



Title	WHAT IS POSTGRESQL AND WHY IS IT SO USEFUL?
Description	Intro
Date	2022 27 Augus
Author	Arashtad
Author URI	https://Arashtad.com



This article focuses on one of the most useful and important databases in the IT industry and that is PostgreSQL. This open-source database is an advanced, enterprise-class relational database that supports SQL (relational) and JSON (non-relational) queries. This database management system has been developed by its community for over 20 years, which has contributed to its high levels of accuracy, integrity, and resilience. Many web, mobile, geospatial, and analytics applications utilize PostgreSQL as their primary data store or warehouse. A key feature of PostgreSQL is that it supports advanced data types, and it offers performance optimizations that are comparable to those found in commercial database systems like Oracle and SQL Server.

WHAT IS POSTGRESQL?

PostgreSQL (a.k.a “Postgres”) is a relational database management system (ORDBMS) that emphasizes extensibility and compliance with standards. The main purpose of a database server is to store data securely and support best practices. And retrieve it later, as requested by other applications, whether they are running on the same computer or on another computer across a network. A recent version also offers replication of the database itself for security and scalability. This makes it suitable for applications ranging from single-machine applications with few concurrent users to large internet-facing applications with many concurrent users.

Using multi-version concurrency control (MVCC), PostgreSQL avoids locking issues, is ACID-compliant, transactional (including most DDL statements), provides immunity to dirty reads, and is fully serializable. SQL:2011 is the most common standard. It is ACID-compliant and transactional (including most DDL statements). It handles complex SQL queries with many indexing techniques not available in other databases. Has triggers, views, and foreign keys.

Postgres has the ability to work with all major proprietary and open source databases. Additionally, it supports migration from them with its extensive standard SQL support and migration tools. Furthermore, if proprietary extensions had been used, the extensibility of this software can emulate many through built-in and third-party open source compatibility extensions, such as Oracle’s.

POSTGRES HISTORY

As a result of Professor Michael Stonebreaker’s leadership at Berkeley, the Postgres project started in 1986. They entitled it POSTGRES in Berkeley, which was the successor of its older version “Ingres” which was another product of the same place. In order to fully support multiple data types, POSTGRES developers designed it with the fewest features possible. In 1996, they renamed it PostgreSQL to illustrate its support for the SQL querying language (although PostgreSQL is still commonly abbreviated as Postgres). PostgreSQL is an open-source database project whose contributors – the PostgreSQL Global Development Group – continue to make regular major and minor releases.

POSTGRESQL USE CASES

Geospatial database

As a geospatial data store, PostgreSQL can be used with PostGIS to support location-based services and geographic information systems (GIS).

LAPP open source stack

The PostgreSQL database can run dynamic websites and apps as an alternative to the LAPP stack (LAPP stands for Linux, Apache, PostgreSQL, PHP, Python, and Perl);

General purpose OLTP database

For internet-scale applications, solutions, and products, startups and large enterprises alike use PostgreSQL as their primary data store.

Federated hub database

The Foreign Data Wrappers and JSON support make Postgres a federated hub for polyglot databases, including NoSQL.

THE FEATURES OF POSTGRESQL

There are many advanced features in PostgreSQL that are not available in other enterprise-class database management systems, including:

1. Sophisticated locking mechanism
2. Nested transactions (savepoints)
3. Multi-version concurrency control (MVCC)
4. Asynchronous replication
5. User-defined types
6. Foreign key referential integrity
7. Table inheritance
8. Native Microsoft Windows Server version
9. Tablespaces
10. Point-in-time recovery
11. Views, rules, subquery

POSTGRESQL VS. MYSQL

MySQL and Postgres both have been around for decades. Despite all their similarities, PostgreSQL and MySQL have a number of characteristics that set them apart from each other, as we will see. They're both secure RDBMSes with clustering capabilities and network fault tolerance.

ARCHITECTURE

The difference between Postgres and MySQL is that Postgres has object-relational capabilities, while MySQL has purely relational capabilities. This makes using PostgreSQL more difficult, but it also makes using PostgreSQL less complex. MySQL provides 16 different storage engines suitable for different use cases. InnoDB is the default storage engine, providing index-organized tables.

PostgreSQL has a single, ACID-compliant storage engine. Each client connection PostgreSQL establishes generates its own memory allocation, so it needs a lot of memory on systems with many connections. MySQL uses a single process and maintains one thread (or execution path) per connection, which is sufficient for most applications that do not exceed enterprise scope. PostgreSQL supports materialized views, which can improve performance for complex queries, and is more robust when it comes to views, triggers, and stored procedures.

SQL INSERT, UPDATE, and DELETE statements are supported by both databases with AFTER and BEFORE triggers. Postgres also provides an INSTEAD OF trigger, which allows complex SQL statements to be executed within a trigger using functions. Postgres offers the ability to call procedures written in languages other than SQL as well as standard SQL stored procedures.

DATA TYPES

A PostgreSQL database offers a wider variety of data types than MySQL. If your application uses any of its unique data types or unstructured data, it may be a better choice. Both databases will work if you only need basic character and numeric data types.



INDEXES

As your data grows, you can fine-tune your database performance by picking from multiple indexing options to get faster query responses from your database and a better user experience for your application users.

SECURITY

Both databases support user and group management and grant SQL privileges to roles. PostgreSQL supports IP-based client filtering and authentication using PAM and Kerberos, while MySQL supports PAM, native windows services, and LDAP for user authentication. In terms of security, the two databases have comparable options.

SUPPORT AND DOCUMENTATION

Due to the fact that Postgres is not developed by a company, it lacks support engineers, but its community support forums are exceptionally good. you can judge each vendor's availability of documentation yourself since it can be found online. MySQL offers similar forums, as well as Oracle's paid support plans.



WRAPPING UP

In this article, you learned about the PostgreSQL database which is also called Postgres. Also, you got familiar with its use cases and features and its comparison to the MySQL database. Generally, choosing a database management system when starting a new project is an afterthought. Frameworks usually come with some kind of object-relational mapping tool (ORM), which makes them all equally slow and hides the differences between the platforms. A good developer must always make informed decisions among the different options, weighing the benefits and disadvantages of each. Don't fall into the trap of familiarity and comfort – don't fall into the trap of comfort and familiarity.

