
Guru99 Provides [FREE ONLINE TUTORIAL](#) on Various courses like

[Java](#) | [MIS](#) | [MongoDB](#) | [BigData](#) | [Cassandra](#) | [Web Services](#)

[SQLite](#) | [JSP](#) | [Informatica](#) | [Accounting](#) | [SAP Training](#) | [Python](#)

[Excel](#) | [ASP Net](#) | [HBase](#) | [Testing](#) | [Selenium](#) | [CCNA](#) | [NodeJS](#)

[TensorFlow](#) | [Data Warehouse](#) | [R Programming](#) | [Live Projects](#) | [DevOps](#)

Top 50 DB2 Interview Questions & Answers

1.) Define DB2.

DB2 is a Database Management System for the MVS Operating System where, DB2 is a subsystem of MVS Operating System.

2.) What is the purpose of using COMMIT?

The data changes can be made permanent by using COMMIT. It also permits data to be accessed by other applications who can reference the committed data

3.) List out the Data types available.

The Data types available here are:
[crayon-5f2d021d6651d661570216/]

4.) What are the uses of DB2 Optimizer?

It processes SQL statements.

It helps to select the access path

5.) Define SQLCA.

SQL Communication Area is a structure of variables, which are updated after every execution of SQL statements. For an application that contains executable SQL statements, only one SQLCA is needed. FORTRAN need more than one SQLCA to be provided .For Java, SQLCA is not applicable.

**6.) Define CHECK constraint.**

It is specified as a condition or criteria to ensure data integrity. A value to be inserted or updated to a table is tested by CHECK constraint. The CHECK constraints are created during the creation of the table.

7.) What is SQLCA's maximum length?

136 is the maximum length of the SQLCA.

8.) Discuss about DB2 bind?

The process that builds access paths to the DB2 table is known as bind. The bind uses Database Request Modules from the DB2 pre-compile step as input and produces an application plan. It checks for user authentication and validates SQL statements in the DBRM(s).

9.) List out the three types of page locks that can be held.

Exclusive
Update
Share.

10.) Define buffer pool.

The buffer pool is a reserved main storage, which is to satisfy the buffering requirements for one or more table spaces or indexes. It can be made up of either 4K or 32K pages.

11.) Explain the function of Data Manager.

The physical database is managed by the DB2 component called Data manager. It invokes other system components to perform logging, locking etc.

12.) What is a storage group (STOGROUP)?

STOGROUP is a named collection of DASD volumes, which is to be used by index spaces and table spaces of the database.

13.) Define predicate?

Predicate is an element of a search condition. It expresses or implies a search condition.

14.) Define Declaration Generator (DCLGEN).

Declaration Generator is a facility that is used to form SQL statements that describes a table or view. At pre-compile time, the table or view descriptions are then used to check the validity of SQL statements.

15.) List out the buffer pools in DB2

There are four buffer pools in DB2 and they are:

BP0
BP1
BP2
BP32

16.) Define clustering index.

Clustering index is a type of index, which locates the table rows and determines how to group the rows together in the tablespace.

17.) What is concurrency?

More than one DB2 application process can access the same data at the same time, is known as concurrency. However, problems can happen such as, lost updates access to unrepeatable reads and uncommitted data.

18.) Explain the Function done by data manager?

Data manager can be considered as a component that is capable of managing the databases that are physically present and is capable of invoking other components associated with the system for performing functionalities like logging, locking and in performing other I/O operations.

19.) Explain about DBRM.

DBRM stands for Database Request Module and is a component inside DB2, which is created by the pre compiler of DB2. This is a module that consists of SQL source statements that get

extracted out of the application program. DBRMs form inputs that are helpful in the binding process.

20.) Define Data page.

Data page can be considered as a unit that is capable of retrieving data from the database. The database from which the data can be retrieved is in the form of 4 kilobytes or 32 kilobytes. The form in which data is retrieved depends on the way the table is defined inside the database. Data page also contains information regarding the catalog or user that are part of the database.

21.) Explain about RCT.

RCT is expanded as Resource – Control Table and is defined in the DB2/CICS region. This is the component that comprises of features that are gathered through macros of DSNCRCT. RCT matches with the transaction ID that of CICS, with the authorization ID that of DB2. This should also be matched with plan ID.

22.) How can tablespace be moved to another DASD volume that is allocated for that tablespace?

Tablespace that you are using is allocated only to STOGROUP, then you can enter the command ALTER STOGROUP for adding as well as deleting volume. REORG TABLESPACE and RECOVER TABLESPACE are statements that are helpful in creating new STOGROUP that can point towards the new volume. ALTER tablespace and REORG and RECOVER are statements used for altering and recovering the tablespace allocated in the memory.

23.) What is the information associated with SYSIBM.SYSLINKS table?

This is the table that contains information on the links that exists between the tables created through referential constraints.

24.) Explain in detail about buffer manager and its functionalities?

Buffer manager can be considered as the component inside DB2 that helps in transferring data between virtual as well as external medium. The buffer manager reduces the quantity of physical input as well as output operations that are actually performed by making use of buffering techniques that are highly sophisticated.

25.) Explain about cursor stability?

Cursor stability is the property that tells the DB2 that the values of database that are read by making use of this application gets protected while the data is used.

26.) Mention a credible reason why SELECT* is never given preference in an SQL program that has been embedded.

There are primarily three reasons why SELECT* is never given preference in an embedded SQL program. These are: -

- In case if there is an alteration in the structure of the table, then the program has to undergo a modification process.
- All columns will be retrieved by the program including those columns , which might not be used.
- In case if the user wishes to scan the index, then that will not be possible.

27.) Explain correlated sub-queries.

Correlated sub-queries are those queries wherein the nester query on the inner side refers directly back to outer query's table. For each and every row that is qualified, the evaluation of correlated sub-query is a must.

28.) Comment whether the cursor is closed during COMMIT or not.

Yes. The cursor is closed during COMMIT.

29.) In an SQL table that is embedded, what is the procedure to retrieve rows that are part of a DB2 table?

Either you can use SELECT statements of single rows or an alternative way is to use CURSOR.

30.) Mention the way of highlighting as well as putting a CURSOR to use in a COBOL program.

The best way of putting a CURSOR to use in a COBOL program is to make use of DECLARE CURSOR, which can be used either in procedure division operation or in working storage. This is being done basically to highlight the SELECT statement. Once DECLARE CURSOR is used, this is followed by OPEN, FETCH and finally CLOSE.

31.) If the CURSOR is kept open followed the issuing of COMMIT, what is the procedure to leave the CURSOR that way?

Inside DECLARE CURSOR, there is a WITH HOLD option, which will come useful in this case. Although, one should note the point that WITH HOLD function has absolutely no effect if considered for CICS pseudo-conventional programs.

32.) Explain PACKAGES.

PACKAGES are units, which consist of executable codes that are meant for SQL statements for one respective DBRM.

33.) Highlight all the advantages that are attached to a PACKAGE.

Following are the advantages attached to a PACKAGE.

- Avoid the cost of a large collection of bind. It is much more advisable to go for a small collection instead of a large one.
- Ensure that you do not have to bring a large collection of members of DBRM together for a particular plan.
- In case if you wish to make changes in the program and these changes eventually lead to errors, then you can decrease the fallback complexities with the help of PACKAGE.
- During the process of automatic binding and rebinding of a particular plan, make sure that the total transaction attached to the process is unavailable during the course.

34.) Mention the definition of COBOL in VARCHAR field.

The REMARKS of VARCHAR column are as follows: -
[crayon-5f2d021d66527744339458/]

35.) Mention the length of physical storage of the given data types of DB2 - DATE, TIMESTAMP, TIME

[crayon-5f2d021d66529408753762/]

36.) For a DB2 column that is being defined as DECIMAL (11, 2), discuss the COBOL picture clause.

PIC S9 (9) V99 COMP – 3

In the expression DECIMAL (11, 2) 2 happens to be the precision whereas 11 is the data type size.

37.) Explain DCLGEN.

DCLGEN basically refers to DeCLarations GENerator whose primary purpose is to generate copy books of the host language for the tables. It is also used to create the DECLARE table.

38.) Mention some fields that are a part of SQLCA.

SQLERRM, SQLCODE, SQLERRD.

39.) Explain the contents that are a part of DCLGEN.

There are primarily two components of DCLGEN. These are: -

- A copy book of the host language, which will give alternative definitions for all the

column names.

- In terms of the data-types of DB2, EXEC SQL DECLARE TABLE shows an aesthetic layout of the table.

40.) Comment whether DCLGEN is mandatorily used. If not, then what is the point of using it?

It is not exactly mandatory to bring DCLGEN into use. The primary use of DCLGEN comes during the pre-compilation phase where it helps in detecting the misspelt column names. Since DCLGEN is just a tool, it will only generate variable definitions for the host and reduces the chances of errors.

41.) In case if at some point of time DB2 is down, would that impact the pre-compilation process of a DB2-COBOL program?

Even if DB2 is down at some point of time, even then it will not impact the pre-compilation process of DB2-COBOL program. The reason for the same is that the pre-compiler never refers to the catalogue tables of DB2.

42.) Following a DB2 update statement, what is the quickest way to compute the total number of updated rows?

All you have to do is check the value that has been stored in SQLERRD (3).

43.) What is meant by EXPLAIN?

EXPLAIN is basically used to show the path of access by the optimizer basically for an SQL statement. Furthermore, EXPLAIN can also be brought to use in SPUFI or even in BIND step.

44.) Before you give the EXPLAIN statement, what are the prerogatives?

Before giving the EXPLAIN statement, we need to make sure that PLAN_TABLE has already been created under AUTHID.

45.) Mention the location where the output received from EXPLAIN statement is stored.

The output from EXPLAIN is stored in userid.PLAN_TABLE

46.) Outputs of EXPLAIN are with MATCHCOLS = 0. What does this signify?

This signifies a non-matching scan of index provided that ACCESSTYPE = I

47.) Mention the various locks that are available in DB2.

EXCLUSIVE, SHARE and UPDATE

48.) What is RELEASE/ACQUIRE in BIND?

There is a certain point in a program at which DB2 acquires or perhaps releases the locks against tables as well as table-spaces. These include intent locks.

49.) Mention the different locking levels that are available in DB2.

TABLE, PAGE and TABLESPACE

50.) Mention the downsides of PAGE level lock.

In case if there are large updates to be done, then the resource utilization is high accordingly.