

ADBMS VIVA Q&A

1. What is a database?

Ans. A database is a structured collection of data for faster and better access, storage, and manipulation of data. A database can also be defined as a collection of tables, schema, views, and other database objects.

2. What is DBMS?

Ans. It is a collection of programs that enables user to create and maintain a database. In other words it is general-purpose software that provides the users with the processes of defining, constructing and manipulating the database for various applications.

3. What is RDBMS?

Ans. An RDBMS or Relational Database Management System is a type of DBMS having relationships between the tables using indexes and different constraints like primary key, foreign key, etc. The use of indexes and constraints helps in faster retrieval and better management of data within the databases.

4. What is the difference between DBMS and RDBMS?

Ans. The primary difference between DBMS and RDBMS is, in RDBMS we have relations between the tables of the database.

Whereas in DBMS there is no relation between the tables (data may even be stored in files).

RDBMS has primary keys and data is stored in tables. DBMS has no concept of primary keys with data stored in navigational or hierarchical form.

RDBMS defines integrity constraints in order to follow ACID properties. While DBMS doesn't follow ACID properties.

5. What is a table?

Ans. A table is a database object that is used for storing related records in the form of rows and columns.

6. What is a field in a table?

Ans. A field is an entity used for storing a particular type of data within a table like numbers, characters, dates, etc.

7. What is a tuple, record, or row in a table?

Ans. A tuple is a single record or single row in a table. For example, in the case of a student table, the data corresponding to a student will be called a tuple or record of that student.

8. What is SQL?

Ans. SQL stands for Structured Query Language. It is a language used for creating, storing, fetching, and updating data and database objects in RDBMS.

9. What are the different types of SQL commands?

Ans. SQL commands are the set of commands used to communicate and manage the data present in the database. The different type of SQL commands are-

1. DDL – Data Definition Language
2. DML – Data Manipulation Language
3. DCL – Data Control Language
4. TCL – Transactional Control Language

10. Explain DDL commands. What are the different DDL commands in SQL?

Ans. DDL refers to Data Definition Language. The DDL commands are used to define or alter the structure of the database. The different DDL commands are-

- CREATE – Used to create a table in the DB
- DROP – Drops the table from the DB
- ALTER – Alters the structure of the DB
- TRUNCATE – Deletes all the records from the DB but not its database structure
- RENAME – Renames a DB object

11. Explain DML commands. What are the different DML commands in SQL?

Ans. DML refers to Data Manipulation Language. These commands are used for managing data present in the database. Some of the DML commands are – select, insert, update, delete, etc.

12. Explain DCL commands. What are the different DCL commands in SQL?

Ans. DCL refers to Data Control Language. These commands are used to create roles, grant permission, and control access to the database objects. The three DCL commands are-

- GRANT – Grants permission to a database user.
- REVOKE – Removes access privileges from a user-provided with the GRANT command.
- Deny – Explicitly prevents a user from receiving particular permission(e.g. preventing a particular user belonging to a group to receive the access controls.

13. Explain TCL commands. What are the different TCL commands in SQL?

Ans. TCL refers to Transaction Control Language. These commands are used to manage the changes made by DML statements. These are used to process a group of SQL statements comprising a logical unit. The three TCL commands are-

- COMMIT – Commit write the changes to the database
- SAVEPOINT – Savepoints are the breakpoints, these divide the transaction into smaller logical units which could be further roll-backed.
- ROLLBACK – Rollbacks are used to restore the database since the last commit.

14. What are SQL constraints?

Ans. SQL constraints are the set of rules that impose some restrictions while inserting, deleting, or updating the data in the databases. In SQL, we have both column level as well as table level constraints which are applied at columns and tables respectively. Some of the constraints in SQL are – Primary Key, Foreign Key, Unique Key, Not NULL, DEFAULT, CHECK, and Index constraint.

15. What is a Unique constraint?

Ans. A unique constraint is used to ensure that a field or column will have only a unique value (no duplication).

16. What is a Primary Key?

Ans. A primary key is a column or a combination of columns that uniquely identifies a record in the database. A

primary key can only have unique and not NULL values and there can be only one primary key in a table.

17. What is the difference between a unique key and a primary key?

Ans. A unique key allows null value (although only one) but a primary key doesn't allow null values. A table can have more than one unique keys columns while there can be only one primary key. A unique key column creates a non-clustered index whereas the primary key creates a clustered index on the column.

18. What is a composite key?

Ans. A composite key is a primary key with multiple columns as in the case of some tables a single field might not guarantee unique and not null values. So a combination of multiple fields is taken as the primary key.

19. What is a NULL value?

Ans. A NULL value in **SQL** is an unknown or blank value. Since NULL is an unknown value so the NULL value cannot be compared with another NULL value. Hence we cannot use the '=' operator in where condition with NULL. For this, we have an IS NULL clause that checks if the value in the field is NULL or not.

20. What is a Not NULL constraint?

Ans. A Not NULL constraint is used for ensuring that the value in the field cannot be NULL.

21. What is a Foreign Key?

Ans. A foreign key is used for enforcing referential integrity in which a field marked as a foreign key in one table is linked with a primary key of another table. With this referential integrity, we can have only the data in the foreign key which matches the data in the primary key of the other table.

22. What is a Check constraint?

Ans. A check constraint is used to limit the value entered in a field. For example, we can ensure that field 'Salary' can only have a value greater than 1000 using check constraint.

```
CREATE TABLE EMP_SALARY(  
EmpID int NOT NULL,  
NAME VARCHAR (30) NOT NULL,  
Salary INT CHECK (Salary > 1000),  
PRIMARY KEY (EmpID)  
);
```

23. What is a Default constraint?

Ans. A Default constraint is used for providing a default value to a column when no value is supplied at the time of insertion of record in the database.

24. What is a clustered index?

Ans. A clustered index defines the order in which data is physically stored in the database tables. A table can have only one clustered index.

25. What is a non-clustered index?

Ans. In the case of a non-clustered index, the data and the indices are stored in different locations. A non-clustered index has a jump table containing key-values with pointers pointing to the location containing the actual rows. There can be multiple clustered indexes in a table.

26. What is the difference between delete, truncate and drop command?

Ans. The difference between the Delete, Truncate and Drop command is –

- Delete command is a DML command. It removes rows from a table based on the condition specified in the where clause, being a DML statement we can rollback changes made by the delete command.
- Truncate is a DDL command. It removes all the rows from the table and also frees the space held, unlike the delete command. It takes a lock on the table while the delete command takes a lock on rows of the table.
- Drop is a DDL command. It removes the complete data along with the table structure (unlike the truncate command that removes only the rows).

27. What are the different types of joins in SQL?

Ans. Joins are used to combine records from multiple tables. The different types of joins in SQL are-

1. Inner Join – To fetch rows from two tables having matching data in the specified columns of both the tables.
`SELECT * FROM TABLE1 INNER JOIN TABLE2 ON TABLE1.columnA = TABLE2.columnA;`
2. Left Join – To fetch all rows from the left table and matching rows of the right table
`SELECT * FROM TABLE1 LEFT JOIN TABLE2 ON TABLE1.columnA = TABLE2.columnA;`
3. Right Join – To fetch all rows from right table and matching rows of the left table
`SELECT * FROM TABLE1 RIGHT JOIN TABLE2 ON TABLE1.columnA = TABLE2.columnA;`
4. Full Outer Join – To fetch all rows of the left table and all rows of right table
`SELECT * FROM TABLE1 FULL OUTER JOIN TABLE2 ON TABLE1.columnA = TABLE2.columnA;`
5. Self Join – Joining a table to itself, for referencing its own data
`SELECT * FROM TABLE1 T1, TABLE1 T2 WHERE T1.columnA = T2.columnB;`

28. What is the difference between cross join and full outer join?

Ans. A cross join returns the cartesian product of the two tables. So there is no condition or on clause as each row of TableA is joined with each row of TableB whereas a full outer join will join the two tables on the basis of the condition specified in the on clause and for the records not satisfying the condition null value is placed in the join result.

29. What is the difference between where and having clause?

Ans. This is one of the commonly asked DBMS interview questions in which the interviewer wants to check your knowledge of the 'where' and 'having' clause along with their usage.

A 'where' clause is used to fetch data from the database that specifies particular criteria (specified after the where clause). Whereas a 'having' clause is used along with 'GROUP BY' to fetch data that meets particular criteria specified by the aggregate function.

For example – for a table with Employee and Project fields. If we want to fetch Employee working on a particular project P2, we will use 'where' clause-

```
Select Employee
From Emp_Project
where Project = P2;
```

Now if we want to fetch Employees who are working on more than one project. We will first have to group the Employee column along with the count of the project and then the 'having' clause can be used to fetch relevant records-

```
Select Employee
From Emp_Project
GROUP BY Employee
Having count(Project)>1;
```

30. What is the difference between Union and Union All command?

Ans. The fundamental difference between Union and Union All command is – Union is by default distinct i.e. it combines the distinct result set of two or more select statements. Whereas, Union All combines all the rows including duplicates in the result set of different select statements.

31. Define the 'Select into' statement.

Ans. Select into statement is used to directly select data from one table and insert it into another. The new table gets created with the same name and type as of the old table-

```
SELECT * INTO newtable FROM oldTable;
```

32. What is a View in SQL?

Ans. A view is a virtual table. It is a named set of SQL statements that can be later referenced and used as a table.

```
CREATE VIEW VIEW_NAME AS
SELECT COLUMN1, COLUMN2
FROM TABLE_NAME WHERE CONDITION;
```

33. Can we use 'where' clause with 'GROUP BY'?

Ans. Yes, we can use 'where' clause with 'GROUP BY'. The rows that don't meet the where conditions are removed first and then the grouping is done based on the GROUP BY column.

```
SELECT Employee, Count(Project )
```

```
FROM Emp_Project
WHERE Employee != 'A'
GROUP BY Project;
```

34. What is Database Normalisation?

Ans. Database normalization is the process of organization of data in order to reduce the redundancy and anomalies in the database. We have different Normalisation forms in SQL like – First Normal Form, Second Normal Form, Third Normal Form, and BCNF.

35. Explain the First Normal Form(1NF).

Ans. According to First Normal Form, a column cannot have multiple values. Each value in the columns must be atomic.

36. Explain the Second Normal Form(2NF).

Ans. For a table to be considered in Second Normal Form. It must follow 1NF and no column should be dependent

on the primary key.

37. Explain the Third Normal Form(3NF).

Ans. For a table to be Third Normal Form, it must follow 2NF and each non-prime attribute must be dependent on the primary key of the table.

For each functional dependency $X \rightarrow Y$ either-

X should be the super key or Y should be the prime attribute (part of one of the candidate keys) of the table.

38. Explain Boyce Codd Normal Form(BCNF).

Ans. BCNF is the advanced or stricter version of 3NF.

For each functional dependency $X \rightarrow Y$, X should be a super key.

39. What are transactions in SQL?

Ans. Transactions are a set of operations performed in a logical sequence. It is executed as a whole. If any statement in the transaction fails, the whole transaction is marked as failed and not committed to the database.

40. What are the ACID properties?

Ans. This is one of the frequently asked database interview questions. ACID properties refer to the four properties of transactions in SQL-

1. Atomicity – All the operations in the transaction are performed as a whole or not performed at all.
2. Consistency – The state of the database changes only on successfully committed transactions.
3. Isolation – Even with concurrent execution of multiple transactions. The final state of the DB would be the same as if transactions got executed sequentially. In other words, each transaction is isolated from one another.
4. Durability – Even in the state of the crash or power loss the state of the committed transaction remains persistent.

41. What are locks in SQL?

Ans. Locks in SQL are used for maintaining database integrity in case of concurrent access to the same piece of data.

42. What are the different types of locks in the database?

Ans. The different types of locks in the database are-

1. Shared locks – Allows data to be read-only(Select operations), and prevents the data to be updated when in the shared lock.
2. Update locks – Applied to resources that can be updated. There can be only one update lock on data at a time.
3. Exclusive locks – Used to lock data being modified (INSERT, UPDATE, or DELETE) by one transaction. Thus ensuring that multiple updates cannot be made to the same resource at the same time.
4. Intent locks – A notification mechanism using which a transaction conveys that intends to acquire a lock on data.
5. Schema locks – Used for operations when the schema or structure of the DB is required to be updated.
6. Bulk Update locks – Used in case of bulk operations when the TABLOCK hint is used.

43. What are the aggregate functions in SQL?

Ans. Aggregate functions are SQL functions that return a single value calculated from multiple values of columns. Some of the aggregate functions in SQL are-

- Count() – Returns the count of the number of rows returned by the SQL expression
- Max() – Returns the max value out of the total values
- Min() – Returns the min value out of the total values
- Avg() – Returns the average of the total values
- Sum() – Returns the sum of the values returned by the SQL expression

44. What are scalar functions in SQL?

Ans. Scalar functions are the functions that return a single value by processing a single value in SQL. Some of the widely used SQL functions are-

- UCASE() – used to convert a string to upper case
- LCASE() – used to convert a string to lowercase
- ROUND() – used to round a number to the decimal places specified
- NOW() – used to fetch the current system date and time
- LEN() – used to find the length of a string
- SUBSTRING() or MID() – MID and SUBSTRING are synonyms in SQL. They are used to extract a substring from a string by specifying the start and end index. Syntax – SUBSTRING(ColumnName, startIndex, endIndex).
- LOCATE() – used to find the index of the character in a string. Syntax – LOCATE(character, ColumnName)
- LTRIM() – used to trim spaces from the left
- RTRIM() – used to trim spaces from the right

45. What is a coalesce function?

Ans. Coalesce function is used to return the first, not NULL value out of the different values or expressions passed to the coalesce function as parameters. Example- COALESCE(NULL, NULL, 5, 'ArtOfTesting') will return the value 5. COALESCE(NULL, NULL, NULL) will return NULL value as no not NULL value is encountered in the parameters list.

46. What are cursors in SQL?

Ans. Cursors are objects in SQL that are used to traverse the result set of a SQL query one by one.

47. What are the stored procedures? Explain their advantages.

Ans. Stored procedures are SQL procedures (a bunch of SQL statements) that are stored in the database and can be called by other procedures, triggers, and other applications.

```
CREATE PROCEDURE
procedureName
AS
Begin
    Set of SQL statements
End
```

The advantages of stored procedure are-

1. Stored procedures improve performance as the procedures are pre-compiled as well as cached.
2. Make queries easily maintainable and reusable as any change is required to be made at a single location.
3. Reduce network usage and traffic.
4. Improve security as stored procedures restrict direct access to the database.

48. What are the triggers in SQL?

Ans. Triggers are special types of stored procedures that get executed when a specified event occurs. Syntax-

CREATE TRIGGER

triggerName

triggerTime{Before or After}
 triggerEvent{Insert, Update or Delete}
 ON tableName
 FOR EACH ROW
 triggerBody

49. What are orphan records?

Ans. Orphan records are records having a foreign key to a parent record that doesn't exist or got deleted.

50. How can we remove orphan records from a table?

Ans. In order to remove orphan records from the database. We need to create a join on the parent and child tables and then remove the rows from the child table where id IS NULL.

51. Which part of the RDBMS takes care of the data dictionary? How?

Ans. Data dictionary is a set of tables and database objects that is stored in a special area of the database and maintained exclusively by the kernel.

52. What is the job of the information stored in data-dictionary?

Ans. The information in the data dictionary validates the existence of the objects, provides access to them, and maps the actual physical storage location.

53. How do you communicate with an RDBMS?

Ans. You communicate with an RDBMS using Structured Query Language (SQL).

54. Define SQL and state the differences between SQL and other conventional programming Languages.

Ans. SQL is a nonprocedural language that is designed specifically for data access operations on normalized relational database structures. The primary difference between SQL and other conventional programming languages is that SQL statements specify what data operations should be performed rather than how to perform them.

55. Name the three major set of files on disk that compose a database in Oracle.

Ans. There are three major sets of files on disk that compose a database. All the files are binary. These are 1.) Database files 2.) Control files 3.) Redo logs. The most important of these are the database files where the actual data resides. The control files and the redo logs support the functioning of the architecture itself. All three sets of files must be present, open, and available to Oracle for any data on the database to be useable. Without these files, you cannot access the database, and the database administrator might have to recover some or all of the database using a backup, if there is one.

56. What is database Trigger?

Ans. A database trigger is a PL/SQL block that can be defined to automatically execute for insert, update, and delete statements against a table. The trigger can be defined to execute once for the entire statement or once for every row that is inserted, updated, or deleted. For any one table, there are twelve events for which you can define database triggers. A database trigger can call database procedures that are also written in PL/SQL.

57. What is Transaction Manager?

Ans. It is a program module, which ensures that the database remains in a consistent state despite system failures and concurrent transaction execution proceeds without conflicting. 51. What is File Manager? It is a program module, which manages the allocation of space on disk storage and data structure used to represent information stored on a disk.

58. What is a query?

Ans. A query with respect to DBMS relates to user commands that are used to interact with a data base. The query language can be classified into data definition language and data manipulation language.

59. What is an Entity?

Ans. It is a 'thing' in the real world with an independent existence.

60. What is an Entity type?

Ans. It is a collection (set) of entities that have same attributes.