

RDBMS Job Interview Questions And Answers



Interview Questions Answers

<https://interviewquestionsanswers.org/>

About Interview Questions Answers

Interview Questions Answers . ORG is an interview preparation guide of thousands of Job Interview Questions And Answers, Job Interviews are always stressful even for job seekers who have gone on countless interviews. The best way to reduce the stress is to be prepared for your job interview. Take the time to review the standard interview questions you will most likely be asked. These interview questions and answers on RDBMS will help you strengthen your technical skills, prepare for the interviews and quickly revise the concepts.

If you find any **question or answer** is incorrect or incomplete then you can **submit your question or answer** directly with out any registration or login at our website. You just need to visit [RDBMS Interview Questions And Answers](#) to add your answer click on the *Submit Your Answer* links on the website; with each question to post your answer, if you want to ask any question then you will have a link *Submit Your Question*; that's will add your question in RDBMS category. To ensure quality, each submission is checked by our team, before it becomes live. This [RDBMS Interview preparation PDF](#) was generated at **Wednesday 29th November, 2023**

You can follow us on FaceBook for latest Jobs, Updates and other interviews material.
www.facebook.com/InterviewQuestionsAnswers.Org

Follow us on Twitter for latest Jobs and interview preparation guides.
<https://twitter.com/InterviewQA>

If you need any further assistance or have queries regarding this document or its material or any of other inquiry, please do not hesitate to contact us.

Best Of Luck.

Interview Questions Answers.ORG Team
<https://InterviewQuestionsAnswers.ORG/Support@InterviewQuestionsAnswers.ORG>



RDBMS Interview Questions And Answers Guide.

Question - 1:

Difference between controlled and uncontrolled redundancy?

Ans:

No Answer is Posted For this Question

Be First To [Post Your Answer Now.](#)

Question - 2:

What is indexing and what are the different kinds of indexing?

Ans:

Indexing is a technique for determining how quickly specific data can be found.

Types:

- * Binary search style indexing
- * B-Tree indexing
- * Inverted list indexing
- * Memory resident table
- * Table indexing

[View All Answers](#)

Question - 3:

What is system catalog or catalog relation? How is better known as?

Ans:

A RDBMS maintains a description of all the data that it contains, information about every relation and index that it contains. This information is stored in a collection of relations maintained by the system called metadata. It is also called data dictionary.

[View All Answers](#)

Question - 4:

What is meant by query optimization?

Ans:

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

[View All Answers](#)

Question - 5:

What is join dependency and inclusion dependency?

Ans:

Join Dependency:

A Join dependency is generalization of Multivalued dependency. A JD $\{R_1, R_2, \dots, R_n\}$ is said to hold over a relation R if $R_1, R_2, R_3, \dots, R_n$ is a lossless-join decomposition of R. There is no set of sound and complete inference rules for JD.

Inclusion Dependency:

An Inclusion Dependency is a statement of the form that some columns of a relation are contained in other columns. A foreign key constraint is an example of inclusion dependency.

[View All Answers](#)

Question - 6:

What is durability in DBMS?

Ans:

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.

[View All Answers](#)

Question - 7:

What do you mean by atomicity and aggregation?

Ans:

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.

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.

[View All Answers](#)

Question - 8:

What is a Phantom Deadlock?

Ans:

In distributed deadlock detection, the delay in propagating local information might cause the deadlock detection algorithms to identify deadlocks that do not really exist. Such situations are called phantom deadlocks and they lead to unnecessary aborts.

[View All Answers](#)

Question - 9:

What is a checkpoint and When does it occur?

Ans:

A Checkpoint is like a snapshot of the DBMS state. By taking checkpoints, the DBMS can reduce the amount of work to be done during restart in the event of subsequent crashes.

[View All Answers](#)

Question - 10:

What are the different phases of transaction?

Ans:

Different phases are

- * Analysis phase
- * Redo Phase
- * Undo phase

[View All Answers](#)

Question - 11:

What do you mean by flat file database?

Ans:

It is a database in which there are no programs or user access languages. It has no cross-file capabilities but is user-friendly and provides user-interface management.

[View All Answers](#)

Question - 12:

What is transparent DBMS?

Ans:

It is one, which keeps its Physical Structure hidden from user.

[View All Answers](#)

Question - 13:

Brief theory of Network, Hierarchical schemas and their properties?

Ans:

Network schema uses a graph data structure to organize records example for such a database management system is CTCTG while a hierarchical schema uses a tree data structure example for such a system is IMS.

[View All Answers](#)

Question - 14:

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.

[View All Answers](#)

**Question - 15:**

What do you mean by Correlated subquery?

Ans:

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.

E.g. Select * From CUST Where '10/03/1990' IN (Select ODATE From ORDER Where CUST.CNUM = ORDER.CNUM)

[View All Answers](#)

Question - 16:

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

Ans:

Addition, deletion and modification.

[View All Answers](#)

Question - 17:

Name the buffer in which all the commands that are typed in are stored?

Ans:

Edit Buffer

[View All Answers](#)

Question - 18:

What are the unary operations in Relational Algebra?

Ans:

PROJECTION and SELECTION.

[View All Answers](#)

Question - 19:

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

Ans:

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.

[View All Answers](#)

Question - 20:

What is RDBMS KERNEL?

Ans:

Two important pieces of RDBMS architecture are the kernel, which is the software, and the data dictionary, which consists of the system-level data structures used by the kernel to manage the database.

You might think of an RDBMS as an operating system (or set of subsystems), designed specifically for controlling data access; its primary functions are storing, retrieving, and securing data. An RDBMS maintains its own list of authorized users and their associated privileges; manages memory caches and paging; controls locking for concurrent resource usage; dispatches and schedules user requests; and manages space usage within its table-space structures.

[View All Answers](#)

Question - 21:

Name the sub-systems of a RDBMS?

Ans:

I/O, Security, Language Processing, Process Control, Storage Management, Logging and Recovery, Distribution Control, Transaction Control, Memory Management, Lock Management

[View All Answers](#)

Question - 22:

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.

[View All Answers](#)

Question - 23:

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.

[View All Answers](#)

Question - 24:

How do you communicate with an RDBMS?

Ans:

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

[View All Answers](#)

Question - 25:

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.

[View All Answers](#)

Question - 26:

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

- * Database files
- * Control files
- * 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.

[View All Answers](#)

Question - 27:

What are the four Oracle system processes that must always be up and running for the database to be useable?

Ans:

The four Oracle system processes that must always be up and running for the database to be useable include DBWR (Database Writer), LGWR (Log Writer), SMON (System Monitor), and PMON (Process Monitor).

[View All Answers](#)

Question - 28:

What are database files, control files and log files. How many of these files should a database have at least? Why?

Ans:

Database Files

The database files hold the actual data and are typically the largest in size. Depending on their sizes, the tables (and other objects) for all the user accounts can go in one database file-but that's not an ideal situation because it does not make the database structure very flexible for controlling access to storage for different users, putting the database on different disk drives, or backing up and restoring just part of the database.

You must have at least one database file but usually, more than one files are used. In terms of accessing and using the data in the tables and other objects, the number (or location) of the files is immaterial.

The database files are fixed in size and never grow bigger than the size at which they were created

Control Files

The control files and redo logs support the rest of the architecture. Any database must have at least one control file, although you typically have more than one to guard against loss. The control file records the name of the database, the date and time it was created, the location of the database and redo logs, and the synchronization information to ensure that all three sets of files are always in step. Every time you add a new database or redo log file to the database, the information is recorded in the control files.

Redo Logs

Any database must have at least two redo logs. These are the journals for the database; the redo logs record all changes to the user objects or system objects. If any type of failure occurs, the changes recorded in the redo logs can be used to bring the database to a consistent state without losing any committed transactions. In the case of non-data loss failure, Oracle can apply the information in the redo logs automatically without intervention from the DBA.

The redo log files are fixed in size and never grow dynamically from the size at which they were created.

[View All Answers](#)

Question - 29:

What is ROWID?

Ans:

The ROWID is a unique database-wide physical address for every row on every table. Once assigned (when the row is first inserted into the database), it never changes until the row is deleted or the table is dropped.

The ROWID consists of the following three components, the combination of which uniquely identifies the physical storage location of the row.

- * Oracle database file number, which contains the block with the rows
- * Oracle block address, which contains the row
- * The row within the block (because each block can hold many rows)

The ROWID is used internally in indexes as a quick means of retrieving rows with a particular key value. Application developers also use it in SQL statements as a quick way to access a row once they know the ROWID



[View All Answers](#)

Question - 30:

What is Oracle Block? Can two Oracle Blocks have the same address?

Ans:

Oracle "formats" the database files into a number of Oracle blocks when they are first created-making it easier for the RDBMS software to manage the files and easier to read data into the memory areas.

The block size should be a multiple of the operating system block size. Regardless of the block size, the entire block is not available for holding data; Oracle takes up some space to manage the contents of the block. This block header has a minimum size, but it can grow.

These Oracle blocks are the smallest unit of storage. Increasing the Oracle block size can improve performance, but it should be done only when the database is first created.

Each Oracle block is numbered sequentially for each database file starting at 1. Two blocks can have the same block address if they are in different database files.

[View All Answers](#)

Question - 31:

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.

[View All Answers](#)

Question - 32:

Name two utilities that Oracle provides, which are used for backup and recovery?

Ans:

Along with the RDBMS software, Oracle provides two utilities that you can use to back up and restore the database. These utilities are Export and Import.

The Export utility dumps the definitions and data for the specified part of the database to an operating system binary file. The Import utility reads the file produced by an export, recreates the definitions of objects, and inserts the data.

If Export and Import are used as a means of backing up and recovering the database, all the changes made to the database cannot be recovered since the export was performed. The best you can do is recover the database to the time when the export was last performed.

[View All Answers](#)

Question - 33:

What are stored-procedures? And what are the advantages of using them?

Ans:

Stored procedures are database objects that perform a user-defined operation. A stored procedure can have a set of compound SQL statements. A stored procedure executes the SQL commands and returns the result to the client. Stored procedures are used to reduce network traffic.

[View All Answers](#)

Question - 34:

Does PL/SQL support overloading? Explain?

Ans:

The concept of overloading in PL/SQL relates to the idea that you can define procedures and functions with the same name. PL/SQL does not look only at the referenced name, however, to resolve a procedure or function call. The count and data types of formal parameters are also considered.

PL/SQL also attempts to resolve any procedure or function calls in locally defined packages before looking at globally defined packages or internal functions. To further ensure calling the proper procedure, you can use the dot notation. Prefacing a procedure or function name with the package name fully qualifies any procedure or function reference.

[View All Answers](#)

Question - 35:

Tables derived from the ERD?

Ans:

- a) Are totally unnormalised
- b) Are always in 1NF
- c) Can be further denormalised
- d) May have multi-valued attributes
- (b) Are always in 1NF

[View All Answers](#)

Question - 36:

Spurious tuples may occur due to?

Ans:

- i. Bad normalization
- ii. Theta joins
- iii. Updating tables from join
- a) i & ii b) ii & iii



c) i & iii d) ii & iii

(a) i & iii because theta joins are joins made on keys that are not primary keys.

[View All Answers](#)

Question - 37:

A B C is a set of attributes. The functional dependency is as follows

AB → B

AC → C

C → B?

Ans:

a) is in 1NF

b) is in 2NF

c) is in 3NF

d) is in BCNF

(a) is in 1NF since (AC)⁺ = { A, B, C } hence AC is the primary key. Since C B is a FD given, where neither C is a Key nor B is a prime attribute, this it is not in 3NF. Further B is not functionally dependent on key AC thus it is not in 2NF. Thus the given FDs is in 1NF.

[View All Answers](#)

Question - 38:

In mapping of ERD to DFD?

Ans:

a) entities in ERD should correspond to an existing entity/store in DFD

b) entity in DFD is converted to attributes of an entity in ERD

c) relations in ERD has 1 to 1 correspondence to processes in DFD

d) relationships in ERD has 1 to 1 correspondence to flows in DFD

(a) entities in ERD should correspond to an existing entity/store in DFD

[View All Answers](#)

Question - 39:

A dominant entity is the entity?

Ans:

a) on the N side in a 1 : N relationship

b) on the 1 side in a 1 : N relationship

c) on either side in a 1 : 1 relationship

d) nothing to do with 1 : 1 or 1 : N relationship

(b) on the 1 side in a 1 : N relationship

[View All Answers](#)

Question - 40:

What is Storage Manager?

Ans:

It is a program module that provides the interface between the low-level data stored in database, application programs and queries submitted to the system.

[View All Answers](#)

Question - 41:

What is Buffer Manager?

Ans:

It is a program module, which is responsible for fetching data from disk storage into main memory and deciding what data to be cache in memory.

[View All Answers](#)

Question - 42:

What is File Manager?

Ans:

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.

[View All Answers](#)

Question - 43:

What is Authorization and Integrity manager?

Ans:

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

[View All Answers](#)

Question - 44:

What are stand-alone procedures?



Ans:

Procedures that are not part of a package are known as stand-alone because they independently defined. A good example of a stand-alone procedure is one written in a SQL*Forms application. 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.

[View All Answers](#)

Question - 45:

What are cursors give different types of cursors?

Ans:

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

- * Implicit
- * Explicit

[View All Answers](#)

Question - 46:

What is cold backup and hot backup (in case of Oracle)?

Ans:

* Cold Backup:

It is copying the three sets of files (database files, redo logs, and control file) when the instance is shut down. This is a straight file copy, usually from the disk directly to tape. You must shut down the instance to guarantee a consistent copy.

If a cold backup is performed, the only option available in the event of data file loss is restoring all the files from the latest backup. All work performed on the database since the last backup is lost.

* Hot Backup:

Some sites (such as worldwide airline reservations systems) cannot shut down the database while making a backup copy of the files. The cold backup is not an available option.

So different means of backing up database must be used - the hot backup. Issue a SQL command to indicate to Oracle, on a tablespace-by-tablespace basis, that the files of the tablespace are to be backed up. The users can continue to make full use of the files, including making changes to the data. Once the user has indicated that he/she wants to back up the tablespace files, he/she can use the operating system to copy those files to the desired backup destination.

The database must be running in ARCHIVELOG mode for the hot backup option.

If a data loss failure does occur, the lost database files can be restored using the hot backup and the online and offline redo logs created since the backup was done. The database is restored to the most consistent state without any loss of committed transactions.

[View All Answers](#)

Question - 47:

What are Armstrong rules? How do we say that they are complete and/or sound?

Ans:

The well-known inference rules for FDs

? Reflexive rule :

If Y is subset or equal to X then X Y.

* Augmentation rule:

If X Y then XZ YZ.

* Transitive rule:

If {X Y, Y Z} then X Z.

* Decomposition rule :

If X YZ then X Y.

* Union or Additive rule:

If {X Y, X Z} then X YZ.

* Pseudo Transitive rule :

If {X Y, WY Z} then WX Z.

Of these the first three are known as Armstrong Rules. They are sound because it is enough if a set of FDs satisfy these three. They are called complete because using these three rules we can generate the rest all inference rules.

[View All Answers](#)

Question - 48:

How can you find the minimal key of relational schema?

Ans:

Minimal key is one which can identify each tuple of the given relation schema uniquely. For finding the minimal key it is required to find the closure that is the set of all attributes that are dependent on any given set of attributes under the given set of functional dependency.

Algo. I Determining X^+ , closure for X, given set of FDs F

1. Set $X^+ = X$

2. Set Old $X^+ = X^+$

3. For each FD Y Z in F and if Y belongs to X^+ then add Z to X^+

4. Repeat steps 2 and 3 until Old $X^+ = X^+$

Algo. II Determining minimal K for relation schema R, given set of FDs F

1. Set K to R that is make K a set of all attributes in R

2. For each attribute A in K

a. Compute $(K - A)^+$ with respect to F

b. If $(K - A)^+ = R$ then set $K = (K - A)^+$

[View All Answers](#)

Question - 49:

What do you understand by dependency preservation?



Ans:

Given a relation R and a set of FDs F, dependency preservation states that the closure of the union of the projection of F on each decomposed relation R_i is equal to the closure of F, i.e.,

$$((\pi_{R_1}(F)) \cup \dots \cup (\pi_{R_n}(F)))^+ = F^+$$

if decomposition is not dependency preserving, then some dependency is lost in the decomposition.

[View All Answers](#)

Question - 50:

What are the different types of JOIN operations?

Ans:

Equi Join: This is the most common type of join which involves only equality comparisons. The disadvantage in this type of join is that there

[View All Answers](#)

Question - 51:

What is meant by Proactive, Retroactive and Simultaneous Update?

Ans:

Proactive Update:

The updates that are applied to database before it becomes effective in real world .

Retroactive Update:

The updates that are applied to database after it becomes effective in real world .

Simultaneous Update:

The updates that are applied to database at the same time when it becomes effective in real world .

[View All Answers](#)

Question - 52:

What is 5NF?

Ans:

A Relation schema R is said to be 5NF if for every join dependency $\{R_1, R_2, \dots, R_n\}$ that holds R, one the following is true

* $R_i = R$ for some i.

* The join dependency is implied by the set of FD, over R in which the left side is key of R.

[View All Answers](#)

Question - 53:

What is 4NF?

Ans:

A relation schema R is said to be in 4NF if for every Multivalued dependency $X \twoheadrightarrow Y$ that holds over R, one of following is true

* X is subset or equal to (or) $XY = R$.

* X is a super key.

[View All Answers](#)

Question - 54:

What is BCNF (Boyce-Codd Normal Form)?

Ans:

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

[View All Answers](#)

Question - 55:

What is 3NF?

Ans:

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

* X is a Super-key of R.

* A is a prime attribute of R.

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

[View All Answers](#)

Question - 56:

What is 2NF?

Ans:

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.

[View All Answers](#)

Question - 57:

What is Fully Functional dependency?

Ans:



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

[View All Answers](#)

Question - 58:

What is 1 NF (Normal Form)?

Ans:

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

[View All Answers](#)

Question - 59:

What is Lossless join property?

Ans:

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

[View All Answers](#)

Question - 60:

What is Multivalued dependency?

Ans:

Multivalued dependency denoted by $X \twoheadrightarrow Y$ specified on relation schema R, where X and Y are both subsets of R, specifies the following constraint on any relation r of R: if two tuples t1 and t2 exist in r such that $t1[X] = t2[X]$ then t3 and t4 should also exist in r with the following properties

- * $t3[X] = t4[X] = t1[X] = t2[X]$
 - * $t3[Y] = t1[Y]$ and $t4[Y] = t2[Y]$
 - * $t3[Z] = t2[Z]$ and $t4[Z] = t1[Z]$
- where $[Z] = (R - (X \cup Y))$

[View All Answers](#)

Question - 61:

When is a functional dependency F said to be minimal?

Ans:

- * Every dependency in F has a single attribute for its right hand side.
- * It cannot replace any dependency $X \twoheadrightarrow A$ in F with a dependency $Y \twoheadrightarrow A$ where Y is a proper subset of X and still have a set of dependency that is equivalent to F.
- * We cannot remove any dependency from F and still have set of dependency that is equivalent to F.

[View All Answers](#)

Question - 62:

What is Functional Dependency?

Ans:

Functional dependency is denoted by $X \twoheadrightarrow Y$ 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.

[View All Answers](#)

Question - 63:

What is normalization?

Ans:

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

- * Minimizing redundancy
- * Minimizing insertion, deletion and update anomalies.

[View All Answers](#)

Question - 64:

How does Tuple-oriented relational calculus differ from domain-oriented relational calculus?

Ans:

The tuple-oriented calculus uses a tuple variables i.e., variable whose only permitted values are tuples of that relation. E.g. QUEL

The domain-oriented calculus has domain variables i.e., variables that range over the underlying domains instead of over relation. E.g. ILL, DEDUCE.

[View All Answers](#)

Question - 65:

What is Relational Calculus?

Ans:

It is an applied predicate calculus specifically tailored for relational databases proposed by E.F. Codd. E.g. of languages based on it are DSL ALPHA, QUEL.

[View All Answers](#)

**Question - 66:**

What is Relational Algebra?

Ans:

It is procedural query language. It consists of a set of operations that take one or two relations as input and produce a new relation.

[View All Answers](#)

Question - 67:

What is Set-at-a-time or Set-oriented?

Ans:

The High level or Non-procedural DML can specify and retrieve many records in a single DML statement. This retrieve of a record is said to be Set-at-a-time or Set-oriented.

[View All Answers](#)

Question - 68:

What is Record-at-a-time?

Ans:

The Low level or Procedural DML can specify and retrieve each record from a set of records. This retrieve of a record is said to be Record-at-a-time.

[View All Answers](#)

Question - 69:

What is DDL Interpreter?

Ans:

It interprets DDL statements and record them in tables containing metadata.

[View All Answers](#)

Question - 70:

What is Query evaluation engine?

Ans:

It executes low-level instruction generated by compiler.

[View All Answers](#)

Question - 71:

What is DML Compiler?

Ans:

It translates DML statements in a query language into low-level instruction that the query evaluation engine can understand.

[View All Answers](#)

Question - 72:

What is DML (Data Manipulation Language)?

Ans:

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

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

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

[View All Answers](#)

Question - 73:

What is Data Storage - Definition Language?

Ans:

The storage structures and access methods used by database system are specified by a set of definition in a special type of DDL called data storage-definition language.

[View All Answers](#)

Question - 74:

What is SDL (Storage Definition Language)?

Ans:

This language is to specify the internal schema. This language may specify the mapping between two schemas.

[View All Answers](#)

Question - 75:

What is VDL (View Definition Language)?

Ans:



It specifies user views and their mappings to the conceptual schema.

[View All Answers](#)

Question - 76:

What is DDL (Data Definition Language)?

Ans:

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

[View All Answers](#)

Question - 77:

What is Relationship?

Ans:

It is an association among two or more entities.

Relationship Set - The collection (or set) of similar relationships.

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

Degree of Relationship Type - It is the number of entity type participating.

[View All Answers](#)

Question - 78:

What is degree of a Relation?

Ans:

It is the number of attribute of its relation schema.

[View All Answers](#)

Question - 79:

What is a Relation Schema and a Relation?

Ans:

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 $t=(v_1, v_2, \dots, v_n)$.

[View All Answers](#)

Question - 80:

What is an attribute?

Ans:

It is a particular property, which describes the entity.

[View All Answers](#)

Question - 81:

What is Weak Entity set?

Ans:

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.

[View All Answers](#)

Question - 82:

What is an Extension of entity type?

Ans:

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

[View All Answers](#)

Question - 83:

What is an Entity set?

Ans:

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

[View All Answers](#)

Question - 84:

What is an Entity type?

Ans:

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

[View All Answers](#)

**Question - 85:**

What is Object Oriented model?

Ans:

This model is based on collection of objects. An object contains values stored in instance variables with in 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.

[View All Answers](#)

Question - 86:

What is E-R model?

Ans:

E-R model stands for Entity-Relationship 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.

[View All Answers](#)

Question - 87:

What is Data Model?

Ans:

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

[View All Answers](#)

Question - 88:

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

Ans:

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 direct 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.

[View All Answers](#)

Question - 89:

What is Data Independence?

Ans:

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 are:

* Physical Data Independence: Modification in physical level should not affect the logical level.

* Logical Data Independence: Modification in logical level should affect the view level.

NOTE: Logical Data Independence is more difficult to achieve.

[View All Answers](#)

Question - 90:

How is the data structure of System R different from the relational structure?

Ans:

Unlike Relational systems in System R

* Domains are not supported

* Enforcement of candidate key uniqueness is optional

* Enforcement of entity integrity is optional

* Referential integrity is not enforced

[View All Answers](#)

Question - 91:

What is System R? What are its two major subsystems?

Ans:

System R was designed and developed over a period of 1974-79 at IBM San Jose Research Center. It is a prototype and its purpose was to demonstrate that it is possible to build a Relational System that can be used in a real life environment to solve real life problems, with performance at least comparable to that of existing system.

Its two subsystems are

* Research Storage

* System Relational Data System.

[View All Answers](#)

Question - 92:

What is extension and intension?

Ans:

Extension - It is the number of tuples present in a table at any instance. This is time dependent.



Intension - It is a constant value that gives the name, structure of table and the constraints laid on it.

[View All Answers](#)

Question - 93:

Define the Integrity Rules?

Ans:

There are two Integrity rules.

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

[View All Answers](#)

Question - 94:

Describe the three levels of data abstraction?

Ans:

The are three levels of abstraction:

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

[View All Answers](#)

Question - 95:

Disadvantage in File Processing System?

Ans:

- * Data redundancy & inconsistency.
- * Difficult in accessing data.
- * Data isolation.
- * Data integrity.
- * Concurrent access is not possible.
- * Security Problems.

[View All Answers](#)

Question - 96:

Advantages of DBMS?

Ans:

Redundancy is controlled. Unauthorised access is restricted. Providing multiple user interfaces. Enforcing integrity constraints. Providing backup and recovery.

[View All Answers](#)

Question - 97:

What is a Database system?

Ans:

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

[View All Answers](#)

Question - 98:

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.

[View All Answers](#)

Question - 99:

What is a Database?

Ans:

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.

[View All Answers](#)

Databases Programming Most Popular & Related Interview Guides

- 1 : [SQL Interview Questions and Answers.](#)
- 2 : [SSRS Interview Questions and Answers.](#)
- 3 : [Sybase Interview Questions and Answers.](#)
- 4 : [Database Administrator \(DBA\) Interview Questions and Answers.](#)
- 5 : [Teradata Interview Questions and Answers.](#)
- 6 : [PostgreSQL Interview Questions and Answers.](#)
- 7 : [MongoDB Interview Questions and Answers.](#)
- 8 : [MS SQL Server Interview Questions and Answers.](#)
- 9 : [IBM DB2 Interview Questions and Answers.](#)
- 10 : [Data Modeling Interview Questions and Answers.](#)

Follow us on FaceBook

www.facebook.com/InterviewQuestionsAnswers.Org

Follow us on Twitter

<https://twitter.com/InterviewQA>

For any inquiry please do not hesitate to contact us.

Interview Questions Answers.ORG Team

[https://InterviewQuestionsAnswers.ORG/
support@InterviewQuestionsAnswers.ORG](https://InterviewQuestionsAnswers.ORG/support@InterviewQuestionsAnswers.ORG)