

The table defines the meaning of basic terms which are widely used within PostgreSQL's documentation and related publications. The purpose is the specification of an unambiguous terminology. Therefore it focuses on those terms, which sound similar but have a distinct meaning, and on those terms, which are different but have the same meaning. Hence it serves as a guideline for beginners as well as for authors of our documentation.

Additionally it addresses a problem within the global database community. People use basic technical terms - mostly - in a consistent and database independent way. But there are exceptions. Some of the terms listed here have a divergent meaning in other database systems. We want to make clear, which understanding the PostgreSQL community has concerning those terms.

Table xyz: Definition and meaning of widely used terms

Part 1: Fundamental Terms		
Term or Accepted Alternative Term (Outdated Term)	Definition	Notes
Server or Node (Database Server)	A server consists of real or virtual hardware where the PostgreSQL software is installed and running.	Don't confuse the term with <u>instance</u> . It is possible to run multiple <u>instances</u> on a single <u>server</u> . The term <u>node</u> is also used in the context of data structures like trees. Considering the context the meaning should become clear.

<p>Database Cluster (Cluster)</p>	<p>A database cluster is a group of 3 or more <u>databases</u> (<i>template0, template1, postgres, ...</i>) plus their meta information.</p> <p>Its control- and runtime-information is stored in files which reside - in analogy to the situation of <u>databases</u> - in one directory and its subdirectories. This directory is often referred to as PGDATA.</p> <p>The database cluster is controlled and managed by one <u>instance</u>.</p>	<p>Depending on the OS the database cluster's default name is "data" or "main", which is identical to the name of the directory at OS level.</p> <p>It is possible to create more <u>databases</u> in the database cluster, e.g. "MyDb". One can create new objects, eg: "MyTable", within the default database "postgres" as well as in new databases.</p> <p>The utility initdb creates a new database cluster (with 3 <u>databases</u>), not a new database within an existing database cluster.</p> <p>Don't confuse the term with the term <u>Replication Cluster</u>.</p> <p>The SQL key word CLUSTER has the meaning "arrange the sequence of rows within a table".</p>
<p>Database</p>	<p>A database is a collection of files. With the possible exception of tablespaces (and targets of links) they reside in PGDATA/base/<oid>/. The files hold the databases data (<u>heap</u>, <u>indexes</u>), its meta data (definitions, views, ...), and system data (functions, operators, ...) of this database.</p>	<p><oid> is the object id of the database within the <u>database cluster's</u> meta data.</p>
<p>Segment</p>	<p>The term "segment" is a synonym for "file". In the case of <u>WAL files</u> its size is 16MB (fix) and in the case of <u>heap</u> or <u>index</u> the size varies and grows up to 1 GB - the <u>instance</u> generates more files if heap or index grows above this limit.</p>	
<p>Instance (Server)</p>	<p>An instance is a group of processes (on a UNIX server) respectively one service (on a Windows server) plus <u>shared buffers</u>, which controls and manages exactly one <u>database cluster</u>.</p>	<p>Please consider the difference between the terms instance (processes, RAM) and database (files).</p>
<p>Schema</p>	<p>A schema is a logical cramp for a collection of objects like tables or functions within one database. (An analogy is a directory within a file system - but "schema" has no recursive nature.)</p>	<p>Every <u>database</u> contains the special schema "public". It holds the names of all system tables, data types, functions, and operators of this database.</p> <p>It is possible to define more tables etc. within the schema "public".</p>

Catalog	In the case of the two <u>schemas</u> “information_schema” and “pg_catalog” the term “schema” is replaced by the term “catalog”.	The catalog (respectively schema) “information_schema” is required by the SQL standard and holds information about the database structure in a standardized form. The catalog (respectively schema) “pg_catalog” holds information about the database structure in a PostgreSQL specific form.
Summary: The hierarchy to any database object is: server.database_cluster.database.schema.object, eg: apollo.main.postgres.public.employee		
Heap	The heap contains the original data (values). He is realized within database files and mirrored in <u>shared buffers</u> .	
Index	The index contains data (as keys) and pointers to the <u>heap</u> . He is realized within database files and mirrored in <u>shared buffers</u> .	
<h2>Part 2: Fundamental Software Components</h2>		
Shared Buffers	Shared buffers are RAM where parts of the <u>database</u> (files) are mirrored by the <u>instance</u> . Its different processes as well as the <u>backend processes</u> have concurrent access to the shared buffers. Shared buffers are constructed within the operating system’s <u>shared memory</u> .	
Shared Memory	RAM which may be accessed by multiple processes. The operating system uses <u>semaphores</u> to grant access to shared memory.	“Shared Memory” is an operating system term, especially in the UNIX family.

Semaphore	Semaphores are provided by the operating system to offer applications undividable (atomic) operations. In conjunction with <u>shared memory</u> they are necessary to implement inter-process communication (IPC).	
Page	A page is a contiguous block of RAM (usually 8 kB) with a PostgreSQL specific layout. This layout is identical in RAM and in files. I/O activities of the <u>instance</u> transfers entire pages between RAM and files.	The operating system page size, the file system block size and PostgreSQL's page size may differ from each other.
Dirty Page	A dirty page is a <u>page</u> within the <u>shared buffers</u> , which is modified since last read from disk but not yet flushed back to disk.	
Query Optimizer or Query Planner (A suggestion to unite the two terms: optimizing query planner)	An algorithm which creates an optimal <u>execution plan</u> for a given SQL statement under consideration of the results of last ANALYZE.	
Execution Plan	An execution plan is the translation of an SQL statement into a tree of nodes. Each single node describes the low level access path to data and the estimated use of resources in terms of time, I/O and RAM.	
Postmaster	Within an <u>instance</u> there is exactly one postmaster process. Every new connection request from a <u>frontend process</u> to the instance reaches this process. As a reaction to such a request, the postmaster creates a new <u>backend process</u> and passes the connection and all further responsibilities to him.	
Backend Process (server process)	There is one backend process per connection. He receives requests from the assigned <u>frontend process</u> , executes them, and sends the result back to him.	Every backend process is assigned to exactly one <u>fronted process</u> .

Frontend Process or Client Process	Frontend processes initiate connections, send requests to the assigned <u>backend process</u> and receive the results.	It is possible that one frontend process communicates with multiple <u>backend processes</u> via different connections. Frontend and backend processes, which are assigned to each other, may run on different or the same <u>server</u> .

Part 3: Multiversion Concurrency Control (MVCC)

Version of a Row	The data values of a single entity are stored in a contiguous memory or disc area. A later data modification does not overwrite these values but copies the complete area to a different area - by taking the modifications into account. Each one of the two (or more) areas are called a “version of a row”.	
Row	A row is that single <u>version of the row</u> , which is - from the perspective of the transaction - the newest one. From the perspective of the complete database and some other transactions there may be newer versions of the same row.	

Part 4: Write Ahead Logging (WAL)

<p>WAL file or WAL (transaction logfile XLog file log segment file WAL segment file segment segment file)</p>	<p>A WAL file is a file in binary format where the <u>log records</u> of the complete <u>database cluster</u> are stored. Multiple WAL files as well as log records within a single WAL file have a certain order which represents the sequence of changes taken place in the database cluster.</p>	<p>The divergent term “logfile” denotes a text file where warn- and error-messages are reported.</p>
<p>Log record (Log entry)</p>	<p>A log record contains information about a single data change. He is stored in a <u>WAL file</u>. The log entry contains not only the data itself but also information about the physical storage location in a file, information about indexes, and VACUUM information.</p>	<p>Log records are the atoms of <u>WAL files</u> as well as of all kind of replication (directly or after a logical decoding). The term “log record” is derived from the verb “to log” - it’s not an abbreviation for “logical”.</p>
<p>LSN: Log Sequence Number</p>	<p>A LSN is a 64-bit integer representing a byte position within the sequence of <u>WAL files</u>.</p>	
<p>Logical Decoding</p>	<p>Logical decoding is a feature to offer consumers a stream of the ongoing data modifications. It is realized by decoding the continuously arising <u>log records</u>.</p>	
<p>Checkpoint</p>	<p>A checkpoint is a point in time when all <u>dirty pages</u> (heap plus index) of the <u>shared buffers</u> are flushed to disk and a special checkpoint record is written to the <u>WAL file</u>.</p>	
<p>Snapshot</p>	<p>A consistent (logical) view to the data in the <u>database cluster</u>, which is shown to a transaction as of its start-time, no matter what other transactions are doing while this transaction runs. So, every transaction has its own snapshot.</p>	
<p>Savepoint</p>	<p>A savepoint is a special mark (a timestamp) inside a transaction. Data modification actions arising after this point in time may be rolled back to the time of the savepoint.</p>	

Part 5: Backup and Recovery

<h2>Part 5: Backup and Recovery</h2>		
Cold Backup or File System Level Backup	A cold backup takes all files of the <u>database cluster</u> with pure operating system commands and without any support of PostgreSQL's tools.	To get a useable backup with this method the instance MUST be shut down!
Hot Backup	A hot backup is a backup which is taken with PostgreSQL tools during the <u>instance</u> is running.	
Logical Backup	A logical backup is a <u>hot backup</u> which creates new files in SQL syntax or in a binary format. Those files are independent from the database clusters directory. A logical backup takes a complete <u>database</u> or <u>database cluster</u> or parts of it.	A logical backup is taken with one of the tools <code>pg_dump</code> or <code>pg_dumpall</code> .
Base Backup	A base backup is a <u>hot backup</u> which copies all files of a <u>database cluster</u> (including <u>WAL files</u>) and creates an additional <u>backup history file</u> . It is not possible to take parts of the database cluster - it must be processed as a whole.	You need the <code>pg_start/stop_backup</code> command or the <code>pg_basebackup</code> command to create a base backup.
Continuous Archiving or Log Shipping	The term denotes the process of copying all <u>WAL files</u> , which are continuously created by the <u>instance</u> , to an additional and reliable storage medium - eg: for archiving, replication, or backup purposes.	You can use archived <u>WAL files</u> during a recovery (but you need more than just this files - namely the files of a <u>base backup</u>). The term "log shipping" is used in replication scenarios.
PITR: Point in Time Recovery	A recovery process can stop at a choosable timestamp which lays between the time of taking the <u>base backup</u> and the time of the last COMMIT before a crash.	To use this feature it is necessary to set certain parameters for the generation of the <u>WAL files</u> .

Backup History File	A backup history file contains meta information and timestamps about the <u>base backup</u> .	A backup history file is created automatically during the generation of a <u>base backup</u> .
<h2>Part 6: Replication</h2>		
Replicating Cluster (Cluster)	A replicating cluster is a set (transitive closure) of <u>instances</u> which interact in any kind of replication.	Don't confuse the term with <u>Database Cluster</u> .
Master Instance (Master server Primary server)	A master instance is an <u>instance</u> which pushes <u>log records</u> to one or more <u>standby instances</u> (in the case of <u>streaming replication</u>) or sends <u>WAL files</u> to them (in the case of <u>log shipping replication</u>).	
Standby Instance (Standby server Slave server Secondary server)	A standby instance is an <u>instance</u> which receives and processes <u>log records</u> or <u>WAL files</u> with the aim that the standby instance contains the same data as their <u>master instance</u> - after a more or less small time gap.	
Warm Standby Instance (Warm Standby Server)	A <u>standby instance</u> is called a "warm standby instance" if it does not accept client connections.	
Hot Standby Instance (Hot Standby Server)	A <u>standby instance</u> is called "hot standby instance" if it accepts client connections in read-only mode.	
Log Shipping Replication (26.2.)	The data transfer from the <u>master instance</u> to the <u>standby instance</u> is realized by the transfer of <u>WAL files</u> .	In a non-replication context transfers like this are called <u>continuous archiving</u> or simply <u>log shipping</u> .

Streaming Replication (26.2.5.)	The data transfer from the <u>master instance</u> to the <u>standby instance</u> is realized by <u>log records</u> which are streamed over an IP connection. ??? relation to logical decoding ???	
Logical Decoding	See above.	
Replication Slot (47.2.2.)	?	
Physical replication Slot	?	
Streaming Replication Slot (26.2.5./26.2.6.)	?	
Logical Replication Slot or Logical Slot (47.2.2.)	?	