

DBMS VIVA QUESTIONS

1. What is database?

A database is a logically coherent collection of data with some inherent meaning, representing some aspect of real world and which is designed, built and populated with data for a specific purpose.

2. What is DBMS?

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 a Database system?

The database and DBMS software together is called as Database system.

4. What are the advantages of DBMS?

1. Redundancy is controlled.
2. Unauthorised access is restricted.
3. Providing multiple user interfaces.
4. Enforcing integrity constraints.
5. Providing backup and recovery.

5. What are the disadvantage in File Processing System?

1. Data redundancy and inconsistency.
2. Difficult in accessing data.
3. Data isolation.
4. Data integrity.
5. Concurrent access is not possible.
6. Security Problems.

6. Describe the three levels of data abstraction?

The are three levels of abstraction:

1. **Physical level:** The lowest level of abstraction describes how data are stored.
2. **Logical level:** The next higher level of abstraction, describes what data are stored in database and what relationship among those data.
3. **View level:** The highest level of abstraction describes only part of entire database.

7. Define the "integrity rules"?

There are two Integrity rules.

1. **Entity Integrity:** States that "Primary key cannot have NULL value"
2. **Referential Integrity:** States that "Foreign Key can be either a NULL value or should be Primary Key value of other relation."

8. What is Data Independence?

Data independence means that "the application is independent of the storage structure and access strategy of data". In other words, The ability to modify the schema definition in one level should not affect the schema definition in the next higher level.

Two types of Data Independence:

1. **Physical Data Independence:** Modification in physical level should not affect the logical level.
2. **Logical Data Independence:** Modification in logical level should affect the view level.

NOTE: Logical Data Independence is more difficult to achieve

9. What is a view? How it is related to data independence?

A view may be thought of as a virtual table, that is, a table that does not really exist in its own right but is instead derived from one or more underlying base table. In other words, there is no stored file that directly represents the view instead a definition of view is stored in data dictionary. Growth and restructuring of base tables is not reflected in views. Thus the view can insulate users from the effects of restructuring and growth in the database. Hence accounts for logical data independence.

10. What is Data Model?

A collection of conceptual tools for describing data, data relationships data semantics and constraints.

11. What is E-R model?

This data model is based on real world that consists of basic objects called entities and of relationship among these objects. Entities are described in a database by a set of attributes.

12. What is Object Oriented model?

This model is based on collection of objects. An object contains values stored in instance variables within the object. An object also contains bodies of code that operate on the object. These bodies of code are called methods. Objects that contain same types of values and the same methods are grouped together into classes.

13. What is an Entity?

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

14. What is an Entity type?

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

15. What is an Entity set?

It is a collection of all entities of particular entity type in the database.

16. What is an Extension of entity type?

The collections of entities of a particular entity type are grouped together into an entity set.

17. What is Weak Entity set?

An entity set may not have sufficient attributes to form a primary key, and its primary key compromises of its partial key and primary key of its parent entity, then it is said to be Weak Entity set.

18. What is an attribute?

It is a particular property, which describes the entity.

19. What is a Relation Schema and a Relation?

A relation Schema denoted by $R(A_1, A_2, \dots, A_n)$ is made up of the relation name R and the list of attributes A_i that it contains. A relation is defined as a set of tuples. Let r be the relation which contains set tuples $(t_1, t_2, t_3, \dots, t_n)$. Each tuple is an ordered list of n -values $t=(v_1, v_2, \dots, v_n)$.

20. What is degree of a Relation?

It is the number of attribute of its relation schema.

21. What is Relationship?

It is an association among two or more entities.

22. What is Relationship set?

The collection (or set) of similar relationships.

23. What is Relationship type?

Relationship type defines a set of associations or a relationship set among a given set of entity types.

24. What is degree of Relationship type?

It is the number of entity type participating.

25. What is DDL (Data Definition Language)?

A data base schema is specifies by a set of definitions expressed by a special language called DDL.

26. What is DML (Data Manipulation Language)?

This language that enable user to access or manipulate data as organised by appropriate data model.

1. **Procedural DML or Low level:** DML requires a user to specify what data are needed and how to get those data.

2. **Non-Procedural DML or High level:** DML requires a user to specify what data are needed without specifying how to get those data.

27. What is Query evaluation engine?

It executes low-level instruction generated by compiler.

28. What is normalization?

It is a process of analysing the given relation schemas based on their Functional Dependencies (FDs) and primary key to achieve the properties

(1).Minimizing redundancy, (2). Minimizing insertion, deletion and update anomalies.

29. What is Functional Dependency?

A Functional dependency is denoted by $X \rightarrow Y$ between two sets of attributes X and Y that are subsets of R specifies a constraint on the possible tuple that can form a relation state r of R. The constraint is for any two tuples t1 and t2 in r if $t1[X] = t2[X]$ then they have $t1[Y] = t2[Y]$. This means the value of X component of a tuple uniquely determines the value of component Y.

30. What is Lossless join property?

It guarantees that the spurious tuple generation does not occur with respect to relation schemas after decomposition.

31. What is 1 NF (Normal Form)?

The domain of attribute must include only atomic (simple, indivisible) values.

32. What is Fully Functional dependency?

It is based on concept of full functional dependency. A functional dependency $X \rightarrow Y$ is full functional dependency if removal of any attribute A from X means that the dependency does not hold any more.

33. What is 2NF?

A relation schema R is in 2NF if it is in 1NF and every non-prime attribute A in R is fully functionally dependent on primary key.

34. What is 3NF?

A relation schema R is in 3NF if it is in 2NF and for every FD $X \rightarrow A$ either of the following is true

1. X is a Super-key of R.
2. A is a prime attribute of R.

In other words, if every non prime attribute is non-transitively dependent on primary key.

35. What is BCNF (Boyce-Codd Normal Form)?

A relation schema R is in BCNF if it is in 3NF and satisfies an additional constraint that for every FD X A, X must be a candidate key.

36. What is meant by query optimization?

The phase that identifies an efficient execution plan for evaluating a query that has the least estimated cost is referred to as query optimization.

37. What is durability in DBMS?

Once the DBMS informs the user that a transaction has successfully completed, its effects should persist even if the system crashes before all its changes are reflected on disk. This property is called durability.

38. What do you mean by atomicity and aggregation?

1. **Atomicity:** Either all actions are carried out or none are. Users should not have to worry about the effect of incomplete transactions. DBMS ensures this by undoing the actions of incomplete transactions.
2. **Aggregation:** A concept which is used to model a relationship between a collection of entities and relationships. It is used when we need to express a relationship among relationships.

39. What is a query?

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.

40. What do you mean by Correlated subquery?

Subqueries, or nested queries, are used to bring back a set of rows to be used by the parent query. Depending on how the subquery is written, it can be executed once for the parent query or it can be executed once for each row returned by the parent query. If the subquery is executed for each row of the parent, this is called a correlated subquery.

A correlated subquery can be easily identified if it contains any references to the parent subquery columns in its WHERE clause. Columns from the subquery cannot be referenced anywhere else in the parent query. The following example demonstrates a non-correlated subquery.

Example: `SELECT * FROM CUST WHERE '10/03/1990' IN (SELECT ODATE FROM ORDER WHERE CUST.CNUM = ORDER.CNUM)`

41. What are the primitive operations common to all record management systems?

Addition, deletion and modification.

42. What are the unary operations in Relational Algebra?

PROJECTION and SELECTION.

43. Are the resulting relations of PRODUCT and JOIN operation the same?

No.

PRODUCT: Concatenation of every row in one relation with every row in another.

JOIN: Concatenation of rows from one relation and related rows from another.

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

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.

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

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

46. How do you communicate with an RDBMS?

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

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

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.

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

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.

49. What is database Trigger?

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.

50. What is Transaction Manager?

It is a program module, which ensures that 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.

52. What is Authorization and Integrity manager?

It is the program module, which tests for the satisfaction of integrity constraint and checks the authority of user to access data.

53. What are stand-alone procedures?

Procedures that are not part of a package are known as stand-alone because they independently defined. These types of procedures are not available for reference from other Oracle tools. Another limitation of stand-alone procedures is that they are compiled at run time, which slows execution.

54. What are cursors give different types of cursors?

PL/SQL uses cursors for all database information accesses statements. The language supports the use two types of cursors

- 1.) Implicit
- 2.) Explicit

SOME MORE QUESTIONS

1. What are data and information, and how are they related in a database?

Data is recorded facts and figures, and information is knowledge derived from data. A database stores data in such a way that information can be created.

2. What is Enterprise Resource Planning (ERP), and what kind of a database is used in an ERP application?

Enterprise Resource Planning (ERP) is an information system used in manufacturing companies and includes sales, inventory, production planning, purchasing and other business functions. An ERP system typically uses a multiuser database.

3. What is a DBMS?

DBMS stands for Database Management System. A DBMS receives requests from applications and translates those requests into actions on a specific database. A DBMS processes SQL statements or uses other functionality to create, process and administer databases.

4. Why is a database considered to be "self-describing"?

In addition to the users' data, a database contains a description of its own structure. This descriptive data is called "metadata."

5. Who is E.F. Codd, and why is he significant in the development of modern database systems?

While working at IBM, E.F. Codd created the relational database model. A paper he published in 1970 presented his ideas to the world at large. His work is the foundation for most of the DBMSs currently in use, and thus forms the basis for database systems as we know and use them today.

6. What is SQL, and why is it important?

SQL stands for Structured Query Language, and is the most important data processing language in use today. It is not a complete programming language like Java or C#, but a data sublanguage used for creating and processing database data and metadata. All DBMS products today use SQL.

7. Write an SQL SELECT statement to display all the columns of the STUDENT table but only those rows where the Grade column is greater than or equal to 90.

```
SELECT * FROM STUDENT WHERE Grade >= 90;
```

CSE-IV Sem A

8. Name and briefly describe the five SQL built-in functions.

COUNT: computes the number of rows in a table. SUM: totals numeric columns. AVG: computes the average value. MAX: obtains the maximum value of a column in a table. MIN: obtains the minimum value of a column in a table.

9. Write an SQL SELECT statement to count the number of rows in STUDENT table and display the result with the label NumStudents.

```
SELECT COUNT(*) AS NumStudents FROM STUDENT;
```

10. What is an SQL subquery?

An SQL subquery is a means of querying two or more tables at the same time. The subquery itself is an SQL SELECT statement contained within the WHERE clause of another SQL SELECT statement, and separated by being enclosed in parenthesis. Some subqueries have equivalent SQL join structures, but correlated subqueries cannot be duplicated by a join..

11. Discuss the alternative terminology that is used in the relational model.

Relations are also called tables, and sometimes by the older data processing term files. A row is known as a tuple in the relational model, but may also be referred to as a record. Finally, relational model attributes are known as table columns and sometimes as fields.

12. Why are functional dependencies not equations?

Equations deal with numerical relationships. A functional dependency deals with the existence of a determinant relationship between attributes, regardless of whether or not there is a numerical relationship between them. Thus, if we know that there is no hot water every Wednesday, No-Hot-Water is functionally dependent on Wednesday. So, if we know it is Wednesday, then we know we will have No-Hot-Water. This is a functional dependency, but not an equation.

13. What is a foreign key, and what is it used for?

A foreign key is used to establish relationships among relations in the relational model. Technically, a foreign key is a column (or columns) appearing in one relation that is (are) the primary key of another table. Although there may be exceptions, the values in the foreign key columns usually must correspond to values existing in the set of primary key values. This correspondence requirement is created in a database using a referential integrity constraint on the foreign key.

14. What are insertion and deletion anomalies?

A deletion anomaly occurs when, by deleting the facts about one entity, we inadvertently delete facts about another entity; with one deletion, we lose facts about two entities. For example, if we delete the tuple for Student 001289 from a table, we may lose not only the fact that Student 001289 is in Pierce Hall, but also the fact that he has \$200 left in his security deposit. An insertion anomaly happens when we encounter the restriction that we cannot insert a fact about one entity until we have an additional fact about another entity. For example, we want to store the fact that the security deposit

for Pierce Hall is \$300, but we cannot enter this data into the Student relation until a student registers for Pierce Hall.

15. What does it mean when we say that a relation is in Boyce-Codd Normal Form (BCNF)?

A relation is in BCNF when every determinant in the relation is a candidate key. This means that any possible primary key can determine all other attributes in the relation. Attributes may not be determined by non-candidate key attributes or part of a composite candidate key. Thus it is said "I swear to construct my tables so that all nonkey columns are dependent on the key, the whole key and nothing but the key, so help me Codd!"

16. You have been given a set of tables with data and asked to create a new database to store them. When you examine the data values in the tables, what are you looking for?

(1) Multivalued dependencies, (2) Functional dependencies, (3) Candidate keys, (4) Primary keys and (5) Foreign keys.

17. Why do normalized tables require more complex SQL when SQL statements are used in application programs?

Tables that are normalized contain data that has been distributed among the tables, but which may need to be recombined to answer queries from an application. To recombine the data, the programmer will have to use subqueries and/or joins. These SQL structures are more complex to write than a simple SELECT statement.

18. What is the multivalued, multicolumn problem? Include an example not used in the text.

The multivalued, multicolumn problem occurs when a table is designed to include multiple columns that hold variations of one type of attribute data. One example is where boat owners have the names of their boats stored as BOAT_01, BOAT_02 and BOAT_03.

19. Why is the multivalued, multicolumn problem another form of the multivalued dependency problem?

Both problems try to store multiple values on an attribute in a table. In the multivalued, multicolumn problem, the multiple values are stored in different columns. In the multivalued dependency problem the multiple values are stored in different rows. In both cases, the solution is the same: store the multiple values in a separate table.

20. What is the inconsistent values problem? Include an example not used in the text.

The inconsistent values problem occurs when different users or data sources use slightly different forms of the same data value. One example is where automobiles are specified as "Ford, 2-door, Red" in one cell and "Red Ford 2-door" in another.

21. Explain the relationship between entity, entity class, and entity instance.

An entity is something that can be identified in the users' work environment, something that the users want to track. Entities of a given type are grouped into entity classes. An entity instance is the representation of a particular entity.

22. Explain the difference between attributes and identifiers.

Entities have attributes. Attributes are properties that describe the entity's characteristics. Entity

instances have identifiers. Identifiers are attributes that name, or identify, entity instances.

23. Name and describe three types of binary relationships.

1:1 - a single entity instance of one type is related to a single-entity instance of another type.

1:N - a single entity instance of one type is related to many-entity instances of another type.

M:N - many-entity instances of one type relate to many-entity instances of another type.

25. What is a recursive relationship? Give an example not used in the text.

A recursive relationship is a relationship between an entity and itself. For example, given the entity PERSON, a recursive relationship could be used to show a PERSON and his or her SIBLINGS (brothers and sisters).

26. What are the steps for transforming an entity into a table?

The steps are: (1) specify the primary key, (2) specify candidate keys, (3) specify column properties including null status, data type, default value (if any), and data constraints (if any), and (4) verifying normalization.

31. What is a SQL view? Briefly explain the use of views.

A SQL view is a virtual table built from other tables or views. Views are used to (1) hide columns or rows, (2) show the results of computed columns, (3) hide complicated SQL syntax, (4) layer built-in functions, (5) provide a level of indirection between application programs and tables, (6) assign different sets of processing permissions to tables, and (7) to assign different sets of triggers to the same table.

32. Explain the "paradigm mismatch" between SQL and application programming languages.

SQL statements return a set of rows, while an application program works on one row at a time. To resolve this mismatch the results of SQL statements are processed as pseudofiles, using a cursor or pointer to specify which row is being processed.

33. Name four applications for triggers.

(1) providing default values, (2) enforcing data constraints, (3) updating views and (4) enforcing referential integrity

34. What are stored procedures, and how do they differ from triggers?

A stored procedure is a program that is stored within the database and is compiled when used. They can receive input parameters and they can return results. Unlike triggers, their scope is database-wide; they can

be used by any process that has permission to use the database stored procedure.

35. What are the advantages of using stored procedures?

The advantages of stored procedures are (1) greater security, (2) decreased network traffic, (3) the fact that SQL can be optimized and (4) code sharing which leads to less work, standardized processing, and specialization among developers.

36. Why is database redesign necessary?

Database redesign is necessary for two reasons. First, redesign is necessary both to fix mistakes made during the initial database design. Second, redesign is necessary to adapt the database to changes in system requirements. Such changes are common because information systems and organizations do not just influence each other they create each other. Thus, new information systems cause changes in systems requirements.

37. What is the difference between a correlated subquery and a regular subquery?

A correlated subquery appears deceptively similar to a regular subquery. The difference is that a regular subquery can be processed from the bottom up. In a regular subquery, results from the lowest query can be determined and used to evaluate the upper-level query. In contrast, in a correlated subquery, the processing is nested; that is, a row from an upper query statement is used in comparison with rows in a lower level query. The key distinction of a correlated subquery is that the lower-level select statements use columns from upper-level statements.

39. Explain how to add a NOT NULL column to a table.

First, add the column as NULL. Then use UPDATE to add data to every row. Finally use an ALTER TABLE . . . ALTER COLUMN statement to change the column constraint to NOT NULL.

41. Explain the difference between an exclusive lock and a shared lock.

An exclusive lock prohibits other users from reading the locked resource; a shared lock allows other users to read the locked resource, but they cannot update it.

Optimistic locking assumes no transaction conflict will occur and deals with the consequences if it does.

Pessimistic locking assumes that conflict will occur and so prevents it ahead of time with locks. In general, optimistic locking is preferred for the Internet and for many intranet applications.

43. What is deadlock? How can it be avoided? How can it be resolved once it occurs?

Deadlock occurs when two transactions are each waiting on a resource that the other transaction holds. Deadlock can be prevented by requiring transactions to acquire all locks at the same time; once it occurs, the only way to cure it is to abort one of the transactions and back out of partially completed work.

44. What are the major functions of the database administrator?

Managing database structure, controlling concurrent processing, managing processing rights and responsibilities, developing database security, providing for database recovery, managing the DBMS and maintaining the data repository.

45. Explain what we mean by an ACID transaction.

An ACID transaction is one that is atomic, consistent, isolated, and durable. Durable means that database changes are permanent. Consistency can mean either statement level or transaction level consistency. With transaction level consistency, a transaction may not see its own changes. There are four transaction isolation levels: read committed, read uncommitted, repeatable read and serialization. Atomic means it is performed as a unit.

46. What are the ways in which an Oracle database can be created?

There are three (3) ways to create an Oracle database. You can create a database using the Database Configuration Assistant, using the Oracle-supplied database creation procedures or using the SQL CREATE DATABASE command.

47. What are sequences, and what are the possible problems when using them

A sequence is an object that generates a sequential series of unique numbers. Sequences are most often used to provide values for surrogate keys. However, there are three problems with using sequences. First, a developer can use a defined sequence for any purpose; and if a sequence is used for purposes other than the surrogate key, some values will be missing. A second problem is that there is nothing in the schema that prevents someone from issuing an INSERT statement that adds data to the table but that does not use the sequence. Finally, it is possible to use the wrong sequence when putting data into a table.

48. Under what conditions should indexes be used?

Indexes can be created to enforce uniqueness, to facilitate sorting, and to enable fast retrieval by column values. A good candidate for an index is a column that is frequently used with equal conditions in WHERE clauses.

50. What are the types of files used in Oracle recovery?

Datafiles, control files and two types of ReDo log files: OnLine Redo and Offline ReDo (which is also known as Archive ReDo).