----- 1 (1 row)</pre>

Description	Command				
View the maximum number of sessions connected to a specific database.	Run the following command to view the upper limit of the number of gaussdb's session connections, where -1 indicates that no upper limit is set for the number of gaussdb's session connections: <pre>SELECT DATNAME,DATCONNLIMIT FROM PG_DATABASE WHERE DATNAME='gaussdb';</pre> <table><tr><th>datname</th><th>datconnlimit</th></tr><tr><td>gaussdb</td><td>-1</td></tr></table> (1 row)	datname	datconnlimit	gaussdb	-1
datname	datconnlimit				
gaussdb	-1				
View the number of session connections that have been used by a specific database.	Run the following command to view the number of session connections that have been used by gaussdb, where 1 indicates the number of session connections that have been used by gaussdb: <pre>SELECT COUNT(*) FROM PG_STAT_ACTIVITY WHERE DATNAME='gaussdb';</pre> <table><tr><th>count</th></tr><tr><td>1</td></tr></table> (1 row)	count	1		
count					
1					
View the number of session connections that have been used by all users.	Run the following command to view the number of session connections that have been used by all users: <pre>SELECT COUNT(*) FROM dv_sessions;</pre> <table><tr><th>count</th></tr><tr><td>10</td></tr></table> (1 row)	count	10		
count					
10					

- gsql: wait xxx.xxx.xxx.xxx:xxxx timeout expired

When **gsql** initiates a connection request to the database, a 5-minute timeout period is used. If the database cannot correctly authenticate the client request and client identity within this period, **gsql** will exit the connection process for the current session, and will report the above error.

Generally, this problem is caused by the incorrect host and port (that is, the xxx part in the error information) specified by the **-h** and **-p** parameters. As a result, the communication fails. Occasionally, this problem is caused by network faults. To resolve this problem, check whether the host name and port number of the database are correct.

- gsql: could not receive data from server: Connection reset by peer.

Check whether CN logs contain information similar to "FATAL: cipher file "/data/coordinator/server.key.cipher" has group or world access". This error is usually caused by incorrect tampering with the permissions for data directories or some key files. For details about how to correct the permissions, see related permissions for files on other normal instances.

- gsql: FATAL: GSS authentication method is not allowed because XXXX user password is not disabled.

In **pg_hba.conf** of the target CN, 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. Contact an administrator

 NOTE

- Do not modify the configurations of database cluster hosts in the **pg_hba.conf** file. Otherwise, the database may become faulty.
- It is recommended that service applications be deployed outside database cluster.

Other Faults

- There is a core dump or abnormal exit due to the bus error.

Generally, this problem is caused by changes in loading the shared dynamic library (.so file in Linux) during process running. Alternatively, if the process binary file changes, the execution code for the OS to load machines or the entry for loading a dependent library will change accordingly. In this case, the OS kills the process for protection purposes, generating a core dump file.

To resolve this problem, try again. In addition, do not run service programs in a cluster during O&M operations, such as an upgrade, preventing such a problem caused by file replacement during the upgrade.

 NOTE

A possible stack of the core dump file contains `dl_main` and its function calling. The file is used by the OS to initialize a process and load the shared dynamic library. If the process has been initialized but the shared dynamic library has not been loaded, the process cannot be considered completely started.