

# FreeBSD

## PostgreSQL



### What is PostgreSQL?

PostgreSQL is an object-relational database management system (ORDBMS). It is freely available and usable without licensing fee. PostgreSQL was originally developed at the University of California. FreeBSD was also developed there and both projects use the same simple BSD-license.


It is a very advanced database system in the open source area. PostgreSQL supports the SQL92, SQL99 and SQL2003 standards and has an array of its own extensions.

In PostgreSQL these are storage of non-atomic data, inheritance and object identities as implemented in object-relational databases. Users can extend the system with self-defined data types, operators and functions. Apart from support for referential integrity and an advanced transaction management, PostgreSQL also offers definition of triggers and rules to control access to database objects.



### Large database server

- ▶ A small to mid-level database server has just small hardware requirements. It is important to put parts of the installation into a jail to protect the PostgreSQL process.



### Further Information

*PostgreSQL (2nd Edition)*, Korry Douglas, Sams Publishing, ISBN: 0672327562

*Beginning Databases with PostgreSQL: From Novice to Professional, Second Edition*, Neil Matthew, Apress, ISBN: 1590594789

*PostgreSQL Developer's Handbook*, Ewald Geschwinde, Sams Publishing, ISBN 0672322609

*Beginning PHP and PostgreSQL 8*, W. Jason Gilmore, Apress, ISBN 1590595475

*PHP and PostgreSQL Advanced Web Programming*, Ewald Geschwinde, Sams Publishing, ISBN 0672323826

PostgreSQL-homepage: [www.postgresql.org](http://www.postgresql.org)

PGCluster-homepage:  
<http://pgcluster.projects.postgresql.org/index.html>

FreeBSD: [www.freebsd.org](http://www.freebsd.org)

- ▶ *Slony (Replication)* – Features: asynchronous mode, single master to multiple slaves, slaves can be cascaded, uses triggers, replications on level of tables, no exclusive locks necessary, can replicate schema changes, all tables require a primary keys, all knots in the network need to be available at all times, cannot replicate large objects, no automatic failover.

### Administration

- ▶ *pgAdmin*: Extensive graphical tool for administration that should fulfill nearly all wishes. It stands out with its ease-of-use and clearly arranged display of the database structure. An installed GTK2 environment is a prerequisite.
- ▶ *phpPgAdmin*: A PHP-based web interface that can be used with the browser of any operating system.


### Client software for PostgreSQL

- ▶ *ERP/CRM*: Compiere, SQL-Ledger
- ▶ *Connectivity*: ODBC/JDBC-driver, Ruby-, Python and Perl-Bindings, C/C++-Libraries, Qt-Plugin
- ▶ *Stored-Languages*: pl/Ruby, pl/Python, pl/Pearl, pl/Java
- ▶ *Groupware*: OpenGroupware, various POP3/IMAP-server, anti-spam administration
- CMS/Shop-Systeme*: LivingShop, XIST4C
- Image Dazabase*: gPictureBase

### Usage Examples

#### Small to mid-level database server

- ▶ A small to mid-level database server has just small hardware requirements. It is important to put parts of the installation into a jail to protect the PostgreSQL process.



## What does PostgreSQL offer?

- ▶ PostgreSQL supports the ANSI-SQL-Standard SQL92. Of the 5 categories of the SQL99 standard, Framework, Foundation, Call Level Interface, Persistent Stored Modules and Host Language Bindings are implemented in the categories 1, 2 and 5. Moreover PostgreSQL offers among many of its own extensions geometric data types.
- ▶ *Language Scope*: Mass operations on one or multiple tables are possible with PostgreSQL just as operations with Subselects, Outer Joins (the Full Outer Join too) or Views.
- ▶ *Referential Integrity*: Tables can be referenced with the definition of primary and foreign keys. That way the relations between tables is represented and the system checks the correctness of these relations by itself.
- ▶ *Transactions*: PostgreSQL supports ACID and Multi Version Concurrency Control (MVCC). In PostgreSQL read access never has to wait for write access and vice versa. And there is Two-phase-commitment and Savepoints.
- ▶ *Concurrently Online Index Builds*: Indexes could be created, without locking an index table. During index creation competing INSERTs, UPDATEs und DELETEs are possible.
- ▶ *Triggers*: In PostgreSQL ‘Before’ or ‘After’ triggers are implemented.
- ▶ *Programming*: offers interfaces for ODBC and JDBC as well as interfaces to multiple other programming languages.
- ▶ *Stored Procedures*: These are functions written in different programming languages as well as SQL or PostgreSQL procedural languages like PL/pgSQL. They are compiled and directly stored in the database server, making them very fast.
- ▶ *Object-relational Characteristics*: PostgreSQL allows the storage of non-atomic data types in arrays or geometric data types. Both are extensions of PostgreSQL not described in the SQL standard. All database objects, whether they are tables, Views, Sequences or single data sets, have their own object identity in PostgreSQL. Apart from that inheritance from table characteristics to derived child tables is supported.
- ▶ *Extendability*: The type system is extensible by custom data types, operators and functions, making the system individually customizable. Functions can be written in SQL or other procedural languages. With PL/PgSQL you have your own procedural language at your disposal.

- ▶ *Large Objects*: Binary data, like sound data or images, can be stored in tables and read in part or whole. To store Large Objects, two data types may be chosen.
- ▶ *GIN – Generalized Inverted Index*: Index structure storing a set of (*key, documentlist*) tuples whereas *documentlist* means a set of data in which a certain field contains *key*. Invers means not the primary key for referencing a record is stored, but expressions that are contain in the field itself.
- ▶ *Character Set*: PostgreSQL understands many international character sets, Multibyte-Coding is standard and allows EUC or Unicode.
- ▶ *SSL*: communication between client and server may be encrypted with SSL.
- ▶ *Kerberos Authentication*: Kerberos is a protocol for authentication in networks. It was developed specifically for secure Client-Server-Communication and can be compiled directly into PostgreSQL.
- ▶ *Operation*: PostgreSQL makes use of multi-processor architectures, supports “Slony”-replication and offers the user the choice of table spaces and Point-In-Time recovery.
- ▶ *Clustering und Loadbalancing*: The software ‘PG-Cluster’ offers all the clustering load balancing features for PostgreSQL.
- ▶ *Warm Standby Log Shipping*: The continuous archiving of transaction logs enables installation and configuration of high availability solutions with one or more standby servers. In case of a breakdown of the main server the standby servers can be used.

## Technical Data

Parameter	Value
Maximal size of database	Unlimited, actual use is currently at 1 TeraByte
Maximal size of table	16 TeraByte, 32 TeraByte with Version 8
Maximal size of row	1,6 TeraByte
Maximal size of field	1 GigaByte
Maximal number of rows in DB	Unlimited, the file system is the limiter
Maximal number of columns	250 to 1600, depending on type
Maximal size of indices in table	Unlimited

## FreeBSD and PostgreSQL

FreeBSD, due to its stability and outstanding integration of PostgreSQL in Ports, is an ideal base for a database server.

- ▶ *Superb operating system stability*: It is very important for a database server to have a reliable and stable operating system. FreeBSD has a well-known history for its stability.
- ▶ *Fast filesystem*: UFS2 with softupdates is a excellent filesystem. FreeBSD offers "softupdates", which is an extension to the internal filesystem code that keeps track of metadata operations, sorts them and always writes metadata out in an order that guarantees that the filesystem may not be in a damageable state ever and that you will lose only data from files you actually modified during the interrupted session. And it is about as fast as asynchronous metadata operation. Tables up to 32 PiB (PebiByte =  $2^{50}$  Bytes) or volumes up to 1YiB (YobiByte =  $2^{80}$  Byte) are possible.
- ▶ *Fast TCP/IP stack*: FreeBSD has a fast and stable TCP/IP stack. This is a precondition to serving clients quickly and efficiently.
- ▶ *Security*: A database server needs to be protected from its environment. FreeBSD offers important features for that:
  1. Packet Filter: With IPFW and PF the administrator has two very powerful network filters at his fingertips to suppress attacks on the server.
  2. Jails: A kind of ‘virtual FreeBSD’. Processes in a Jail have no way of interacting with processes on the host system or another jail.
- ▶ *All in all*: FreeBSD is the ideal operating system base to run a PostgreSQL database. They complement each other nearly perfectly.

## Software for PostgreSQL and FreeBSD

### System software

- ▶ *PGCluster (Loadbalancing and Clustering)* – Features: synchronous, replication on table level, multi-master, two or more database servers can simultaneously serve client requests, replication on demand, sent within a group of server, replication of sequences and large objects, serial data types are synchronised.