

[By OnlineInterviewQuestions.com](http://OnlineInterviewQuestions.com)

PostgreSQL Interview Questions

Postgresql is a widely used system for Object-Relational Database Management that is mainly used for large web applications. Every year there are some vacancies, in big companies like Oracle, SQLite, Amazon web services and such others, to fill up in this field. People with technical expertise in the computer can easily crack the interview.

Here, we have listed few important and expected questions that may come up in your interview. With a little bit of practice and hard work and thoroughness, the questions can be understood. Going through the questions, you may find some questions that come up almost in every interview.

These predictable questions and answers will help you get through the interview without breaking a sweat.

Here are few Postgresql Interview questions with their answers:

Q1. List some of the features of Postgresql ?

Following are some of the major features of Postgresql :

1. Object-relational database
2. Supports major Operating systems
3. Support Extensibility for SQL and Complex SQL queries
4. Nested transactions
5. Flexible API and Database validation
6. Multi-version concurrency control (MVCC) and Procedural languages
7. WAL and Client server
8. Table inheritance & Asynchronous replication

Q2. What is the option that can be used in PostgreSQL to make transactions see rows affected in previous parts of the transaction?

The SQL standard is defined by four levels of transaction isolation basically regarding three phenomena. The three phenomenon must be prevented between concurrent transactions. The unwanted phenomena are:

- **Phantom read:** A transaction that re-executes a query, returning a set of rows that satisfy a search condition and then finds that the set of rows that have been satisfying the condition has changed due to another recently-committed transaction.
- **Non-repeatable read:** A transaction that re-reads the data that it has previously read and then finds that data has already been modified by another transaction (that committed since the initial read).

- **Dirty read** : A transaction when reads data that is written by a concurrent uncommitted transaction is the dirty read.

Q3. Put some light on Multi-Version concurrency control?

MVCC or better known as Multi-version concurrency control is used to avoid unwanted locking of the database. The time lag for the user is removed so that one can easily log into his database. All the transactions are well- kept as a record. The time lag occurs when someone else is on the content.

Q4. How are the stats updated in Postgresql?

It is not that hard as it seems. To get your statistics updated in PostgreSQL a special function called explicit 'vacuum' call is made. The method to do is to create a Vacuum where the option of Analyze is used to update statistics in Postgresql

```
VACUUM ANALYZE ;
```

is the syntax.

Q5. Provide a brief explanation of the functions in Postgresql.

Anywhere, functions are an important part because they help in executing the code on the server. Some of the languages to program functions are PL/pgSQL, a native language of PostgreSQL, and other scripting languages like Perl, Python, PHP, etc. statistical language named PL/R can also be used to increase the efficiency of the functions.

Q6. List different datatypes of Postgresql?

There are new, different data-types supported by Postgresql. Following are those data-types:

- UUID
- Numeric types
- Boolean
- Character types
- Temporal types
- Geometric primitives
- Arbitrary precision numeric
- XML
- Arrays etc.

Users can also create their indexes and get them indexed.

Q7. Which are different types of database administration tools used in Postgresql?

There are the number of data administration tools, and they are

- Phppgadmin
- Psql
- Pgadmin

Out of these, phppgadmin is the most popular one. Most of these tools are front-end administration tools and web-based interfaces.

Q8. Do provide an explanation for pgadmin? (100% asked Postgresql Interview Questions)

Pgadmin is a feature that is known to form a graphical front-end administration tool. This feature is available under free software released under Artistic License. Pgadmin iii is the new database administration tool released under artistic license.

Q9. What is PostgreSQL?

Postgres or simply known as **Postgresql** in the SQL world is one of the widely and popularly used for Object-Relational Database Management System that is used mainly in large web applications. It is one of the open-source object-relational database systems which also powerful. It provides additional and substantial power by incorporating four basic concepts in such a way that the user can extend the system without any problem. It extends and uses the SQL language that is combined with various features to safely scale and store the intricate data workloads.

Q10. List different advantages of Postgresql?

Following are some of the advantages of PostgreSQL :

- Stable
- Reliable
- Extensible
- Easy to learn
- Open source
- Designed for High Volume Environments
- Cross Platform
- Better Support
- Flexible

Q11. What is multi-version control in PostgreSQL?

Multi-version concurrency control or MVCC in PostgreSQL is used to avoid unnecessary locking of the database. This removes the time lag for the user to log into his database. This feature or time lag occurs when someone else is accessing the content. All the transactions are kept as a record.

Q12. What are string constants in PostgreSQL?

A **string constant** in PostgreSQL is a sequence of some character that is bounded by single quotes (').

Example

'This is a string Constant'

Q13. What Is A CtId?

CTIDs is a field, which exists in every PostgreSQL table and is known to identify specific physical rows according to their block and offset positions within a particular table. They are used by index entries to point to physical rows. It is unique for each record in the table and easily denotes the location of a tuple. A logical row's CTID changes when it is updated, so the CTID cannot be used as a long-term row identifier. However, it is sometimes useful to identify a row within a transaction when no competing update is expected.

Q14. Provide an explanation About Write Ahead Logging?

WAL or write-ahead logging is a standard method to ensure data integrity. It is a protocol or the correct rule to write both actions and changes into a transaction log. This feature is known to increase the reliability of the database by logging changes before any changes or updating to the database. This provides the log of the database in case of a database crash. This helps to start the work from the point it was discontinued.

Q15. What is the command enable-debug in PostgreSQL?

The command enable-debug is used to enable the compilation of all the applications and libraries. The execution of this procedure usually impedes the system, but it also amplifies the binary file size. Debugging symbols which are present generally assist the developers for spotting the bugs and other problems which may arise associated with their script.

Q16. What are the Indices of PostgreSQL?

Indices of PostgreSQL are inbuilt functions or methods like **GIST Indices**, **hash table** and **B-tree (Binary tree)** which can be used by the user to scan the index in a backward manner. Users can also define their indices of PostgreSQL.

Q17. What are tokens in PostgreSQL?

Tokens in PostgreSQL are the building blocks of any source code. They are known to comprise many of the special character symbols. These can be regarded as constants, quoted identifiers, other identifiers, and keywords. Tokens which are keywords consist of pre-defined SQL commands and meanings. Identifiers are used to represent variable names like columns, tables, etc.

Q18. How to start database server in PostgreSQL?

Before you can have access to the database, you must be able to start the database server. The server program of the database is called Postgres. The Postgres program must know where to find the data it is supposed to use. This is done with the -D option. Thus, the simplest way to start the server is:

1. /usr/local/etc/rc.d/010.pgsql.sh start
2. /usr/local/etc/rc.d/PostgreSQL start

Q19. What will be the new characteristics of Postgre 9.1?

During the process of updating the project, one can never be certain what features will go in and which ones won't make the cut. The project has precise and stringent standards for quality, and some patches may or may not match them before the set deadline. Currently, the 9.1 version is working on some important features which include JSON support, synchronous replication, nearest-neighbor geographic searches, collations at the column level, SQL/MED external data connections, security labels as well as index-only access. However, this list has a high chance of changing completely by the time Postgre 9.1 is released.

Q20. Compare 'PostgreSQL' with 'NoSQL'

The expression 'NoSQL' encompasses a wide collection of implementations which are part of the non-relational database. This includes tiny embedded databases such as TokyoCabinet, massive bunched data processing platforms such as Hadoop and everything in between. In short, it's practically impossible to comment on the range comprised by NoSQL as a typical class.

Choosing between the non-relational and relational databases is also quite commonly debated as both have existed alongside each other for over forty years. In fact, users should opt for the features, community support and implementation of the database according to their current application needs. Additionally, use of multiple various databases for sizeable projects is becoming more of a norm than a trend. Moreover, the users of PostgreSQL are no exception.

Q21. Explain the history of PostgreSQL.

The **origin of PostgreSQL** dates back to 1986 as part of the POSTGRES project at the University of California at Berkeley and has more than 30 years of active development on the core platform. It runs on all the major operating systems and has been ACID-compliant since 2001. It also has add-on like PostGIS database extender. In MAC OS Postgresql is the default database. Michel Stonebraker is Father of Postgresql who has started the Post Ingres project for supporting Contemporary Database systems. PostgreSQL's developers pronounce PostgreSQL as It is abbreviated as *Postgres* because of ubiquitous support for the SQL Standard among most relational databases. PostgreSQL, originally called Postgres, was created at UCB by a computer science professor named Michael Stonebraker, who went on to become the CTO of Informix Corporation.

Stonebraker started Postgres in 1986 as a followup project to its predecessor, Ingres, now owned by Computer Associates. The name Postgres thus plays off of its predecessor (as in "after Ingres"). Ingres, developed from 1977 to 1985, had been an exercise in creating a database system according to classic RDBMS theory. Postgres, developed in 1986-1994, was a project meant to break new ground in database concepts such as exploration of "object-relational" technologies. An enterprise-class database, PostgreSQL boasts sophisticated features such as Multi-Version Concurrency Control (MVCC), point in time recovery, tablespaces, asynchronous replication, nested transactions (savepoints), online/hot backups, a sophisticated query planner/optimizer, and write-ahead logging for fault tolerance.

Q22. How can we change the columns datatype in PostgreSQL?

Use **change column type statement** with ALTER TABLE command to change a column type in PostgreSQL.

Example

```
ALTER TABLE table_name  
ALTER COLUMN column_name [SET DATA] TYPE new_data_type;
```

Q23. What is Cube Root Operator (||/) in PostgreSQL?

PostgreSQL Cube Root Operator (||/) is used to get the cube root of a number.

Example

```
SELECT ||/40 AS "Cube Root of 40";
```

Q24. What is use of pgadmin in PostgreSQL?

It is a free open source GUI tool PostgreSQL database administration tool for Windows, Mac OS X, and Linux system. It is used for information retrieval, development, testing, and ongoing maintenance of Databases.

Q25. What is table partitioning in PostgreSQL?

Table partitioning in PostgreSQL is the process of splitting a large table into smaller pieces. A partitioned table is a logical structure used to divide a large table into smaller pieces called partitions.

Please Visit [OnlineInterviewquestions.com](https://www.onlineinterviewquestions.com) to download more pdfs