



# PostgreSQL Interview Questions

[Click here](#) to view the live version of the page

# Top PostgreSQL Interview Questions and Answers

PostgreSQL is the most advanced, open-source relational database management system that uses SQL (Structured Query Language) as its main query language. Several big tech companies such as Apple and Cisco have embraced the technology in their back-end applications. PostgreSQL supports multiple operating systems like Windows, Linux, and macOS.

Read along as we cover all the top industry questions related to PostgreSQL and help you ace your next interview.

These PostgreSQL Interview questions can be classified into these different categories, namely:

[PostgreSQL Basic Interview Questions for Freshers](#)

[PostgreSQL Developer Interview Questions for Experienced](#)

[PostgreSQL Interview Questions for 3 Years Experience](#)

[PostgreSQL Interview Questions for 5 Years Experience](#)

[PostgreSQL Advanced Interview Questions](#)

[PostgreSQL Interview Questions for 7 Years of Experience](#)

[PostgreSQL Interview Questions for 10 Years of Experience](#)

[PostgreSQL DBA Interview Questions For Experienced Professionals](#)

[PostgreSQL Replication Interview Questions](#)

[PostgreSQL Backup Interview Questions](#)

[PostgreSQL Developer Salary Trends](#)

[PostgreSQL Developer Job Trends](#)

## [PostgreSQL Developer Roles and Responsibilities](#)

### [Conclusion](#)

#### Did You Know?

- Postgres is the undisputed most admired and desired database in Stack Overflow's 2023 developer survey, according to the article published on [EDB](#).
- Over 60% of Fortune 500 companies and [Fortune 50 banks](#) trust PostgreSQL for their data management needs.
- PostgreSQL also has a feature that includes being relational and object-oriented in design and is also called an [ORDBMS](#).

## PostgreSQL Basic Interview Questions for Freshers

### 1. What is PostgreSQL?

PostgreSQL, commonly referred to as Postgres, powers applications with its robust and open-source relational database management system. Developers leverage its extensive features, such as Function Overloading and Table Inheritance, to create advanced applications. PostgreSQL seamlessly operates across major operating systems, including Windows, UNIX, macOS, and Linux.

### 2. What are the advantages of PostgreSQL?

The advantages of PostgreSQL include:

- PostgreSQL is highly fault-tolerant, owing to its feature of write-ahead logging.

- It is flexible and easy to learn.
- It supports a variety of replication methods.
- It can be used for large-scale web applications because of its powerful and robust nature.
- As the source code of PostgreSQL is available for free due to its open-source license, users can edit and modify it easily according to their business requirements.

### 3. Define a non-clustered index.

In a non-clustered index, the order of the index rows differs from the physical order of the real data. The leaf pages of a non-clustered index instead contain pointers to the real data rather than the actual data itself. Its main advantage is that it provides faster access to data.

### 4. Which data types are used in PostgreSQL?

The following data types are used in PostgreSQL:-

- Numeric data type (Integer, Float)
- Geometric primitives
- Boolean data type
- Character data type (varchar, char, text)
- Monetary data type
- Array
- Document data type (JSON, XML, Key-value, etc.)
- Date/Time data type
- Customization data type (Composite, custom types, etc.)

### 5. What do you mean by a parallel query?

Parallel query in PostgreSQL is an advanced feature. It allows the arrangement of query plans in such a way that they can exploit multiple CPUs. This helps in answering user queries in a much faster and quicker manner.

## 6. What is the meaning of PgAdmin?

PgAdmin is a free open-source graphical front-end PostgreSQL database administration tool. This web-based GUI tool is prominently used to manage PostgreSQL databases. It assists in monitoring and managing numerous complex PostgreSQL and EDB database systems. PgAdmin is used to accomplish tasks like accessing, developing, and carrying out quality testing procedures.

## 7. Define Write-Ahead logging.

Write-Ahead Logging is a technique used to ensure the data integrity of PostgreSQL databases. It helps in maintaining the resilience or the reliability of the database. Write-ahead logging is a method wherein any changes and actions in the database are logged in a transaction log prior to the updating or modification of the database. In case there is a database crash, this feature helps in providing the log of the database changes. In addition, it also helps the user in resuming work from where it was discontinued, after the crash.

## 8. What is the full form of MVCC?

The full form of MVCC is Multi-version Concurrency Control.

## 9. Why do companies use PostgreSQL?

Numerous high-profile organizations, such as Apple, Spotify, IMDb, Instagram, and Skype, make use PostgreSQL database, owing to its excellent features:

- PostgreSQL is extremely easy to use.
- It is a powerful and robust open-source tool.

- PostgreSQL follows and supports the ACID properties.
- It supports MVCC (Multiversion Concurrency Control).
- It is highly fault-tolerant.
- It runs on almost all different operating systems.

## 10. What is the full form of GEQO?

The full form of GEQO is Genetic Query Optimization. It enables non-exhaustive search to efficiently manage large join queries in PostgreSQL.

## 11. What do you mean by index in PostgreSQL?

An index in PostgreSQL is a way of increasing the speed and efficiency of the database. Databases use indexes as special lookup tables that help them retrieve data in a much quicker manner. Indexes enable the user to find specific rows in a database. They act like pointers to the data in the database, thereby enhancing the overall performance.

## 12. What is the main query language of PostgreSQL?

SQL or Structured Query Language is the main query language of PostgreSQL.

## 13. What do you think is the latest PostgreSQL version in the market?

As of 2022, the latest version of PostgreSQL in the market is PostgreSQL 15. It was launched on 13 October, 2022.

## 14. What is the full form of ORDBMS?

The full form of ORDBMS is Object-[Relational Database Management System](#).

## 15. What do you mean by a string constant in PostgreSQL?

A string constant is defined as the sequence of characters that are bounded by single quotes i.e., ('). It can be used during insertion or while passing the characters to the database objects. This is an important feature when performing the parsing of data. In the case of PostgreSQL, string constant is allowed with single quotes but embedded by a C-style backslash.

Example: 'This is an example of a string constant bound by single quotes.'

## PostgreSQL Developer Interview Questions for Experienced

### 16. What is Multi-version Control?

Multi-version Concurrency Control or MVCC is a technique to enhance database performance by handling concurrency in PostgreSQL databases. It prevents the locking of databases. MVCC reduces the delay time that users face while logging into their accounts and comes into action when someone else is accessing the contents of the account.

Inconsistency occurs when numerous transactions attempt to access the same data. To preserve data consistency, concurrency control is necessary.

Let's take an example of an ATM machine. If concurrency is not applied in this case, different users won't be able to access their accounts and draw money at the same time. Whereas if concurrency control is enabled, then multiple users can do so easily.

### 17. Explain table partitioning in PostgreSQL.

Table Partitioning in PostgreSQL is the process wherein a large table is split into smaller pieces. These smaller pieces are known as partitions. List and range partitioning is supported by PostgreSQL through its table inheritance feature.

Table partitioning helps in increasing the query performance of PostgreSQL, as it is much easier to select data from these partitions rather than selecting from one main table.

Each partition can store data according to how frequently it is used, allowing low-use data to be stored on media that may be slower or less expensive.

## 18. Explain the use of PostgreSQL triggers.

A trigger can be defined as a function that is called automatically when the insertion, updation, or deletion event occurs. They serve as a way to check the data integrity. Triggers are capable of handling any errors that occur in the database. Another advantage of triggers is: Any table that is present in a PostgreSQL database can be forced to receive security approvals with the use of PostgreSQL triggers.

## 19. Is PostgreSQL compatible with Cloud?

Yes, PostgreSQL is compatible and can be run on Cloud. PostgreSQL is highly portable. Moreover, similar to other open-source databases, PostgreSQL can be effortlessly executed on virtual containers.

*Are you looking forward to building a career in AWS? Enroll in this professional course to start your journey today! [AWS Certification Training Course for Solutions Architect](#)*





# PostgreSQL Interview Questions for 3 Years Experience



## 20. Name the different types of operators that are used in PostgreSQL.

Operators are the special characters or words that are used mainly in the WHERE clause in PostgreSQL. These operators can be used to perform a variety of functions and operations.

### Operators used in PostgreSQL

-  Arithmetic operators
-  Logical operators
-  Comparison operators
-  Bitwise operators

The different types of operators that are used in PostgreSQL are as follows:-

- Arithmetic operators
- Logical operators
- Comparison operators
- Bitwise operators

## 21. What do you mean by the CTID field in PostgreSQL?

In PostgreSQL, CTIDs serve as unique identifiers for each record within a table. The CTID field enables the precise location of physical rows based on their offset and block positions, facilitating the effective distribution of data across the table. By utilizing CTID fields, users can gain insights into the actual storage positions of rows within the database table.

## 22. How do you start a database server in PostgreSQL?

To access the database data, the first step is to start the database server. The database server application, known as Postgres, plays a crucial role in managing the database. It actively seeks information about the data it needs to utilize, which is essential for its effective functioning. By initiating the database server and providing the necessary data location details, users can actively enable access and utilization of the database data through the Postgres software. The `-D` option is used to accomplish this.

Execute these commands to start the database server:

- `usr/local/etc/rc.d/010.pgsql.sh start`
- `/usr/local/etc/rc.d/postgresql start`

Another way to start a database server in PostgreSQL is:

- Start by pressing the Windows key + R simultaneously to enter the Run Window.
- To find the PostgreSQL services, type `services.msc` next.
- Using the version that is installed, search the Postgres service.
- Click on Start to start the database server.

## PostgreSQL Interview Questions for 5 Years Experience

### 23. State the maximum size of a table on PostgreSQL.

The maximum number of blocks in a table decides the limit of the table. As the number of blocks is  $2^{32}$  and 8192 bytes is the default size of the block, therefore, the maximum size of a table on PostgreSQL is 32TB.

## 24. What are the differences between PostgreSQL and MongoDB?

PostgreSQL	MongoDB
PostgreSQL is a relational database management system.	MongoDB is a non-relational database management system.
PostgreSQL was created using the C language.	MongoDB was created using the C++ language.
PostgreSQL is object-oriented.	MongoDB is document-oriented.
PostgreSQL stores data in the form of different tables.	MongoDB stores data in the form of key-value pairs as one record.
PostgreSQL is faster than MongoDB.	MongoDB is relatively slower than PostgreSQL.

## 25. State the role of tokens in PostgreSQL.

Tokens in PostgreSQL have an important role in the parsing and interpretation of SQL statements. When SQL queries are executed, the PostgreSQL server breaks down the statements into smaller units known as tokens. These tokens represent various elements like keywords, identifiers, literals, operators, and punctuation marks. This tokenization process ensures accurate parsing and evaluation of SQL queries, enabling the server to handle database operations efficiently.

## PostgreSQL Advanced Interview Questions

### 26. When should a developer use PostgreSQL?

Developers should choose PostgreSQL when they require a dependable and feature-rich database management system. PostgreSQL is suitable for diverse applications, including web development, data analysis, and enterprise solutions. It provides advanced features like JSON support and geospatial capabilities. With its cross-platform compatibility and active community, developers can rely on PostgreSQL for scaling and achieving optimal performance while managing complex data.

### 27. Describe the history of PostgreSQL in brief.

PostgreSQL originated as a vital component of the POSTGRES project initiated by Professor Michael Stonebraker in 1986 at the University of California, Berkeley. It boasts compatibility with major operating systems such as macOS, Windows, Linux, and UNIX.

PostgreSQL has upheld ACID properties since 2001, with continuous core platform development spanning over three decades. It encompasses noteworthy additions like the PostGIS database extender and has become the standard [database](#) for macOS. commonly known as Postgres, due to its widespread adoption of the SQL Standard among relational databases.

### 28. Explain the procedure to set up PgAdmin in PostgreSQL.

PgAdmin is a web-based management tool that interacts with the PostgreSQL database. It can be used to perform any database administration operations on PostgreSQL.

To set up PgAdmin in PostgreSQL, follow these steps:

- Start and launch pgAdmin 4.
- Then select “Add new Server” from the “Quick Link” section under the “Dashboard” menu.
- Choose the “Connection” tab in the “Create-Server” box after clicking “Add new Server” in the window.
- Put your server’s IP address in the “Hostname/Address” column to configure the connection.
- Finally, you must define “Port” as “5432,” which is the PostgreSQL server’s default port.

## 29. Explain the role of table space in PostgreSQL.

Table spaces in PostgreSQL are defined as the directories where data files can be stored. They are used to store various databases as well as database objects.

Using table spaces, the disk layout of a PostgreSQL installation can be easily handled and managed.

In addition to that, tablespaces give administrators the ability to enhance performance by making use of their knowledge of the usage patterns of database objects.

## PostgreSQL Interview Questions for 7 Years of Experience

### 30. List the disadvantages of PostgreSQL.

Despite its many advantages, PostgreSQL has numerous disadvantages. Some of these include:

- PostgreSQL may have a slower speed compared to MySQL.
- It supports fewer open-source applications compared to MySQL.
- Market recognition for PostgreSQL has been challenging due to its lack of specific ownership.
- Its performance rate may be lower than that of MySQL in certain situations.

*Check out the blog to know what are the differences between the popular databases i.e. [Postgresql vs. MySQL](#).*

### 31. Explain the term 'Sequence' in PostgreSQL.

The Sequence is a generator that produces a progressive number that can help synchronize the keys across multiple rows or tables and construct a single primary key automatically.

A sequence in PostgreSQL can be defined as a user-defined schema-bound object that generates an integer sequence based on a specific requirement.

## PostgreSQL Interview Questions for 10 Years of Experience

### 32. Differentiate between clustered and non-clustered indexes.

Clustered Index	Non-Clustered Index
It is faster than the non-clustered index.	It is relatively slower as compared to the clustered index.

Index is considered the main data in the clustered index.	In the case of a non-clustered index, the index is the copy of data.
The clustered index has the ability to store data naturally on the disk.	The non-clustered index cannot naturally store data on the disk.
It requires less memory for operations as compared to the non-clustered index.	The non-clustered index requires more memory to perform operations.
A table can consist of only one clustered index.	A table can contain multiple non-clustered indexes.

### 33. What is the procedure for storing binary data in PostgreSQL?

The users can store binary data in PostgreSQL in two distinct ways:

- By using the data type BYTEA.
- By using the Large Object feature.

### 34. What do you understand about the enable-debug command in PostgreSQL?

The enable-debug command in PostgreSQL is the command that assists in compiling all libraries and applications.

It has a few debugging symbols that make it easier for developers to find flaws and other issues that can arise during the script's execution. This process can slow

down or impede the system when it is being used, increasing the size of the binary file.

### 35. Describe the method by which you can change the column data type in PostgreSQL.

The data type of one or more columns in PostgreSQL can be changed by using the following commands along with the TYPE keyword:

```
1 ALTER TABLE
```

```
2 ALTER COLUMN
```

Example:

```
1 ALTER TABLE tab_name
```

```
2 ALTER COLUMN col_name TYPE new_data_type;
```

## PostgreSQL DBA Interview Questions For Experienced Professionals

### 36. How can the first 5 records be selected in PostgreSQL?

The LIMIT keyword can be used to select the first N records in PostgreSQL.

Example:



```
1 SELECT * FROM Employee ORDER BY Salary DESC LIMIT 5
```

Here, the Employee is the name of the table that contains employee data.

ORDER BY command arranges the data in descending order based on the salary of employees.

LIMIT keyword used with the number 5 prints the first 5 or the top 5 records present in the Employee table.

## 37. What are the features of PostgreSQL?

PostgreSQL or Postgres is an object-relational database management system or ORDBMS. Some of its prominent features are as follows:-

- Extremely high fault-tolerance
- Free to download
- Reliable and secure
- Robust and powerful
- Easy recovery process
- Low maintenance cost
- Easily compatible with a wide variety of platforms and languages.
- High availability
- Easy to use

## 38 Does PostgreSQL support Full-Text Search?

When a search is conducted on a portion of text contained in a large body of electronically recorded text, it is referred to as a full-text search, the results that are returned may include all or some of the search terms. Traditional searches, however, would only produce exact matches.

Yes, PostgreSQL supports the Full-Text Search feature. It is a powerful tool in PostgreSQL and can be enhanced by incorporating functions like result highlighting or by creating your own unique dictionaries or functions.

## PostgreSQL Replication Interview Questions

### 39. What are the physical methods used in PostgreSQL for replication?

A few replication methods are:

Physical Replication: This method includes two ways:

- Streaming Replication: In this replication, there is a constant streaming of changes from a primary server to one or more standby PostgreSQL servers.
- File-based Replication (pg\_basebackup): In this replication, a base backup is created of the primary server's directory and forwarded to standby servers.

### 40. How many types of replications are present in PostgreSQL?

There are majorly four types of PostgreSQL replications, namely:

- Physical Replication: It is of two types:
  1. Streaming Replication
  2. File-based Replication (pg\_basebackup)
- Logical Replication: It consists of:
  1. Build-In Logical Replication
  2. Third-party Tools

- Bi-Directional Replication (BDR)
- Custom Replication Solutions.

## 41. How can replication conflicts be addressed in PostgreSQL?

When data changes occur at the same time in different places, such as multi, which may lead to replication conflicts, the software automatically ensures the change that should be allowed with the help of “conflict handlers” in PostgreSQL. These conflict handlers automatically resolve the changes applicable to avoid conflicts.

## 42. What are the advantages of logical replication as compared to physical replication methods?

Advantages of logical replications are:

1. Selective replication is possible in logical replication, which makes it more flexible to choose which data to replicate.
2. Physical replication requires identical PostgreSQL versions on both primary servers and standby servers, whereas logical replication can work with different versions of PostgreSQL.
3. Logical replication offers more flexibility as compared to physical replication in handling the changes in Schema.
4. Logical replication is compatible with multiple applications and database maintenance tasks.

## 43. What is the role of replication slots in PostgreSQL's streaming replication mechanism?

Replication slots within the streaming replication mechanism of PostgreSQL guarantee that the primary server will keep the Write Ahead Logging (WAL) files

until they are received and applied by all servers. This ensures that the primary server will not delete these files prematurely.

## PostgreSQL Backup Interview Questions

### 44. How can you take the backup of a database?

PostgreSQL permits the user to take a backup of the database by using “pg\_dump”.

To perform a backup on a plain-text SQL file, login into your database server and implement the following command:

```
1 pg_dump database_name > filename.sql
```

The database can be reconstructed using the commands available in the SQL file.

Another way to backup the database is:

```
1 /usr/local/bin/pg_dump mydatabase > mydatabase.pgdump
```

### 45. How can you stop a PostgreSQL Server? Can you stop a particular database in the PostgreSQL cluster?

To stop a PostgreSQL server implement the following steps and commands:

- For Windows

The first step is to locate the PostgreSQL database directory.

After that, open the command prompt and execute the following command-

```
1 pg_ctl -D "C:\Program Files\PostgreSQL\9.6\data" stop
```

An alternative way to stop the PostgreSQL server on Windows is:

- Press the Windows key + R simultaneously to enter the Run Window.
- Type services.msc to find the PostgreSQL services.
- Using the installed version, locate the Postgres service.
- Click Stop to stop the database server.
- For Linux

Use the following command on Linux to stop the server-

```
1 sudo service postgresql stop
```

- For macOS

Use the following command on macOS to stop the server manually-

```
1 pg_ctl -D /usr/local/var/postgres stop
```

No, PostgreSQL does not allow the user to stop a specific database in the cluster.

## 46. Discuss the differences between file-level and logical backups in PostgreSQL

When it comes to backing up a PostgreSQL database cluster, there are two approaches: file-level backups and logical backups.

File-level backups involve making copies of the files that make up the database cluster. This includes data files, configuration files, and transaction logs. This method is great for disaster recovery situations where you need to restore the

database cluster. However, because it copies all the files, the backup files can be quite large. Require storage space.

On the other hand, logical backups export data in a format using tools like `pg_dump`. These backups are smaller in size. Offer flexibility when it comes to selectively restoring specific databases, tables, or subsets of data. However, they can be slower compared to file-level backups when it comes to both backup and restore operations

## 47. Discuss the advantages and limitations of using `pg_dump` for backups.

The advantages of using `pg_dump` for backups are:

- Portability: Backup files can be effortlessly run on PostgreSQL installations.
- Selective Backup: It enables the backup of databases, schemas, tables, or rows as needed.
- Comprehensive: It captures both the database schema and data, ensuring a backup.
- Integration: `pg_dump` can be seamlessly integrated into automation scripts for operation.

The limitations of using `pg_dump` for backups are:

- Impact on performance: It is possible that there could be some load on the database server, particularly when dealing with large databases.
- Storage requirements: Backups might take up an amount of space and may also require more time for transferring.
- Time required for restoration: Restoring from `pg_dump` backups can be a time-consuming process, potentially resulting in downtime.
- Limited parallelism: By default, the operation is performed using a thread, which could potentially prolong the duration.

## 48. How can you perform a point-in-time recovery in PostgreSQL? How can you perform a point-in-time recovery in PostgreSQL?

To perform a point-in-time recovery in PostgreSQL, you need:

1. Make sure that WAL archiving is enabled.
2. Restore the base backup.
3. Apply the WAL files either by using recovery.conf or via a server.
4. Stop the recovery process when you reach the recovery point you desire.
5. Allow read-write operations on the database.

## 49. How can you perform a selective restore of data from a PostgreSQL backup?

To selectively restore data from a backup of PostgreSQL, follow these steps:

1. Utilize the `pg_restore` command, Include the `t` option to indicate the table(s) you wish to restore.
2. If necessary, you can also use the `n` option to specify schemas if the tables belong to them.
3. Execute the `pg_restore` command by providing the file and any additional options that may be required.
4. Keep an eye on the restoration process. Ensure that you verify the restored data once it is completed.

## 50. What is the purpose of the `pg_archivecleanup` utility in PostgreSQL?

The main objective of the `pg_archivecleanup` tool in PostgreSQL is to get rid of WAL (Write Ahead Logging) files from the directory. This ensures that only the necessary WAL files required for Point-In-Time-Recovery (PITR) are kept intact.

## Global Postgre Salary Trends

Experience Level	Average Salary (India: !NR)	US Salaries (in USD)
Entry Level	₹4- 6 lakhs	65k-85k
Intermediate	₹6- 12 lakhs	85k-110k
Senior	₹12- 20 lakhs	110k-160k

## PostgreSQL Job Trends

Global Scenario PostgreSQL: The world job landscape for PostgreSQL experts is increasing, boasting over **18,000 active job openings** globally. This surge underlines the rising demand for professionals with robust database management capabilities. In India, the market is also flourishing with openings on platforms like **Indeed** and **LinkedIn**.

India's PostgreSQL Boom; PostgreSQL is making waves, with a notable uptrend in job openings. The rise is attributed to local demand and the global recognition of **India's tech prowess**, making it a hotspot for PostgreSQL professionals who are keen on capitalizing on diverse opportunities.

PostgreSQL is not just a skill; it's your gateway to an evolving global job market.

## Postgre SQL : Roles and Responsibilities



Job Role	Description
Database Administrator (DBA)	The DBA is responsible for the overall security and maintenance of databases.
Database Developer	They are specialized in designing and implementing database structures.
Data Engineer	They are involved in the design and maintenance of data pipelines and ETL.
Database Architect	They ensure that the database architecture aligns with business requirements.
Systems Administrator	They are focused on installation configuration and PostgreSQL Servers.
Application Developer	They integrate database functionality with software applications.
DevOps Engineer	They are responsible for streamlined software development and deployment further in the context of PostgreSQL they automate database deployments and scalability.
Data Analyst	They are responsible for extracting meaningful insights for better decision-making.

Business Intelligence (BI) Developer:	They use PostgreSQL as a data source to create interactive dashboards and reports for business stakeholders.
Technical Support Engineer:	They assist the users in resolving the issues by guiding best practices.

The roles and responsibilities of a PostgreSQL are quite diverse and vary for different departments. You can check it with JD of [LTIMindtree](#). Some of the most important roles and responsibilities of a PostgreSQL are as follows:

- Designing the optimal database schemas and models based on application requirements and queries. Choose suitable data types, keys, relations, partitions, etc.
- Write efficient, high-performance SQL queries using best practices around joins, subqueries, CTEs, etc.
- Improve pre-existing query execution plans using indexes, optimizer hints, and better data models to improve runtime and scalability.
- Track key metrics like connection utilization, IOPS, RAM usage, and WAL throughput to plan infrastructure sizing
- Set up partitions, and retention policies to manage tremendous data volumes
- Discover and resolve bottlenecks due to locks, memory limits
- Taking care of user access control, security policies, backups, recovery, standard production
- Upgrade PostgreSQL versions to leverage new optimizations with minimal downtime

## Conclusion

Whether you're a database administrator, developer, data analyst, or work in a DevOps role, mastering PostgreSQL can enhance your skill set and contribute to the improvement of workflows in a variety of professional settings. The ability to work with databases effectively is a valuable skill in today's data-driven world. Salary differs according to the role and key responsibilities. A major factor in deciding salary is your skillset and years of experience in the respective domains.

*Join our [SQL community](#) to check out our community page to see what other learners are learning.*