

Mindtree SQL Interview Questions

1. What is SQL?

SQL stands for Structured Query Language. It is used to communicate with databases. SQL allows you to query, insert, update, and delete data.

2. What are the different types of SQL statements?

SQL statements can be categorized into:

- DDL (Data Definition Language): CREATE, ALTER, DROP
- DML (Data Manipulation Language): SELECT, INSERT, UPDATE, DELETE
- DCL (Data Control Language): GRANT, REVOKE
- TCL (Transaction Control Language): COMMIT, ROLLBACK, SAVEPOINT

3. What is a JOIN in SQL?

A JOIN clause is used to combine rows from two or more tables. It is based on a related column between them. Common types include INNER JOIN, LEFT JOIN, RIGHT JOIN, and FULL JOIN.

4. What is the difference between INNER JOIN and LEFT JOIN?

- INNER JOIN: Returns only the matching rows from both tables.
- LEFT JOIN: Returns all rows from the left table and matching rows from the right table. If no match is found, NULL values are returned.

5. How do you optimize a SQL query?

- Use Indexes: Indexes speed up the retrieval of rows.
- **Avoid SELECT* : Specify only the columns you need.
- Use WHERE clause to filter data: Reduces the number of rows processed.
- Consider query execution plan: It helps you understand how the query is executed.

6. What is a subquery?

A subquery is a query within another query. It is enclosed in parentheses. Subqueries can be used in **SELECT**, **INSERT**, **UPDATE**, and **DELETE** statements.

7. What is the difference between **WHERE** and **HAVING** clauses?

- **WHERE:** Filters rows before the grouping is done.
- **HAVING:** Filters groups after the grouping is done. Used with **GROUP BY**.

8. What is a primary key?

A primary key uniquely identifies each record in a table. It must contain unique values and cannot contain **NULLs**.

9. What is a foreign key?

A foreign key is a field in one table that links to the primary key in another table. It helps maintain referential integrity between the tables.

10. How do you handle **NULL** values in SQL?

- **IS NULL:** Check for **NULL** values.
- **COALESCE():** Replace **NULL** with a specified value.
- **NULLIF():** Returns **NULL** if two expressions are equal.

11. What is normalization?

Normalization is the process of organizing data in a database. The goal is to reduce redundancy and improve data integrity.

12. Explain the different normal forms.

- **1NF (First Normal Form):** Ensures that the table is free of repeating groups.
- **2NF (Second Normal Form):** No partial dependency of any column on the primary key.
- **3NF (Third Normal Form):** No transitive dependency for non-prime attributes.

13. What are indexes in SQL?

Indexes are used to speed up the retrieval of data. They work like a table of contents in a book. Indexes can be created on one or more columns of a table.

14. Explain the concept of ACID properties.

- **Atomicity:** Transactions are all-or-nothing.
- **Consistency:** Transactions bring the database from one valid state to another.
- **Isolation:** Concurrent transactions do not interfere with each other.
- **Durability:** Once a transaction is committed, it remains so even in case of a failure.

15. What is a stored procedure?

A stored procedure is a prepared SQL code that

16. What is a trigger in SQL?

A trigger is a special kind of stored procedure that automatically runs when specific events occur in a database. Triggers are commonly used to maintain the integrity of the data by enforcing rules or restrictions.

17. What is the difference between DELETE and TRUNCATE?

- **DELETE:** Removes specific rows from a table based on a condition. It can be rolled back, and it maintains the integrity of the table's structure.
- **TRUNCATE:** Removes all rows from a table without logging individual row deletions. It cannot be rolled back and is faster than DELETE.

18. What is a view in SQL?

A view is a virtual table based on the result-set of a query. It does not store data but fetches data dynamically from the underlying tables when queried. Views can simplify complex queries and provide a layer of abstraction.

19. What is a cursor in SQL?

A cursor is a database object used to retrieve, manipulate, and navigate through a result set one row at a time. Cursors are useful when row-by-row operations are necessary, but they can be slower compared to set-based operations.

20. What is the difference between UNION and UNION ALL?

- **UNION:** Combines the result sets of two queries and removes duplicate rows.
- **UNION ALL:** Combines the result sets of two queries but includes duplicate rows.

21. What is the difference between CHAR and VARCHAR data types?

CHAR and VARCHAR are string data types in SQL

- **CHAR:** Fixed length; padded with spaces.
- **VARCHAR:** Variable length; no padding.

22. What is a database index?

A database index enhances the speed of data retrieval operations. It is created on columns of a table and functions like a lookup table to facilitate quicker access to data.

23. How do you create a table in SQL?

- Use the **CREATE TABLE** statement.
- Define columns and data types.
- Specify primary keys and constraints.

24. What is the purpose of the GROUP BY clause?

- **GROUP BY** groups rows with the same values.
- Used with aggregate functions.
- Helps summarize data.

[Click here to ace Mindtree SQL Interview Questions by learning the fundamentals!](#)

25. How do you delete all records from a table?

To delete all records from a table, use the **DELETE** statement without a **WHERE** clause. This will remove all rows but keep the table structure intact.

26. What is a self-join?

A self-join involves joining a table with itself. This type of join is used to combine rows from the same table, often requiring table aliases to differentiate between the instances.

27. What is the difference between UNION and INTERSECT?

- **UNION:** Combines all unique results from queries.
- **INTERSECT:** Returns common results between queries.

28. What is the purpose of the HAVING clause?

- **HAVING** filters results after **GROUP BY**.
- Used to apply conditions on aggregated data.
- Works with **GROUP BY** in queries.

29. How do you update a record in SQL?

- Use the **UPDATE** statement.
- Specify the table and columns to update.
- Example: **UPDATE table_name SET column1 = value WHERE condition;**

30. What is a composite key?

- A composite key is a primary key made of multiple columns.
- It uniquely identifies records in a table.
- Used when a single column is insufficient.

31. What are aggregate functions in SQL?

- Aggregate functions perform calculations on a set of values.
- Examples include **COUNT**, **SUM**, **AVG**, **MIN**, **MAX**.
- They return a single value based on the set.

32. How do you use the CASE statement in SQL?

- The CASE statement provides conditional logic.
- It returns values based on conditions.
- Example: CASE WHEN condition THEN result ELSE default END

33. What is the difference between RANK() and ROW_NUMBER()?

- RANK(): Assigns rank with gaps for ties.
- ROW_NUMBER(): Assigns unique sequential numbers.

34. How do you find duplicate rows in SQL?

To find duplicate rows, use **GROUP BY** along with the **HAVING** clause. Group by the columns of interest and count occurrences. Rows with a count greater than one are duplicates.

Example: **SELECT column1, column2, COUNT(*) FROM table_name GROUP BY column1, column2 HAVING COUNT(*) > 1;**

35. What is the difference between an implicit join and an explicit join?

- Implicit Join: Uses comma-separated tables in **FROM** clause.
- Explicit Join: Uses **JOIN** keywords to specify the type.

36. What is a recursive query?

A recursive query refers to itself to process hierarchical or nested data. It is often implemented using Common Table Expressions (CTEs) to handle recursive relationships within the data.

37. How do you prevent SQL injection attacks?

- Use parameterized queries or prepared statements.
- Validate and sanitize user inputs.
- Employ proper error handling and permissions.

38. What is a SQL view?

A SQL view is a virtual table created from a query's result set. It does not physically store data but provides a dynamic and simplified view of the underlying tables, which can be used to simplify complex queries.

39. What is a materialized view?

A materialized view stores the results of a query physically, improving performance by precomputing and caching the data. Unlike regular views, materialized views need to be refreshed periodically to stay updated with the underlying data.

40. What is a transaction in SQL?

- A transaction is a sequence of SQL operations.
- It is treated as a single unit of work.
- Ensures all operations succeed or none at all (ACID properties).

41. How do you handle exceptions in SQL?

Exception handling in SQL can be managed using TRY...CATCH blocks. This approach allows you to capture and handle errors, providing alternate code execution or logging when an error occurs. In SQL Server, it is implemented using BEGIN TRY...END TRY and BEGIN CATCH...END CATCH.

42. What is a common table expression (CTE)?

- A CTE is a temporary result set.
- It is defined using the WITH keyword.
- Useful for simplifying complex queries.

43. What is a pivot table in SQL?

- A pivot table rotates data to present it in a summarized form.
- Converts rows into columns for easier analysis.
- Implemented using PIVOT operations or CASE statements.

44. How do you perform a full outer join in SQL?

- FULL OUTER JOIN returns all rows from both tables.

- Includes rows with no match in both tables.
- Example: `SELECT * FROM table1 FULL OUTER JOIN table2 ON table1.id = table2.id;`

45. What is a temporary table in SQL?

A temporary table is a table that exists only for the duration of a session or transaction. It is used for storing intermediate results or temporary data, and it is automatically dropped when the session or transaction ends.

46. What is the purpose of the SQL COALESCE() function?

- COALESCE() returns the first non-NULL value.
- Useful for handling NULL values in queries.
- Example: `COALESCE(column_name, 'default_value')`

47. What is the difference between a correlated subquery and a non-correlated subquery?

A correlated subquery references columns from the outer query and is executed once for each row processed by the outer query. A non-correlated subquery is independent of the outer query and is executed only once.

48. How do you use the LIMIT clause in SQL?

- LIMIT restricts the number of rows returned.
- It is used to paginate results.
- Example: `SELECT * FROM table_name LIMIT 10;`

49. What is the purpose of the DISTINCT keyword?

- DISTINCT removes duplicate rows from results.
- It ensures unique values in the output.
- Used with SELECT statements.

50. What is a SQL alias, and how is it used?

- An alias is a temporary name for a table or column.
- It simplifies query writing and improves readability.



- Created using the AS keyword (optional).

