

**openGauss**  
**3.0.0**

# **Glossary**

<b>Issue</b>	01
<b>Date</b>	2022-03-31



**Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.**

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

## **Trademarks and Permissions**



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

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

## **Notice**

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

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

## **Huawei Technologies Co., Ltd.**

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

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

Email: [support@huawei.com](mailto:support@huawei.com)

---

# Contents

1 Glossary.....1

# 1 Glossary

**Table 1-1** Glossary

Terms	Description
<b>A - E</b>	
ACID	Atomicity, consistency, isolation, and durability (ACID) is a set of properties that guarantee that transactions are processed reliably in a database management system (DBMS).
AZ	Acronym for available zone which usually refers to an equipment room.
bgwriter	A background write thread created when the database starts. The thread is used to flush dirty pages out of the database to a permanent device (such as a disk).
bit	Short for binary digit. The smallest unit of information handled by a computer. One bit expresses a 1 or a 0 in a binary numeral, or a true or false logical condition, and is represented physically by an element such as a high or low voltage at one point in a circuit or a small spot on a disk magnetized one way or the other. A single bit conveys little information a human would consider meaningful. A group of 8 bits, however, makes up a byte, which can be used to represent many types of information, such as a letter of the alphabet, a decimal digit, or other character.
bloom filter	Bloom filter is a space-efficient binary vectorized data structure, conceived by Burton Howard Bloom in 1970, that is used to test whether an element is a member of a set. False positive matches are possible, but false negatives are not. In other words, a query returns either "possibly in set (possible error)" or "definitely not in set". In the cases, Bloom filter sacrificed the accuracy for time and space.
CEK	Acronym for column encryption key.

Terms	Description
CIDR	Acronym for classless inter-domain routing. Whereas classical network design for IPv4 sized the network prefix as one or more 8-bit groups, resulting in the blocks of Class A (8-bit), B (16-bit), or C (24-bit) addresses, CIDR allocates address space on any address bit boundary. A CIDR address is in the format of <i>IP address/Number of bits in a network ID</i> . For example, in <b>192.168.23.35/21</b> , <b>21</b> indicates that the first 21 bits are the network prefix and others are the host ID.
CLI	Acronym for command-line interface. A means of communication between a program and its user, based solely on textual input and output. Commands are input with the help of a keyboard or similar device and are interpreted and executed by the program. Results are output as text or graphics to the terminal.
CM	Acronym for cluster manager, a database management module. It manages and monitors the running status of functional units and physical resources in a system, ensuring stable running of the entire system.
CMK	Acronym for client master key in the full encryption scenario.
CU	Acronym for compression unit. It is the smallest storage unit in a column-store table.
core file	<p>A file that is created for further analysis when memory overwriting, assertion failures, or access to invalid memory occurs in a program, causing a process to fail.</p> <p>A core file stores memory dump data, and supports binary mode and specified ports. The name of a core file consists of the word "core" and the OS process ID.</p> <p>The core file is available regardless of the type of platform.</p>
core dump	A core dump, memory dump, or system dump consists of the recorded state of the working memory of a computer program at a specific time, generally when the program stops abnormally. In practice, other key pieces of program state are usually dumped at the same time, including the processor registers, which may include the program counter and stack pointer, memory management information, and other processor and OS flags and information. Core dumps are often used to assist in diagnosing and debugging errors in computer programs.
DBA	Acronym for database administrator who instructs or executes database maintenance operations.
DBLINK	An object of the path from one database to another database. Using DBLINK, you can query a remote database object.

Terms	Description
DBMS	Acronym for database management system. It is a piece of system management software that allows users to access information in a database. It is a collection of programs that allows users to access, manage, and query data in a database. A DBMS can be classified as memory DBMS or disk DBMS based on the location of the data.
DCF	Acronym for distributed consensus framework which implements data synchronization based on the Paxos algorithm.
DCL	Acronym for data control language.
DDL	Acronym for data definition language.
DEK	Acronym for data encryption key.
DML	Acronym for data manipulation language.
backup	A backup copy or an act of creating a backup refers to copying and archiving computer data for purposes of recovery in case the original copy of data is lost.
backup and restoration	A collection of concepts, procedures, and strategies to prevent data loss caused by invalid media or misoperations.
standby node	A node in the openGauss primary/standby solution. It functions as a backup of the primary node. If the primary node fails, the standby node is promoted to primary, ensuring uninterrupted data services.
crash	A crash (or system breakdown) is an event that a computer or program, such as a software application or an operating system, fails to function correctly. If a program experiences a crash, it often automatically exits. The program responsible may appear to freeze or hang until a crash reporting service reports the crash and any details relating to it. If the program is a critical part of the operating system, the entire system may crash or hang, often resulting in a kernel panic or fatal system error.
coding	Coding is representing data and information using code so that it can be processed and analyzed by a computer. Characters, digits, and other objects can be converted into digital code, or information and data can be converted into the required electrical pulse signals based on predefined rules.
encoding technology	A technology that presents data using a specific set of characters, which can be identified by computer hardware and software.

Terms	Description
table	A table consists of columns and rows. Each column is referred to as a field. The value in each field represents a data type. For example, if a table contains three fields <b>Name</b> , <b>City</b> , and <b>State</b> , it has three columns <b>Name</b> , <b>City</b> , and <b>State</b> . In every row of the table, the <b>Name</b> column contains a name, the <b>City</b> column contains a city, and the <b>State</b> column contains a state.
tablespace	A tablespace is a logical storage structure that contains tables, indexes, large objects, and long data. A tablespace provides an abstract layer between physical data and logical data, and provides storage space for all database objects. When you create an object, you can specify which tablespace it belongs to.
concurrency control	A DBMS service that ensures data integrity when multiple transactions are concurrently executed in a multi-user environment. In a multi-threaded openGauss environment, concurrency control ensures that database operations are safe and all database transactions remain consistent at any given time.
query	A request sent to the database with the purpose of updating, modifying, querying, or deleting information.
query operator	An iterator or a query tree node, which is a basic unit for the execution of a query. Execution of a query can be split into one or more query operators. Common query operators include scan, join, and aggregation.
durability	One of the ACID properties of a database transaction. After a transaction is complete, the changes made by the transaction to the database are permanently stored in the database and will not be rolled back.
stored procedure	A collection of SQL statements compiled and stored on a server in a large database system that can be executed using an interface (specifying the procedure name and parameters if any) to perform a specific operation.
operating system	An operating system (OS) is loaded by a boot program to a computer to manage other programs in the computer. Other programs are called applications or application programs.
Blob	Acronym for binary large object, a collection of binary data stored in a database. Blobs are typically videos, audio or other multimedia objects.
segment	A segment is a set of extents in a database. The smallest space scope of a database is an extent, which consists of data blocks. One or more segments comprise a tablespace.
<b>F – J</b>	

Terms	Description
failover	Automatic switchover from a faulty primary node to its standby node. Reversely, automatic switchback from the standby node to the primary node is called failback.
FDW	Acronym for foreign data wrapper. It is an SQL interface provided by openGauss. It is used to access big data objects stored in remote data so that DBAs can integrate data from unrelated data sources and store them in public schema in the database.
freeze	An operation automatically performed by the AutoVacuum Worker process when transaction IDs are exhausted. openGauss records transaction IDs in the row heading. When a transaction reads a row, the transaction ID in the row heading and the actual transaction ID are compared to determine whether this row is explicit. Transaction IDs are integers containing no symbols. If exhausted, transaction IDs are re-calculated outside of the integer range, causing the explicit rows to become implicit. To prevent such a problem, the freeze operation marks a transaction ID as a special ID. Rows marked with these special transaction IDs are explicit to all transactions.
GDB	Acronym for GNU debugger which is used to monitor the internal situation of a running program or rewind a crashed program to see what happened. GDB can perform the following operations (strengthening PDK functions) to detect bugs: <ul style="list-style-type: none"><li>• Starts a program and specifies anything that might affect its behavior.</li><li>• Stops a program under a specific condition.</li><li>• Checks what happens when a program stops.</li><li>• Modifies the program content to rectify the fault and proceeds with the next one.</li></ul>
GIN index	Acronym for generalized inverted index. It is used for handling cases where the items to be indexed are composite values, and the queries to be handled by the index need to search for element values that appear within the composite items.
GNU	The GNU Project was publicly announced on September 27, 1983 by Richard Stallman, aiming at building an OS composed wholly of free software. GNU is a recursive acronym for "GNU's Not Unix!". Stallman announced that GNU should be pronounced as Guh-NOO. Unix is a widely used commercial OS. GNU's design is Unix-like, but GNU differs from Unix by being free software.



Terms	Description
gsq	openGauss interaction terminal. It enables you to interactively enter and issue queries to openGauss, and view the query results. Queries can also be entered from files. In addition, <b>gsq</b> supports many meta commands and shell-like commands, allowing you to conveniently compile scripts and automate tasks.
GUC	Acronym for grand unified configuration which includes parameters for running databases, and the values of which affect database system behavior.
HA	Acronym for high availability which helps to minimize the duration of service interruptions caused by routine maintenance (planned) or sudden system breakdowns (unplanned), improving the system and application usability.
HBA	Acronym for host-based authentication which allows hosts to authenticate on behalf of all or some of that particular host's users. Those accounts can be all of the accounts on a system or a subset designated by the <b>Match</b> directive. This type of authentication can be useful for managing computing and other fairly homogenous pools of machines. In all, three files on the server and one file on the client must be modified to prepare for host-based authentication.
IV	Acronym for initialization vector. An IV is a block used for random encryption in many encryption modes. Therefore, different ciphertexts may be generated by using the same plaintext and key.
server	A combination of hardware and software designed for providing clients with services. This word alone refers to a computer running a server OS, or software or dedicated hardware providing services.
isolation	One of the ACID properties of a database transaction. Isolation means that operations and data used in a transaction are isolated from those in other concurrent transactions. Concurrent transactions are independent of each other.
relational database	A database that conforms to the relational model. It processes data using mathematical concepts and methods such as the set algebra.
archive thread	A thread started when the archive function is enabled on a database. The thread is used to archive database logs to a specified path.
failover	Automatic substitution of a functionally equivalent system component for a failed one. The system component can be a processor, server, network, or database.

Terms	Description
environment variable	Environment variables are used to define part of the environment in which a process runs. For example, it can be used to define a home directory, command search path, terminal being used, or the current time zone.
checkpoint	A mechanism that stores data at a certain time in the database memory to disks. openGauss periodically stores the data of committed transactions and data of uncommitted transactions to disks. The data and redo logs can be used for database restoration if a database restarts or breaks down.
encryption	A function hiding information content during data transmission to prevent the unauthorized use of the information.
node	Database nodes (or nodes) are physical and virtual servers that make up the openGauss database environment.
error correction	A technique that automatically detects and corrects errors in software and data flows to improve system stability and reliability.
process	An instance of a computer program that is being executed. A process consists of one or more threads. A process cannot use a thread occupied by another process.
PITR	Acronym for point-in-time recovery, a backup and restoration feature of openGauss. Data can be restored to a specified point in time if backup data and WAL logs are normal.
record	In a relational database, a record corresponds to data in each row of a table.
<b>K - O</b>	
KMC	Acronym for key management component.
KMS	Acronym for key management service.
KSF	Acronym for key store file.
logical replication	Data synchronization mode between primary and standby database nodes or between two databases. Different from physical replication which replays physical logs, logical replication transfers logical logs between two databases or synchronizes data through SQL statements in logical logs.
logical log	Logs recording database changes made through SQL statements. Generally, the changes are logged at the row level. Logical logs are different from physical logs that record changes of physical pages.
logical decoding	Logic decoding is a process of extracting all permanent changes in database tables into a clear and easy-to-understand format by decoding Xlogs.

Terms	Description
logical replication slot	In a logical replication process, logic replication slots are used to prevent Xlogs from being reclaimed by the system or <b>VACUUM</b> . A logical replication slot in openGauss is an object that records logical decoding positions. It can be created, deleted, read, and pushed by invoking SQL functions.
MVCC	Acronym for multi-version concurrency control. It is a protocol that allows a tuple to have multiple versions, on which different query operations can be performed. One advantage is that read and write operations do not conflict.
NameNode	NameNode is the centerpiece of an HDFS, managing the namespace of the file system and client access to files.
OM	Acronym for operations management. It provides management interfaces and tools for routine maintenance and configuration management of the database.
client	A computer or program that connects to or requests the services of another computer or program.
free space management	A mechanism for managing free space in a table. This mechanism enables the database system to record free space in each table and establish an easy-to-find data structure, accelerating operations (such as INSERT) performed on the free space.
junk tuple	A tuple that is deleted using the <b>DELETE</b> and <b>UPDATE</b> statements. When deleting a tuple, openGauss only marks the tuples that are to be cleared. The <b>VACUUM</b> thread will then periodically clear these junk tuples.
column	An equivalent concept of field. A database table consists of one or more columns.
logical node	Multiple logical nodes can be installed on the same physical node. A logical node is a database instance.
schema	Collection of database objects, including logical structures, such as tables, views, sequences, stored procedures, synonyms, indexes, and database links.
schema file	An SQL file that determines the database structure.
<b>P – T</b>	
page	Smallest memory unit for row storage in the relational object structure in openGauss. The default size of a page is 8 KB.
Paxos	Distributed consistency protocol.
PostgreSQL	An open-source relational DBMS developed by volunteers all over the world. PostgreSQL is not controlled by any companies or individuals. Its source code can be used for free.

Terms	Description
postmaster	<p>A thread started when the database service is started. It listens on connection requests from other nodes in the database or from clients.</p> <p>After receiving and accepting a connection request from the standby node, the primary node creates a WAL sender thread to interact with the standby node.</p>
publication	<p>A publication can be defined on any primary physical replication server. The node where a publication is defined is called the publisher. A publication is a set of changes generated from a table or a group of tables. It can also be described as a change set or replication set. Each publication exists in only one database.</p>
RHEL	<p>Acronym for Red Hat Enterprise Linux.</p>
redo log	<p>A log that records operations on the database. Redo logs contain the information required for performing these operations again. If a database is faulty, redo logs can be used to restore the database to its pre-fault state.</p>
RK	<p>Acronym for root key.</p>
SCTP	<p>Acronym for stream control transmission protocol. It is a transport-layer protocol defined by Internet Engineering Task Force (IETF) in 2000. The protocol ensures the reliability of datagram transport based on unreliable service transmission protocols by transferring SCN narrowband signaling over IP network.</p>
savepoint	<p>A savepoint marks the end of a sub-transaction (also known as a nested transaction) in a relational DBMS. The process of a long transaction can be divided into several parts. After a part is successfully executed, a savepoint will be created. If later execution fails, the transaction will be rolled back to the savepoint instead of being totally rolled back. This is helpful for recovering database applications from complicated errors. If an error occurs in a multi-statement transaction, the application can be restored by rolling back to the savepoint without terminating the entire transaction.</p>
session	<p>If a database receives a connection request from an application, a task is created for the connection. Sessions are managed by the session manager. They execute initial tasks and perform all user operations.</p>

Terms	Description
SMP	Acronym for symmetric multiprocessing. A group of processors (multiple CPUs) is integrated on a computer. These CPUs share the memory subsystem and bus structure. The OS must support multitasking and multithreading to ensure an SMP system achieves high performance. In databases, SMP means to concurrently execute queries using the multi-thread technology, efficiently using all CPU resources and improving query performance.
SQL	Acronym for structured query language. A standard database query language, consisting of data definition language (DDL), data manipulation language (DML), and data control language (DCL).
SSL	Acronym for secure sockets layer. It is a network security protocol introduced by Netscape. It is based on the TCP/IP and uses public key technology. SSL supports a wide range of networks and provides three basic security services, all of which use the public key technology. SSL ensures the security of service communication through the network by establishing a secure connection between a client and a server and then being able to securely send any data through this connection.
subscription	A subscription is the downstream side of logical replication. The node where a subscription is defined is called the subscriber. A subscription defines the connection to another database and the set of publications (one or more) that it wants to subscribe to.
oversubscription ratio	The ratio of downlink bandwidth to uplink bandwidth of a switch. A high oversubscription ratio indicates a highly oversubscribed traffic environment and severe packet loss.
Table Access Method	The table access method layer decouples the execution engine from the storage engine to implement the pluggable capability of the storage engine.
TCP	Acronym for transmission control protocol. It splits data into packets which are sent through the Internet protocol (IP), and checks and reassembles packets received through IP to obtain original information. TCP is a connection-oriented, reliable protocol that ensures information correctness in transmission.
trace	A specialized use of logging to record information about the way a program is executed. Programmers use the information for debugging. System administrators and technical support personnel can diagnose common problems by using software monitoring tools and based on this information.
strong consistency	A query cannot see any instantaneous intermediate state of a transaction.
full backup	Backup of the whole database.

Terms	Description
full synchronization	A data synchronization mechanism specified in the openGauss primary/standby solution, which is used to synchronize all data from the primary node to a standby node.
log file	A file containing a record of to activities made in a computer.
transaction	A logical unit of work performed within a database management system against a database. A transaction consists of a limited database operation sequence, and must have ACID features.
data	A representation of facts or directives for manual or automatic communication, explanation, or processing. Data includes constants, variables, arrays, and strings.
data partition	The action of dividing a table into parts (partitions) whose data does not overlap within a database instance. Tables can be partitioned by range, where the target storage location is mapped based on the range of the values in the column that is specified in the tuple.
database	A collection of data that is stored together and can be accessed, managed, and updated. Data in a view in the database can be classified into the following types: numerals, full text, digits, and images.
database instance	A database instance consists of a process in openGauss and files controlled by the process. openGauss allows multiple database instances to be installed on one physical node. A database instance is also called a logical node.
database primary/standby solution	openGauss provides a highly reliable HA solution. In this solution, each openGauss is called a primary or standby node. At the same time, only one openGauss is identified as the primary node. When the primary/standby system is deployed for the first time, the primary node performs full synchronization on the standby node, and then performs incremental synchronization on the standby node. When the primary/standby system is running, the primary node can receive data read and write operation requests and the standby node only synchronize logs.
database file	A binary file that stores user data and the data inside the database system.
data dictionary	A read-only collection of database tables and views containing reference information about the database. The information includes database design information, stored procedure information, user rights, user statistics, database process information, database increase statistics, and database performance statistics.
deadlock	A situation where different transactions are unable to proceed, because each holds a lock that the other needs.

Terms	Description
index	An ordered data structure in DBMS to help quickly query and update data in database tables.
statistics	Information that is automatically collected by databases, including table-level information (number of tuples and number of pages) and column-level information (column value range distribution histogram). Statistics in databases are used to estimate the cost of query plans to find a plan with the lowest cost.
stop word	In computing, stop words are words which are filtered out before or after processing of natural language data (text), saving storage space and improving search efficiency.
<b>U - Z</b>	
Ustore	Alias for the in-place update storage engine which solves the problems of space expansion and large tuples of the Append update storage engine. The design of efficient rollback segments is the basis of the in-place update storage engine.
Undo Record	An undo record can be inserted, queried, and organized. It connects to the ustore through the northbound interface and connects to the buffer pool through the southbound interface.
Undo Space	Manages physical resources of undo records, including adding and deleting undo files.
Undo Zone	Be bound to service threads and manages the undo logical resources of each service thread.
TransactionSlot	Records undo records by transaction granularity for transaction rollback and undo record recycling.
TIMECAPSULE	Flashback keyword. After the flashback technology is used, it takes only seconds to restore the submitted data before the database is modified. The restoration time is irrelevant to the database size.
RECYCLE BIN	After the recycle bin function is enabled, DROP TABLE can move a table and its sub-objects to the recycle bin.
PURGE	Clears objects in the recycle bin.
VACUUM	A thread that is periodically started by a database to remove junk tuples. Multiple <b>VACUUM</b> threads can be started concurrently by setting a parameter.
verbose	A verbose option specifies the information to be displayed.
WAL	Acronym for write-ahead logging, a standard method for logging a transaction. Corresponding logs must be written into a permanent device before a data file (carrier for a table and index) is modified.

Terms	Description
WAL receiver	A thread created by a standby node during database replication. The thread is used to receive data and commands from the primary node and to tell the primary node that the data and commands have been acknowledged. Only one WAL receiver thread can run on one standby node.
WAL sender	A thread created on the primary node when the primary node has received a connection request from a standby node during database replication. This thread is used to send data and commands to the standby node and to receive responses from the standby node. Multiple WAL sender threads may run on one primary node. Each WAL sender thread corresponds to a connection request initiated by a standby node.
WAL writer	A thread for writing redo logs that are created when a database is started. This thread is used to write logs in the memory to a permanent device, such as a disk.
Xlog	A transaction log. A logical node can have only one .xlog file.
xDR	The x detail record is a general term that refers to call detail records (CDRs), user flow data records (UFDRs), transaction detail records (TDRs), and statistics detail records (SDRs) on the user and signaling planes.
physical node	A physical machine.
system catalog	A system catalog stores meta information about a database, including user tables, indexes, columns, functions, and data types.
pushdown	In openGauss, a primary database node can send a query plan to multiple DNs for parallel execution. This behavior is called pushdown. It achieves better query performance than extracting data to a primary database node for query.
compression	Data compression, source coding, or bit-rate reduction involves encoding information that uses fewer bits than the original representation. Compression can be either lossy or lossless. Lossless compression reduces bits by identifying and eliminating statistical redundancy. No information is lost in lossless compression. Lossy compression reduces bits by identifying and removing less noticeable information. The process of reducing the size of a data file is commonly referred as data compression, although its formal name is source coding (coding done at the source of the data, before it is stored or transmitted).
consistency	One of the ACID properties of a database transaction. Data in the database must comply with integrity constraints.



Terms	Description
metadata	Data that provides information about other data. Metadata describes the source, size, format, or other characteristics of data. In the data field, metadata helps to explain the content of a data warehouse.
atomicity	One of the ACID properties of a database transaction. A transaction is composed of an indivisible unit of work. Operations performed in a transaction must be all finished or have not been performed. If an error occurs during transaction execution, the transaction is rolled back to the state when it was not performed.
dirty page	A page that has been modified, where the changes are not yet written to a permanent device.
incremental backup	Incremental backup only saves data changed since the last valid backup.
incremental synchronization	A data synchronization mechanism in the openGauss primary/standby solution, which is used to synchronize inconsistent data from the primary node to a standby node.
primary node	A node that allows read and write operations and works with all standby nodes in the openGauss primary/standby system. At the same time, only one node in the primary/standby system is identified as the primary node.
subject term	A standardized word or phrase that describes the subject of an article.
dump file	A specific type of trace file. A dump is typically a one-time output of diagnostic data in response to an event, whereas a trace tends to be continuous output of diagnostic data.
minimum restoration point	A method used by openGauss to ensure data consistency. During startup, openGauss checks consistency between the latest WAL logs and the minimum restoration point. If the record location of the minimum restoration point is greater than that of the latest WAL logs, the database fails to start.