

DB2 INTERVIEW QA

1. What do the initials DDL and DML stand for? What does they mean?

A. DDL is data definition language and DML is data manipulation language. DDL statements are CREATE, ALTER, TRUNCATE. DML statements are SELECT, INSERT, DELETE and UPDATE.

2. What is the difference between static SQL and dynamic SQL?

A. The static SQL statement is prepared before the program is executed and the operational form of the statement persists beyond the execution of the program. The source form of a static SQL statement is embedded within an application program written in a host language such as COBOL. Unlike static SQL, the dynamic statements are constructed and prepared at run time. The source form of a dynamic statement is a character string that is passed to DB2 by the program using the static SQL statement PREPARE or EXECUTE IMMEDIATE. Whether the operational form of the statement is persistent depends on whether dynamic statement caching is enabled.

3. What is deferred embedded SQL?

A. A deferred embedded SQL statement is neither fully static nor fully dynamic. Like a static statement, it is embedded within an application, but like a dynamic statement, it is prepared during the execution of the application. Although prepared at run time, a deferred embedded SQL statement is processed with bind-time rules such that the authorization ID and qualifier determined at bind time for the plan or package owner are used. Deferred embedded SQL statements are used for DB2 private protocol access to remote data.

4. What is DB2 ODBC?

A. DB2 Open Database Connectivity (DB2 ODBC) is an alternative to using embedded static or dynamic SQL. DB2 ODBC is an application-programming interface in which functions are provided to application programs to process SQL statements. The function calls are available only for C and C++ application programs. Through the interface, the application invokes a C function at execution time to connect to the data source, to issue SQL statements, and to get returned data and status information. Unlike using embedded SQL, no pre-compilation is required. Applications developed using this interface might be executed on a variety of data sources without being compiled against each of the databases. Note that only C and C++ applications can use this interface.

5. What are JDBC and SQLJ?

A. JavaSoft(TM) Java Database Connectivity (JDBC) and SQLJ are two methods for accessing DB2 data from the Java programming language. In general, Java applications use JDBC for dynamic SQL and SQLJ for static SQL.

6. What is a schema?

A. A schema is a collection of named objects. The objects that a schema can contain include distinct types, functions, stored procedures, and triggers. An object is assigned to a schema when it is created.

The schema name of the object determines the schema to which the object belongs. When a distinct type, function, or trigger is created, it is given a qualified, two-part name. The first part is the schema name (or the qualifier), which is either implicitly or explicitly specified. The second part is the name of the object. When a stored procedure is created, it is given a three-part name. The first part is a location name, which is implicitly or explicitly specified, the second part is the schema name, which is implicitly or explicitly specified, and the third part is the name of the object.

7. What are the differences between base table, auxiliary table and temporary table?

A. A base table is a table created with the SQL statement `CREATE TABLE` and used to hold persistent user data. An auxiliary table is a table created with the SQL statement `CREATE AUXILIARY TABLE` and used to hold the data for a column that is defined in a base table. A temporary table is a table described by the SQL statement `CREATE GLOBAL TEMPORARY TABLE` and used to hold data temporarily, such as the intermediate results of SQL transactions. Temporary tables persist as long as the application supports them. Table space and database operations, locking, logging, and recovery do not apply.

8. What is an index?

A. An index is an ordered set of pointers to rows of a base table or an auxiliary table. Each index is based on the values of data in one or more columns. An index is an object that is separate from the data in the table. When you define an index using the `CREATE INDEX` statement, DB2 builds this structure and maintains it automatically.

9. Define a composite key

A. A key composed of more than one column is called a composite key.

10. Define unique key

A. A unique key is a key that is constrained so that no two of its values are equal (no duplicates).

11. How can the uniqueness of a key be enforced? When does DB2 enforce this constraint?

A. The mechanism used to enforce the uniqueness of a key is a unique index. Thus, every unique key is a key of a unique index. Such an index is also said to have the `UNIQUE` attribute. A unique key can be defined using the `UNIQUE` clause of the `CREATE TABLE` statement. A table can have an arbitrary number of unique keys.

12. When does DB2 enforce the various constraints.

A. DB2 enforces the constraint during the execution of the `LOAD` utility and the `SQL INSERT` and `UPDATE` statements.

13. How many primary keys are possible for a table?

A. One

14. Describe primary index.

A. The unique index on a primary key is called a primary index. When a primary key is defined in a CREATE TABLE statement, the table is marked unavailable until the primary index is created by the user unless the CREATE TABLE statement is processed by the schema processor. In that case, DB2 automatically creates the primary index.

15. Does every table need a primary key?

A. No, primary keys are optional.

16. Define a trigger.

A. A trigger defines a set of actions that are executed when a delete, insert, or update operation occurs on a specified table. When such an SQL operation is executed, the trigger is said to be activated.

17. Which statement is used to create a trigger?

A. CREATE TRIGGER

18. How many indexes can be stored in an index space?

A. One

19. How many tables can be stored in a table space?

A. One or more

20. How many tables can be stored in a partitioned table space?

A. One

21. What is a view? Why use it?

A. A view is a named specification of a result table. The specification is an SQL SELECT statement that is effectively executed whenever the view is referenced in an SQL statement. In other words, a view is a virtual table made up of data from base tables and other views, but not stored separately.

22. Can an index be defined for a view? How does an index improve the performance of a view?

A. An index cannot be created for a view. However, an index created for a table on which a view is based might improve the performance of operations on the view.

23. How will you create a Read-only view? Can you use COMMIT and ROLLBACK in the application program when the program is executed under IMS or CICS?

A. No, under CICS and IMS environment, CICS or IMS performs commit and rollback

24. What is a DB2 package?

A. A package contains control structures used to execute SQL statements. Packages are produced during program preparation. During the program preparation, the pre-compiler generates DBRM (database request module) that contains SQL statements extracted from the source program. From DBRM, the bind operation generates operational form of SQL or internal control structures to access the data. All control structures in a package are derived from the SQL statements embedded in a single source program.

25. What is a DB2 plan?

A. An application plan relates an application process to a local instance of DB2, specifies processing options, and contains one or both of the following elements: A list of package names. The bound form of SQL statements taken from one or more DBRMs. Every DB2 application requires an application plan. Plans and packages are created using the DB2 subcommands BIND PLAN and BIND PACKAGE

26. What is the difference between a package and a plan? How does one bind 2 versions of a CICS transaction with the same module name in two different CICS regions that share the same DB2 subsystem?

A. Package and plan contain optimized code for SQL statements - a package for a single program, module or subroutine contained in the database request module (DBRM) library. A plan may contain multiple packages and pointers to packages. The one CICS module would then exist in a package that could be referenced in two different plans.

27. What is a distributed operation?

A. A DB2 application program can use SQL to access data at other database management systems (DBMSs) other than the DB2 at which the application's plan is bound. This DB2 is known as the local DB2. The local DB2 and the other DBMSs are called application servers. Any application server other than the local DB2 is considered a remote server, and access to its data is a distributed operation.

28. What is meant by local DB2?

A. The DB2 application server in which the application's plan is bound is known as local DB2.

29. How many connections are possible for an application program at a time?

A. At a time, only one connection is possible for an application program.

A) The local DB2 is connected to a remote a DBMS. The application program issues a SELECT statement. Then it tries to connect to another DBMS. What can you expect?

DB2 will not allow it. An application program can have only one connection at a time and cannot connect to new application server until it executes a commit or rollback operation.

B) What if the program closes the connection and tries to connect to another DBMS?

30. What are the different types of methods by which a local DB2 connect to a remote DBMS?

A. DB2 provides two methods of accessing data at remote application server

1. DRDA: This protocol allows the application program to connect to DB2 as well as other types of DBMS.

2. DB2 private protocol access: It allows DB2 connect to another DB2 running on a different application server

31. What does Remote unit of work mean?

A. A unit of work is a transaction made up of a related set of SQL statements. Under remote UOW, DB2 on one computer can send multiple, related set of SQL statements to another DBMS on a remote computer. The remote site performs the processing for the statements, but the sending site controls whether to commit or rollback the UOW.

32. What does an SQL connection mean?

A. An SQL connection is an association between an application process and a local or remote application server. SQL connections can be managed by the application or by using bind options. At any time:

33. Max number of columns in a DB2 table ?

A. 224. Tables without foreign keys and parent tables can have up to 750 columns. Tables with foreign keys can have up to 749 columns.

34. What is the maximum length of a column name?

A. 18

35. How will you set default values for columns? What is the minimum and max value of n in CHAR (n)?

A. PRIMARY_ID CHAR (8) WITH DEFAULT USER,
SQL_ID CHAR (8) WITH DEFAULT CURRENT SQLID

The min and max value of n in char(n) of var-char is 1 and 254.

36. What is a thread?

A thread is the connection between DB2 and some other subsystem, such as CICS or IMS/DC.

37. If not specified what are the default values for numbers, fixed-length strings, variable length strings dates, time and timestamp?

A. Numbers - 0 (SMALLINT, INTEGER, DECIMAL, or FLOAT); Fixed-length strings - blanks (CHAR or GRAPHIC); Varying-length strings - empty strings (VARCHAR, LONG VARCHAR, VARGRAPHIC, or LONG VARGRAPHIC); DATE - CURRENT DATE; TIME - CURRENT TIME; TIMESTAMP - CURRENT TIMESTAMP

37. What is RUNSTATS command in DB2? When will you need to run the RUNSTATS utility?

The RUNSTATS utility collects statistics about DB2 objects. These statistics can be stored in the DB2 catalog, and are used during the bind process by optimizer to choose the path in accessing data. If you never use RUNSTATS and subsequently rebind your packages or plans DB2 will not have the information that it needs to choose the most efficient access path. This can result in unnecessary I/O operations and excessive processor consumption. It also collects statistics used for space management.

Run RUNSTATS at least once against each table and its associated indexes, After a load, or after mass updates, inserts, deletes, or after REORG...

39. I need to view the number of tables owned by one particular Owner. Is it possible? If so, please give the SQL query for this?

A. The query `SELECT * FROM SYSIBM.SYSTABLES WHERE CREATOR = 'owner id'` this displays the table names with that If you want only the number of tables give the following query. `SELECT COUNT(*) FROM SYSIBM.SYSTABLES WHERE CREATOR = 'owner id'` Make sure that you are in correct subsystem.

40. What is JOIN and different types of JOIN.

A. The ability to join rows and combine data from two or more tables is one of the most powerful features of relational system. Three type of joins:1. Equi-join2.Non-equi-join3.self-join

41. Can I alter a table (e.g. adding a column) when other user is selecting some columns or updating some columns from the same table?

A. Yes possible. until the updating or selection is committed db2 table will not be restructured. new column definition will be there but it will not be included until all the tasks on the table are committed.

42. How many sub-queries can you combine together ?

A Total 16 queries and sub-queries are 15

43. What are the different methods of accessing db2 from TSO?

A. There are three ways in establishing tso/db2 connection 1. SPUFI 2. QMF 3. CATALOG VISIBILITY

44. How is the connection established between TSO & DB2?

A. A thread between TSO & DB2 is established while attempting to make connection between TSO & DB2.

45. What are buffer pools?

A. Buffer pools, also known as virtual buffer pools, are areas of virtual storage used temporarily to store pages of table spaces or indexes. When an application program needs to access a row of a table, DB2 retrieves the page containing that row and places the page in a buffer. If the row is changed, the buffer must be written back to the table space. If the needed data is already in a buffer, the application program will not have to wait for it to be retrieved from DASD. The result is faster performance. The sizes of virtual buffer pools can be changed while DB2 is running. The result is greater flexibility.

46. How many buffer pools are available in db2?

A. Ten 32k size buffer-pools and fifty 4k size buffer-pools (bp0 to bp49) default buffer pools are bp0,bp1,bp2 & bp32

47. How many Bufferpools are there in DB2 and what are they?

A. There are 4 Bufferpools. They are BP0,BP1,BP2 and BP32.

48. What is normalization? Explain the different rules?

A. Normalization rules help to avoid redundancies and inconsistencies in the data. It protects against update and delete anomalies and provide smaller tables and fewer total bytes. The rules of normalization are: 1st Normal Form:- No data item (repeating groups) should not be repeated within a given record.

2nd Normal Form: Each column that is not in the key provides a fact that depends on the entire key.

3rd Normal Form: Each non-key column provides a fact that it's independent of other non-key columns and depends only on key columns. 4th Normal Form: No row contains two or more independent multi-valued facts about an entity. 5th Normal Form: Sub-relations that cannot be reconstructed.

49. Explain de-normalization technique?

A. The rules of normalization do not consider performance. De-normalization concentrates on performance, but not on redundancy. What you have to consider is the trade-off--whether duplication, in several tables, of often-requested columns is less expensive than the time it takes to perform joins. This duplication of columns in multiple tables is de-normalization, and increases redundancy.

50. What are the possible reasons to use views instead of tables?

A. Some of your users might find that no single table contains all the data they need; rather, the data might be scattered among several tables. A view is an alternative way of describing data that exists in one or more tables. To limit access to certain kinds of data. One table might contain more data than the users want to see, or more than they should be authorized to see. For those situations, you can create views. To allow you to alter tables without affecting application programs

51. What is referential integrity?

A. Referential integrity refers to the consistency that must be maintained between primary and foreign keys.

It is the state in which all values of all foreign keys at a given DB2 are valid. Rules of referential integrity:

Every foreign key value must have a matching primary key value.

An insert into a primary key table can not violate referential integrity and no checking is required.

A deletion of foreign key value can not violate referential integrity and no checking is required.

Changes in primary key values are allowed only for those values that don't have a matching foreign key.

Insert and update rules: The insertion of any given foreign key value (non-null value) or an update to that value is allowed only if the matching value exists in the primary key.

Delete rules for primary key value:

If RESTRICT is specified on foreign key definition, deletion of primary key value is not allowed.

If CASCADE is specified, both primary key and the foreign key values will be deleted.

IF SET NULL is specified, the foreign key value will be set to null, provided nulls are allowed in the foreign key column.

52. When can an insert of a new primary key value threaten referential integrity?

A. Never. New primary key values are not a problem. However, the values of foreign key inserts must have corresponding primary key values in their related tables. And updates of primary key values may require changes in foreign key values to maintain referential integrity.

53. What is the self-referencing constraint?

A. A single table can include both a primary key and a related foreign key. The limitations to changes in the primary key that the foreign key defines are called self-referencing constraints. The foreign key in a self-referencing table must specify the DELETE CASCADE rule.

54. What happens if either the DELETE RESTRICT or DELETE SET NULL is specified on a self-referencing table during the execution of an ALTER TABLE statement?

A. DB2 issues error message defining an invalid constraint is attempted.

55. What is table check constraint?

Table check constraints designate the values that specific columns of a base table can contain, providing you a method of controlling the integrity of data entered into tables. You can create tables with table check constraints using the CREATE TABLE statement, or add the constraints with the ALTER TABLE statement. E.g.,

```
CREATE TABLE EMPSAL (ID INTEGER NOT NULL,  
                      SALARY INTEGER CHECK (SALARY >= 15000)).
```

56. What is Check integrity and check pending state?

A. If the check integrity is compromised, or cannot be guaranteed for a table, the table space or partition that contains the table is placed in a check pending state. Check integrity is the condition that exists when each row of a table conforms to the check constraints defined on that table.

57. What are delete-connected tables?

A. Tables related with a foreign key are called delete-connected because a deletion in the primary key table can affect or be affected by the contents of the foreign key table.

58. What is a cycle in DB2 mean?

A cycle is a situation in which tables are related to each other through both primary and foreign key relationships. I.e., A set of referential constraints in which each associated table is a descendent of itself.

59. Which statement would you use to find out whether DB2 has used index to access the data?

A. EXPLAIN statement or the EXPLAIN option of bind.

60. What are the difference types of indexes and how will you create them?

A. There are two types of indexes: TYPE 1 & TYPE 2. Use the statement CREATE INDEX to create either type. TYPE 2 index comes with DB2V4 onwards. With TYPE 2 index data can be retrieved faster as only the data pages are locked and not the index pages. Hence TYPE 2 index is recommended.

61. What is the default value of index when you use CREATE INDEX?

A. The default value depends on the value of LOCKSIZE for the associated table space.

1. If LOCKSIZE is ROW then the default index type is type 2. You can't use row locking with a type 1 index.

2. If LOCKSIZE is not ROW, then the default for CREATE INDEX is the type specified in field DEFAULT INDEX TYPE of installation panel DSNTIPE.

62. What is a Leaf Page?

A. Index page that points directly to the data in the table is called a leaf page. The leaf page contains the key of the record and a pointer to the record.

63. What is a root page?

A. The opposite of a leaf page; it is the highest-level index page. An index can contain only one root page; all other index pages are associated to the root. Each record of a root page contains a pointer to another index page and the highest of key of that page.

64. What is B-tree Index. Explain.

A. The simplest DB2 index is the B-tree and the B-tree's top page is called the root page. The root page entries represent the upper range limits of the index and are referenced first in a search.

65. How many indexes a table can have?

A. A table can have more than one index

66. What is the function of UNIQUE WHERE NOT NULL clause?

A. The clause is used with CREATE INDEX command. This ensures the uniqueness of the key column on which the index is defined, but it allows null values in the column. Or, DB2 will not allow duplication of non-null values in the key column.

67. What advantage does clustering index provide?

A. When a table has a clustering index, an INSERT statement inserts records as nearly as possible in the order of their index values. These clustered inserts can

provide a significant performance advantage in some operations, particularly those that involve many records, such as comparisons other than equal, grouping and ordering. Although a table can have several indexes, only one can be a Clustering index.

68. How many clustering indexes are allowed for a table?

A. Only one.

69. How will you specify a clustering index?

To specify a clustering index, use the CLUSTER clause in the CREATE INDEX statement.

70. You have loaded data using a non-clustering index. Now, You are defining a clustering index for the table. What is its effect on the new inserts?

A. It does not have any effect on the new insert. The data will still be organized using the non-clustering index. However, when the table space is reorganized by REORG utility, DB2 clusters data in accordance with the clustering index.

71. What is a partitioning index?

A. When you store a table in a partitioned table space, you tell DB2 how to divide the data among partitions by using the PART clause of a CREATE INDEX statement. The index that divides the data is called a partitioning index. It is also a clustering index, because the data is clustered by the index key values. Thus, your PART clause must be preceded by the CLUSTER clause.

72. Can you compress data in LOB space?

A. No

73. Name the different types of Table spaces.

- A. 1. Simple Table Space
2. Segmented Table Space
3. Partitioned Table Space (some partitioned table spaces can also be EA-enabled table spaces)
4. LOB Table Space

74. How will compress data in a table space or partition?

A. To compress data in a table space or partition, specify COMPRESS YES on CREATE TABLESPACE or ALTER TABLESPACE, then run LOAD or REORG. When you compress data, bit strings that occur frequently are replaced by shorter strings. Information about the mapping of bit strings to their replacements is stored in a compression dictionary. Computer processing is required to compress data before it is stored and to decompress the data when it is retrieved from storage.

75. B37 abend during SPUFI

A. The B37 ABEND in the SPUFI is because of space requirements, the query has resulted in so many rows that the SPUFI.OUT file is not large enough to handle it, increase the space allocation of spufi.out.

76. What is the command used by TSO users to invoke DB2?

A. DSN RUN

77. What is the error code -803?

A. Unique index violation (Duplicate rows are not allowed)

78. How do you filter out the rows retrieved from a Db2 table?

A. One way is to use The SQL WHERE clause.

79. What is a collection?

A. A collection is something that every programmer should assign/Specify for every package. The Collection name is about 1-18 characters long. It's a user defined name that is the anchor for packages. It has no physical existence. Main usage is to group packages.

80. What is Skeleton cursor table (SKCT)?

A. When you bind a plan, DB2 creates a skeleton cursor table in the skeleton cursor table space (SCT02). IT is the executable form of a Plan. This is stored in sysibm.sct02 table.

81. What is a record in a DB2 environment?

A. In DB2, a record is the storage representation of a row.

82. What is the size constrain of a row?

A. In DB2, records or rows are stored in 4K or 32K pages and a single record cannot occupy more than one page. Therefore, a table cannot be created with a maximum record size greater than the page size.

Why is it preferable to place variable length columns at the end of the row? When you use ALTER to add a new column to an existing table, where will the new column be placed?

A. When you use alter to add a new column to an existing table, it is added a the end of the row.

84. What's the equivalent Cobol Data type for Decimal(x,y) in DB2?

A. PIC S9(x-y)V9(Y) Comp-3;

85. What does the CURRENT SQLID register contain?

The current SQLID contains the current authorization ID.

86. Can we declare DB2 HOST variable in COBOL COPY book?

A. NO. If we declare DB2 host variable in COBOL COPY book, at the time of Pre-compilation we get the host variable not defined, because pre-compiler will not expand COBOL COPY book. So we declare it either in DCLGEN with EXEC SQL INCLUDE Dclgennam END-EXEC or we directly hardcode it in the working storage section.

87. What should be specified along with a cursor in order to continue updating process after commit?

A. With Hold option.

88. What is the name of the default DB2 Catalog Database?

A. DSNDB06

89. When can you be sure that a query will return only one row?

A. When you use the primary key and only the primary key in the where clause.

90. What is the difference between join and union?

A. Join is used to retrieve data from different tables using a single SQL statement. Union is used to combine the results of two or more SQL queries.

91. What is the difference between a UNION and UNION ALL?

A. UNION: eliminates duplicates. UNION ALL: retains duplicates. Both are used to combine the results from different SELECT statements.

92. What is a correlated sub-query?

Answer: In a sub-query, if the outer query refers back to the outcome of inner-query it is called correlated sub-query. That's why the outer query is evaluated first unlike an ordinary sub-query

93. What are the functions of Bind?

A. BIND mainly performs two things syntax checking and authorization checking. It binds together all packages into an application plan hence the name BIND. Apart from this bind has optimizer as a sub-component. Its function is to determine the optimum access strategy.

94. MAX. NO OF ROWS PER PAGE

A. 127

95. The only place of VSAM KSDS in DB2 is?

A. BSDS is a VSAM KSDS.

96. Question: Can all Users have the privilege to use the SQL Statement SELECT * (DML)?

A. No, the user should be granted privilege to use it.

97. Question: what's the best locksize that you could use when you create a tablespace?

A. The answer is Locksize = ANY. Unless you are Sure what's the Purpose of tablespace ie., Read-only or R/W. If you use lock size =any, Db2 would automatically determine what type of locks it should use.

98. What's the error code for Unique Index Violation:

A. -803

99. What's the percentage free space for

A. ZERO

100.Can you define an Index if the table size less than 10 PAGES?

The Answer is : NO

1.What's the Maximum Length of SQLCA and what's the content of SQLCABC?

A. The Max length is 136. and the SQLCABC has the Value of SQLCA.

2. What's the percentage free space for

A. The answer is ZERO.

3. What's the maximum number of volumes that can be added to a STOGROUP?

A. The answer is 133.Usually it will be difficult monitor more than 3 or 4 volumes to a STOGROUP.

4. What's the maximum number of characters that a table name can have?

A. The answer is 18 characters.

5. What is the meaning of -805 SQL return code?

A. Program name not in plan. Bind the plan and include the DBRM for the program named as part of plan.

6. When does the SQL statement gets executed when you use cursor in the application programming?

A. SQL statement gets executed when we open cursor

7. What does CURRENTDATA option in bind indicate?

A. CURRENTDATA option ensures block fetch while selecting rows from a table. In DB2V4 the default has been changed to NO. Therefore it is necessary to change all the bind cards with CURRENTDATA(YES), which is default in DB2V3 & earlier to CURRENTDATA(NO).

8. What are the levels of isolation available with DB2V4 ?

CS, RR, UR (added new for DB2V4 which stands for uncommitted read which allows to retrieve records from the space which has exclusive locks also but data integrity will be affected if this option is used)The best available option for data integrity & data concurrency is CS.

9. How do you achieve record locking in DB2 in the versions, which don't support record level locking?

A. Yesterday I had posted this queue. The answer should have read as follows: By having the record length more than half of the page size !Sorry again & Thanks

10. How do u achieve record level locking in DB2 versions when record level locking is not allowed?

A. By having the length of the record greater than that of a page!

11. In a DB2-CICS program, which is, acts as co-coordinator and which is participant?

A. DB2 - participant CICS- coordinator

12. What does DML stand for and what are some examples of it?

A. Data Manipulation Language. Some examples are SELECT, INSERT, DELETE, REPLACE.

13. How to define the data items to receive the fetch items for the SQL?

A. Using the DSECT, followed by lines of - 'dataitems DS datatype'.

14. How will you delete duplicate records from a table?

A. Delete From Table1 Where Id In (Select Id From Table1 As Temp Group By Id Having Count(*) >1)

15. What is the difference between Where and Having Clause?

A. WHERE is for Rows and HAVING is for Groups.

16. How to see the structure of db2 table?

A. Using QMF.

17. How do you declare a host variable (in COBOL) for an attribute named EMP-NAME of type VARCHAR(25) ?

A. 01 EMP-GRP. 49 E-LEN PIC S9(4) COMP. 49 E-NAME PIC X(25).

18. What is the maximum number of tables that can be stored on a Partitioned Table Space ?

A. ONE .

19. What are the max. & min. no. of partitions allowed in a partition tablespace?

A. Minimum is 4. maximum is 64.

20. What is the maximum number of tables that can be joined ?

A. Fifteen

21. What technique is used to retrieve data from more than one table in a single SQL statement?

A. The Join statement combines data from more than two tables

22. Question: What is a foreign key?

A. It identifies a related row in another table and establishes a logical relationship between rows in two tables.

23. Explain the use of the WHERE clause.

It directs DB2 to extract data from rows where the value of the column is the same as the current value of the host variable.

24. What is a DB2 bind?

Answer: A DB2 bind is a process that builds an access path to DB2 tables.

25. What is a DB2 access path?

A. An access path is the method used to access data specified in DB2 SQL statements.

26. What is normalization and what are the five normal forms?

A. Normalization is a design procedure for representing data in tabular format. The five normal forms are progressive rules to represent the data with minimal redundancy.

27. What are foreign keys?

A. These are attributes of one table that have matching values in a primary key in another table, allowing for relationships between tables.

28. Describe the elements of the SELECT query syntax.

SELECT element FROM table WHERE conditional statement.

29. Explain the use of the WHERE clause.

A. WHERE is used with a relational statement to isolate the object element or row.

30. What techniques are used to retrieve data from more than one table in a single SQL statement?

A. Joins, unions and nested selects are used to retrieve data.

31. Explain an outer join.

A. An outer join includes rows from tables when there are no matching values in the tables.

32. What is a subselect? Is it different from a nested select?

A. A subselect is a select which works in conjunction with another select. A nested select is a kind of subselect where the inner select passes to the where criteria for the outer select.

33. What is the difference between group by and order by?

A. Group by controls the presentation of the rows, order by controls the presentation of the columns for the results of the SELECT statement.

34. Explain the EXPLAIN statement.

A. The explain statement provides information about the optimizer's choice of access path of the SQL. It can be used in SPUFI (for single SQL statement) or in BIND step (for embedded SQL).

35. What is a tablespace?

A. Tables are stored in tablespaces (hence the name)! There are three types of tablespaces: simple, segmented and partitioned.

36. What is a cursor and what is its function?

A. An embedded SQL statement may return a number of rows while the programming language can only access one row at a time. Cursor is a programming device that allows the SELECT to find a set of rows but return them one at a time. Cursor should be used because the host language can deal with only one row at a time.

37. Usually, which is more important for DB2 system performance - CPU processing or I/O access?

A. I/O operations are usually most critical for DB2 performance (or any other database for that matter).

38. Is there any advantage to de-normalizing DB2 tables?

A. Denormalizing DB2 tables reduces the need for processing intensive relational joins and reduces the number of foreign keys.

39. What is the database descriptor?

A. The database descriptor, DBD is the DB2 component that limits access to the database whenever objects are created, altered or dropped.

40. What is lock contention?

A. To maintain the integrity of DB2 objects the DBD permits access to only one object at a time. Lock contention happens if several objects are required by contending application processes simultaneously.

41. What is SPUFI?

A. SPUFI stands for SQL processing using file input. It is the DB2 interactive menu-driven tool used by developers to create database objects.

42. What is the significance of DB2 free space and what parameters control it?

A. The two parameters used in the CREATE statement are the PCTFREE which specifies the percentage of free space for each page and FREEPAGE which indicates the number of pages to be loaded with data between each free page. Free space allows room for the insertion of new rows.

43. What is a NULL value? What are the pros and cons of using NULLS?

A. A NULL value takes up one byte of storage and indicates that a value is not present as opposed to a space or zero value. It's the DB2 equivalent of TBD on an organizational chart and often correctly portrays a business situation. Unfortunately, it requires extra coding for an application program to handle this situation.

44. What is a synonym? How is it used?

A synonym is used to reference a table or view by another name. The other name can then be written in the application code pointing to test tables in the development stage and to production entities when the code is migrated. The synonym is linked to the AUTHID that created it, so it is accessible only by creator.

45. What is an alias and how does it differ from a synonym?

A. An alias is an alternative to a synonym, designed for a distributed environment to avoid having to use the location qualifier of a table or view. The alias is not dropped when the table is dropped.

46. What is a LIKE table and how is it created?

A. A LIKE table is created by using the LIKE parameter in a CREATE table statement. LIKE tables are typically created for a test environment from the production environment.

47. If the base table underlying a view is restructured, e.g. attributes are added, does the application code accessing the view need to be redone?

A. No. The table and its view are created anew, but the programs accessing the view do not need to be changed if the view and attributes accessed remain the same.

48. Under what circumstances will DB2 allow an SQL statement to update more than one primary key value at a time?

A. Never. Such processing could produce duplicate values violating entity integrity. Primary keys must be updated one at a time.

49. What is the cascade rule and how does it relate to deletions made with a subselect?

A. The cascade rule will not allow deletions based on a subselect that references the same table from which the deletions are being made.

Q1. What is a DB2 bind?

A1. A DB2 bind is a process that builds an access path to DB2 tables.

Q2. What is a DB2 access path?

A2. An access path is the method used to access data specified in DB2 sql statements.

Q3. What is a DB2 plan?

A3. An application plan or package is generated by the bind to define an access path.

Q4. What is normalization and what are the five normal forms?

A4. Normalization is a design procedure for representing data in tabular format. The five normal forms are progressive rules to represent the data with minimal redundancy.

Q5. What are foreign keys?

A5. These are attributes of one table that have matching values in a primary key in another table, allowing for relationships between tables.

Q6. Describe the elements of the SELECT query syntax.

A6. SELECT element FROM table WHERE conditional statement.

Q7. Explain the use of the WHERE clause.

A7. WHERE is used with a relational statement to isolate the object element or row.

Q8. What techniques are used to retrieve data from more than one table in a single SQL statement?

A8. Joins, unions and nested selects are used to retrieve data.

Q9. What do the initials DDL and DML stand for and what is their meaning?

A9. DDL is data definition language and DML is data manipulation language. DDL statements are CREATE, ALTER, TRUNCATE. DML statements are SELECT, INSERT, DELETE and UPDATE.

Q10. What is a view? Why use it?

A10. A view is a virtual table made up of data from base tables and other views, but not stored separately.

Q11. Explain an outer join.

A11. An outer join includes rows from tables when there are no matching values in the tables.

Q12. What is a subselect? Is it different from a nested select?

A12. A subselect is a select which works in conjunction with another select. A nested select is a kind of subselect where the inner select passes to the where criteria for the outer select.

Q13. What is the difference between group by and order by?

A13. Group by controls the presentation of the rows, order by controls the presentation of the columns for the results of the SELECT statement.

Q14. Explain the EXPLAIN statement.

A14. The explain statement provides information about the optimizer's choice of access path of the sql.

Q15. What is tablespace?

A15. Tables are stored in tablespaces (hence the name)! There are three types of tablespaces: simple, segmented and partitioned.

Q16. What is a cursor and what is its function?

A16. An embedded sql statement may return a number of rows while the programming language can only access one row at a time. The programming device called a cursor controls the position of the row.

Q17. What is referential integrity?

A17. Referential integrity refers to the consistency that must be maintained between primary and foreign keys, ie every foreign key value must have a corresponding primary key value.

Q18. Usually, which is more important for DB2 system performance - CPU processing or I/O access?

A18. I/O operations are usually most critical for DB2 performance (or any other database for that matter).

Q19. Is there any advantage to denormalizing DB2 tables?

A19. Denormalizing DB2 tables reduces the need for processing intensive relational joins and reduces the number of foreign keys.

Q20. What is the database descriptor?

A20. The database descriptor, DBD is the DB2 component that limits access to the database whenever objects are created, altered or dropped.

Q21. What is lock contention?

A21. To maintain the integrity of DB2 objects the DBD permits access to only one object at a time. Lock contention happens if several objects are required by contending application processes simultaneously.

Q22. What is SPUFI?

A22. SPUFI stands for SQL processing using file input. It is the DB2 interactive menu-driven tool used by developers to create database objects.

Q23. What is the significance of DB2 free space and what parameters control it?

A23. The two parameters used in the CREATE statement are the PCTFREE which specifies the percentage of free space for each page and FREEPAGE which indicates the number of pages to be loaded with data between each free page. Free space allows room for the insertion of new rows.

Q24. What is a NULL value? What are the pros and cons of using NULLS?

A24. A NULL value takes up one byte of storage and indicates that a value is not present as opposed to a space or zero value. It's the DB2 equivalent of TBD on an organizational chart and often correctly portrays a business situation. Unfortunately, it requires extra coding for an application program to handle this situation.

Q25. What is a synonym? How is it used?

A25. A synonym is used to reference a table or view by another name. The other name can then be written in the application code pointing to test tables in the development stage and to production entities when the code is migrated. The synonym is linked to the AUTHID that created it.

Q26. What is an alias and how does it differ from a synonym?

A26. An alias is an alternative to a synonym, designed for a distributed environment to avoid having to use the location qualifier of a table or view. The alias is not dropped when the table is dropped.

Q27. What is a LIKE table and how is it created?

A27. A LIKE table is created by using the LIKE parameter in a CREATE table statement. LIKE tables are typically created for a test environment from the production environment.

Q28. If the base table underlying a view is restructured, eg. attributes are added, does the application code accessing the view need to be redone?

A28. No. The table and its view are created anew, but the programs accessing the view do not need to be changed if the view and attributes accessed remain the same.

Q29. Under what circumstances will DB2 allow an SQL statement to update more than one primary key value at a time?

A29. Never. Such processing could produce duplicate values violating entity integrity. Primary keys must be updated one at a time.

Q30. What is the cascade rule and how does it relate to deletions made with a subselect.

A30. The cascade rule will not allow deletions based on a subselect that references the same table from which the deletions are being made.

Q31. What is the self-referencing constraint?

A31. The self-referencing constraint limits in a single table the changes to a primary key that the related foreign key defines. The foreign key in a self referencing table must specify the DELETE CASCADE rule.

Q32. What are delete-connected tables?

A32. Tables related with a foreign key are called delete-connected because a deletion in the primary key table can affect the contents of the foreign key table.

Q33. When can an insert of a new primary key value threaten referential integrity?

A33. Never. New primary key values are not a problem. However, the values of foreign key inserts must have corresponding primary key values in their related tables. And updates of primary key values may require changes in foreign key values to maintain referential integrity.

Q34. In terms of DB2 indexing, what is the root page?

A34. The simplest DB2 index is the B-tree and the B-tree's top page is called the root page. The root page entries represent the upper range limits of the index and are referenced first in a search.

Q35. How does Db2 use multiple table indexes?

A35. DB2 use the multiple indexes to satisfy multiple predicates in a SELECT statement that are joined by an AND or OR.

Q36. What are some characteristics of columns that benefit from indexes?

A36. Primary key and foreign key columns; columns that have unique values; columns that have aggregates computed frequently and columns used to test the existence of a value.

Q37. What is a composite index and how does it differ from a multiple index?

A37. A multiple index is not one index but two indexes for two different columns of a table. A composite index is one index made up of combined values from two columns

in a table. If two columns in a table will often be accessed together a composite index will be efficient.

Q38. What is meant by index cardinality?

A38. The number of distinct values for a column is called index cardinality. DB2's RUNSTATS utility analyzes column value redundancy to determine whether to use a tablespace or index scan to search for data.

Q39. What is a clustered index?

A39. For a clustered index DB2 maintains rows in the same sequence as the columns in the index for as long as there is free space. DB2 can then process that table in that order efficiently.

Q40. What keyword does an SQL SELECT statement use for a string search?

A40. The LIKE keyword allows for string searches. The % sign is used as a wildcard.

Q41. What are some sql aggregates and other built-in functions?

A41. The common aggregate, built-in functions are AVG, SUM, MIN, MAX, COUNT and DISTINCT.

Q42. How is the SUBSTR keyword used in sql?

A42. SUBSTR is used for string manipulation with column name, first position and string length used as arguments. Eg. SUBSTR (NAME, 1 3) refers to the first three characters in the column NAME.

Q43. What are the three DB2 date and time data types and their associated functions?

A43. The three data types are DATE, TIME and TIMESTAMP. CHAR can be used to specify the format of each type. The DAYS function calculates the number of days between two dates. (It's Y2K compliant).

Q44. Explain transactions, commits and rollbacks in DB2.

A44. In DB2 a transaction typically requires a series of updates, insertions and deletions that represent a logical unit of work. A transaction puts an implicit lock on the DB2 data. Programmers can use the COMMIT WORK statement to terminate the transaction creating smaller units for recovery. If the transaction fails DB2 uses the log to roll back values to the start of the transaction or to the preceding commit point.

Q45. What is deadlock?

A45. Deadlock occurs when transactions executing at the same time lock each other out of data that they need to complete their logical units of work.

Q46. What are the four lockable units for DB2?

A46. DB2 imposes locks of four differing sizes: pages, tables, tablespace and for indexes subpage.

Q47. What are the three lock types?

A47. The three types are shared, update and exclusive. Shared locks allow two or more programs to read simultaneously but not change the locked space. An exclusive lock bars all other users from accessing the space. An update lock is less restrictive; it allows other transactions to read or acquire shared locks on the space.

Q48. What is isolation level?

A48. SQL statements may return any number of rows, but most host languages deal with one row at a time by declaring a cursor that presents each row at a unique isolation level.

Q49. What is an intent lock?

A49. An intent lock is at the table level for a segmented tablespace or at the tablespace level for a nonsegmented tablespace. They indicate at the table or tablespace level the kinds of locks at lower levels.

Q50. What is the difference between static and dynamic sql?

A50. Static sql is hard-coded in a program when the programmer knows the statements to be executed. For dynamic sql the program must dynamically allocate memory to receive the query results.

Q51. What is cursor stability?

A51. Cursor stability means that DB2 takes a lock on the page the cursor is accessing and releases the lock when the cursor moves to another page.

Q52. What is the significance of the **CURSOR WITH HOLD** clause in a cursor declaration?

A52. The clause avoids closing the cursor and repositioning it to the last row processed when the cursor is reopened.

Q53. What is the SQL Communications Area and what are some of its key fields?

A53. It is a data structure that must be included in any host-language program using SQL. It is used to pass feedback about the sql operations to the program. Fields are return codes, error messages, handling codes and warnings.

Q54. What is the purpose of the **WHenever** statement?

A54. The **WHenever** statement is coded once in the host program to control program actions depending on the SQL-CODE returned by each sql statement within the program.

Q55. What is **DCLGEN**?

A55. **DCLGEN** stands for declarations generator; it is a facility to generate DB2 sql data structures in COBOL or PL/I programs.

Q56. What is the **FREE** command?

A56. The FREE command can be used to delete plans and/or packages no longer needed.

Q57. DB2 can implement a join in three ways using a merge join, a nested join or a hybrid join. Explain the differences.

A57. A merge join requires that the tables being joined be in a sequence; the rows are retrieved with a high cluster ratio index or are sorted by DB2. A nested join does not require a sequence and works best on joining a small number of rows. DB2 reads the outer table values and each time scans the inner table for matches. The hybrid join is a nested join that requires the outer table be in sequence.

Q58. Compare a subselect to a join.

A58. Any subselect can be rewritten as a join, but not vice versa. Joins are usually more efficient as join rows can be returned immediately, subselects require a temporary work area for inner selects results while processing the outer select.

Q59. What is the difference between IN subselects and EXISTS subselect?

A59. If there is an index on the attributes tested an IN is more efficient since DB2 uses the index for the IN. (IN for index is the mnemonic).

Q60. What is a Cartesian product?

A60. A Cartesian product results from a faulty query. It is a row in the results for every combination in the join tables.

Q61. 4/99 Mail from Joseph Howard: 'Q: DB2 What is the difference between a package and a plan? How does one bind 2 versions of a CICS transaction with the same module name in two different CICS regions that share the same DB2 subsystem?'

A61. Package and plan are usually used synonymously, as in this site. Both contain optimized code for SQL statements - a package for a single program, module or subroutine contained in the database request module (DBRM) library. A plan may contain multiple packages and pointers to packages. The one CICS module would then exist in a package that could be referenced in two different plans.

Q62. What is an asynchronous write?

A62. It is a write to disk that may occur before or long after a commit. The write is controlled by the buffer manager.

Q63. What is a lock?

A63. A lock is the mechanism that controls access to data pages and tablespaces.

Q64. What is meant by isolation level?

A64. This is a key concept for any relational database. Isolation level is the manner in which locks are applied and released during a transaction. For DB@ a 'repeatable read' holds all locks until the transaction completes or a syncpoint is issued. For

transactions using 'cursor stability' the page lock releases are issued as the cursor 'moves', i.e. as the transaction releases addressability to the records.

Q65. What are leaf pages?

A65. They are the opposite of root pages. Leaf pages are the lowest level index pages - the pages that contain index entries and information to the corresponding table rows.

Q66. What is a precompiler?

A66. It is a DB2 facility for static SQL statements - it replaces these statements with calls to the DB2 language interface module.

Q67. What is a root page?

A67. The opposite of a leaf page; it is the highest level index page. An index can contain only the one root page; all other index pages are associated to the root.

Q68. What is a thread?

A68. A thread is the connection between DB2 and some other subsystem, such as CICS or IMS/DC.

Question: How do you debug a DB2 STORED PROCEDURE

Answer: I want someone to answer it.

Question: maxx number of columns in a db2 table

Answer: 224

Question: What is RUNSTATS command in DB2?

Answer: .

Question: Question: Question: I need to view the number of tables existing under one particular Owner. Is it possible? If so, pl give the SQL query for this?

Answer: The query `SELECT * FROM SYSIBM.SYSTABLES WHERE CREATOR = 'owner id'` This displays the table names with that If you want only the number of tables give the following query. `SELECT COUNT(*) FROM SYSIBM.SYSTABLES WHERE CREATOR = 'owner id'` Make sure that you are in correct subsystem.

Question: Question: I need to view the number of tables existing under one particular Owner. Is it possible? If so, pl give the SQL query for this?

Answer: The query `SELECT * FROM SYSTABLES WHERE OWNER=` should work.

Question: Question: I need to view the number of tables existing under one particular Owner. Is it possible? If so, pl give the SQL query for this?

Answer: Answer to Ravi's query: Db2 records information for its operation in a catalog which is actually a group of tables. So we can use the SYSTABLES to get answer to ur query. So

Question: I need to view the number of tables existing under one particular Owner. Is it possible? If so, pl give the SQL query for this?

Answer: Can any one give me the answer for this?

Question: I need to view the number of tables existing under one particular Owner. Is it possible? If so, pl give the SQL query for this?

Answer: Can any one give me the answer for this?

Question: What is JOIN and different types of JOIN.

Answer: The ability to join rows and combine data from two or more tables is one of the most powerful features of relational system. Three type of joins: 1. Equi-join 2. Non-equi-join 3. self-join

Question: can I alter a table (e.g. adding a column) when other user is selecting some columns or updating some columns from the same table?

Answer: yes possible. until the updation or selection is committed db2 table will not be restructured. new column definition will be there but it will not be included until all the tasks on the table are committed.

Question: How many subqueries can you combine together ?

Answer: Total 16 queries and subqueries are 15

Question: A. What are the different methods of accessing db2 from tso? B. How is the connection established between TSO & DB2?

Answer: A. There are three ways in establishing tso/db2 connection 1. SPUFI 2. QMF 3. CATALOG VISIBILITY B. A thread between TSO & DB2 is established while attempting to make connection between tso & db2.

Question: How many buffer pools are available in db2?

Answer: ten 32k size bufferpools and fifty 4k size buffer pools (bp0 to bp49) default buffer pools are bp0, bp1, bp2 & bp32

Question: B37abend?

Answer: the b37 abend in the spufi is because of space requirements, the query has resulted in so many rows that the spufi.out file is not large enough to handle it, increase the space allocation of spufi.out file.

Question: How many Bufferpools are there in DB2 and what are they?

Answer: There are 4 Bufferpools. They are BP0, BP1, BP2 and BP32.

Question: What is the command used by TSO users to invoke DB2?

Answer: DSN RUN

Question: what is the error code -803 ?

Answer: unique index violation

Question: How do you install DB2

Answer: Install DB2 according to the procedure in the manual

Question: how do you filter out the rows retrieved from a Db2 table ?

Answer: one way is to use The Sql WHERE clause.

Question: what is a collection?

Answer: A collection is something that every programmer should assign/Specify for every package. this about 1-18 characters long.

Question: What is Skeleton cursor table (SKCT)?

Answer: The Executable form of a Plan. This is stored in sysibm.sct02 table.

Question: what's the equivalent Cobol Data type for Decimal(x,y) in DB2? what does the current SQLID register contain?

Answer: Pic s9(x-y)V9(Y) Comp-3; the current SQLID contains the current authorization ID.

Question: Can we declare DB2 HOST variable in COBOL COPY book?

Answer: NO.If we declare DB2 host variable in COBOL COPY book, at the time of Pre-compilation we get the host variable not defined, because pre-compiler will not expand COBOL COPY book. So we declare it either in DCLGEN with EXEC SQL INCLUDE Dclgennname END-EXEC or we directly hardcode it in the working storage section.

Question: What should be specified along with a cursor in order to continue updating process after commit?

Answer: With Hold option.

Question: WHAT IS THE NAME OF THE DEFAULT db2 CATALOG DATABASE?

Answer: DSNDB06

Question: When Can you be sure that a query will return only one row?

Answer: When you use the primary key and only the primary key in the where clause.

Question: what is the difference between join and union?

Answer: join is used to retrieve data from different tables using a single sql statement.union is used to combine the results of two or more sql queries.

Question: What is a corelated subquery?

Answer: In a subquery, if the outer query reffers back to the outcome of innerquery it is called corelated subquery. That's why the outer query is evaluated first unlike an ordinary subquery

Question: What are the functions of Bind?

Answer: BIND mainly performs two things syntax checking and authorization checking.It binds together all packages into an application plan hence the name BIND.Apart from this bind has optimiser as a subcomponent.Its function is to determine the optimum access strategy.

Question: MAX. NO OF ROWS PER PAGE

Answer: 127

Question: The only place of VSAM KSDS in DB2 is?

Answer: BSDS is a VSAM KSDS.

Question: CAN ALL USERS HAVE THE PRIVILAGE TO USE THE SQL STATEMENT SELECT * (DML)?

Answer: NO THE USER SHOULD BE GRANTED PRIVILAGE TO USE IT.

Question: What is the size of a data page?

Answer: 4K to 8K

Question: what's the best locksize that you could use when you create a tablespace?

Answer: The answer is Locksize = ANY. Unless you are Sure what's the Purpose of tablespace ie., Read-only or R/W. If you use lock size = any, Db2 would automatically determine what type of locks it should use.

Question: what's the error code for Unique Index Voilation:

Answer: -803

Question: what's the percentage free space for

Answer: ZERO

Question: Can you define an Index if the table size less than 10 PAGES?

Answer: the Answer is : NO

Question: What's the Maximum Length of SQLCA and what's the content of SQLCABC?

Answer: The Max length is 136. and the SQLCABC has the Value of SQLCA.

Question: what's the percentage free space for

Answer: The answer is ZERO.

Question: What's the maximum number of volumes that can be added to a STOGROUP?

Answer: The answer is 133. Usually it will be difficult monitor more than 3 or 4 volumes to a Stogroup.

Question: What's the maximum number of characters that a tablename can have?

Answer: The answer is 18 characters.

Question: What is the meaning of -805 sql return code?

Answer: Program name not in plan. Bind the plan and include the DBRM for the program named as part of the plan.

Question: when does the sql statement gets executed when you use cursor in the application programming ?

Answer: sql statement gets executed when we open cursor

Question: What does CURRENTDATA option in bind indicate

Answer: CURRENTDATA option ensures block fetch while selecting rows from a table. In DB2V4 the default has been changed to NO. Therefore it is necessary to change all the bind cards with CURRENTDATA(YES) which is default in DB2V3 & earlier to CURRENTDATA(NO).

Question: What is the difference between TYPE 1 index & TYPE 2 index

Answer: TYPE 1 & TYPE 2 are specified when an index is created on the table. TYPE 2 index is the option which comes with DB2V4. With TYPE 2 index data can be retrieved faster as only the data pages are locked and not the index pages. Hence TYPE 2 index is recommended.

Question: What are the levels of isolation available with DB2V4

Answer: CS RR UR(added new for DB2V4 which stands for uncommitted read which allows to retrieve records from the space which has exclusive locks also but data integrity will be affected if this option is used)The best available option for data integrity & data concurrency is CS.

Question: How do you achieve record locking in DB2 in the versions which donot support record level locking?

Answer: Y'day I had posted this que. The answer shud hv read as follows:By having the record length more than half of the page size !Sorry again& Thanx

Question: How do u achieve record level locking in DB2 versions when record level locking is not allowed?

Answer: By having the length of the record greater than that of a page!

Question: In a DB2-CICS program which is acts as co-ordinator and which is participant?

Answer: DB2 - participant CICS- coordinator

Question: What does DML stand for and what are some examples of it?

Answer: Data Manipulation Language. Some examples are SELECT, INSERT, DELETE, REPLACE.

Question: How to define the dataitems to receive the fetch items for the SQL?

Answer: Using the DSECT, followed by lines of - 'dataitems DS datatype'.

Question: Re:How will you delete duplicate records from a table?

Answer: Delete From Table1Where Id In (Select Id From Table1 As Temp Group By Id Having Count(*) >1)

Question: What is the difference between Where and Having Clause

Answer: WHERE is for Rows and HAVING is for Groups

Question: How to see the structure of db2 table??

Answer: Using QMF.

Question: How do you declare a host variable (in COBOL) for an attribute named EMP-NAME of type VARCHAR(25) ?

Answer: 01 EMP-GRP. 49 E-LEN PIC S9(4) COMP. 49 E-NAME PIC X(25).

Question: What is the maximum number of tables that can be stored on a Partitioned Table Space ?

Answer: ONE

Question: Name the different types of Table spaces.

Answer: 1. Simple Table Space 2. Segmented Table Space and 3. Partitioned Table Space

Question: what are the max. & min. no. of partitions allowed in a partition tablespace?

Answer: minimum is 4, maximum is 64.

Question: what is the maximum number of tables that can be joined ?

Answer: fifteen

Question: What technique is used to retrieve data from more than one table in a single SQL statement?

Answer: The Join statement combines data from more than two tables

Question: What is a foreign key?

Answer: It identifies a related row in another table and establishes a logical relationship between rows in two tables.

Question: Explain the use of the WHERE clause.

Answer: It directs DB2 to extract data from rows where the value of the column is the same as the current value of the host variable.

Q1: How do you find the second max and second min in a given table?

Ans-1.

```
SELECT MAX(COLA) FROM TABLENAME WHERE COLA < (SELECT MAX(COLA)
FROM TABLENAME) ;
SELECT MIN(COLA) FROM TABLENAME WHERE COLA > (SELECT MIN(COLA)
FROM TABLENAME) ;
```

Q2: What is Cursor Stability(CS)?

Ans-2. A page lock is held only while the cursor is positioned on that page. When the cursor moves to another page, the lock is released. When a page is locked concurrent application programs cannot update or delete a row of the locked page. The current lock is not released until a new lock is acquired. If an application program updates or deletes data, the lock is held until the data is committed. CS applies only to data that is read. All changed data remains locked until COMMIT or ROLLBACK

Q3: What are the three lock types?

Ans-3. The three types are shared, update and exclusive. Shared locks allow two or more programs to read simultaneously but not change the locked space. An exclusive lock bars all other users from accessing the space. An update lock is less restrictive; it allows other transactions to read or acquire shared locks on the space.

INTRODUCION TO DATABASE MANAGEMENT SYSTEMS

1. What is a collection of data designed to be used by different people called?

- Table
- Tuple
- Database
- File

[Ans](#)

2. DB2 is the product of

- Microsoft
- International Business Machines Corporation
- Powersoft
- None of the Above

[Ans](#)

3. DB2 is a

- Database/Data Communication System
- Database
- Relational Database Management System
- Programming Language

[Ans](#)

4. True or False.

- Using a database will reduce data redundancy.
- True
- False

[Ans](#)

5. What are the characteristics of a Database Management System?

- Data Independence
- Non-redundancy
- Security protection
- All of the above

[Ans](#)

6. True or False

- The three different types of database management systems are relational, hierarchical and network models.
- True
- False

[Ans](#)

7. In relational model rows are called
- Rows
 - Relations
 - Tuples
 - None of the above

[Ans](#)

8. In relational model columns are called
Columns
Relations
Attributes
None of the above

[Ans](#)

9. What is a pool of values from which the actual values appearing in a given column are drawn called?
Domain
Superset
Relations
None of the above

[Ans](#)

10. True or False
In a hierarchical model data is represented as a tree structure.
True
False

[Ans](#)

11. True or False
In a network model data is represented by records and links.
True
False

[Ans](#)

12. Which of the following is a Relational Database Management System?
DB2
ORACLE
SYBASE
All of the above

[Ans](#)

13. Which of the following is a hierarchical database management system?
IMS
DB2
IDMS
None of the above

[Ans](#)

14. Which of the following is a Network database Management system?
IDMS
IMS
DB2
SYSTEM R

[Ans](#)

RELATIONAL DATABASE MANAGEMENT SYSTEMS (RDBMS)

15. Who is called the father of Relational Database Management Systems?

- C.J. Date
- Mullins
- E.F. Codd
- None of the above

[Ans](#)

16. The database language SEQUEL was introduced by

- E.F. Codd
- Chamberlin and Boyce
- Mullins
- None of the above

[Ans](#)

17. What is the name of the first prototype of SEQUEL introduced by IBM in 1974?

- SQL
- SEQUEL-XRM
- SEQUEL
- None of the above

[Ans](#)

18. Which was the first successful prototype of a relational database management system ?

- XRM
- SYSTEM-R
- DB2/2
- None of the above

[Ans](#)

19. Which is the first successful commercial RDBMS

- DB2
- ORACLE
- SYBASE
- None of the above

[Ans](#)

20. True or False

The relational model is concerned with three aspects of data: data structure, data integrity, and data manipulation

- True
- False

[Ans](#)

21. What is the number of attributes in a relation called?

- Degree

Domain
Number
None of the above

[Ans](#)

22. What is the number of tuples or rows in a relation called?

- Degree
- Cardinality
- Rowcount
- None of the above

[Ans](#)

23. True or False

The primary key is the unique identifier of the rows in a table.

- True
- False

[Ans](#)

24. True or False

There can be more than one primary key for a table.

- True
- False

[Ans](#)

25. True or False

The value of a primary key can be a null.

- True
- False

[Ans](#)

26. What is an attribute or attribute combination of one table whose values are required to match those of the primary key of some other table called?

- Primary key
- Foreign key
- Composite key
- Alternate key

[Ans](#)

27. True or False

The foreign key and the primary key should be defined on the same underlying domain

- True
- False

[Ans](#)

28. A given foreign key value must have matching primary key value in some tuple of the referenced relation if that foreign key value is non null. This principle is called

- Entity Integrity
- Referential Integrity
- Primary Integrity
- None of the above

[Ans](#)

29. A key made of more than one column is called
- Composite key
 - Complex key
 - Foreign key
 - None of the above

[Ans](#)

OVERVIEW OF DB2

30. True or False

'DB2' is an abbreviation for 'IBM DATABASE 2'.

True

False

[Ans](#)

31. The operating system in which DB2 run is

OS/2

UNIX

XENIX

MVS

[Ans](#)

32. True or False

DB2 does not support SQL.

True

False

[Ans](#)

33. What is the interactive SQL interface of DB2 called ?

DB2PM

DB2SQL

DB2I

None of the above

[Ans](#)

34. True or False

In DB2, the SQL statements can be used in all of the following host languages: C, COBOL, FORTRAN. PL/I and System /370 Assembler language.

True

False

[Ans](#)

35. True or False

The DB2 system can be used by more than one user at a time.

True

False

[Ans](#)

36. What are the two kinds of tables in DB2?

Tables and Files

Tables and Indexes

Base tables and Views

All of the above

[Ans](#)

37. True or False

A view is a virtual table, a table which does not physically exist.

True

False

[Ans](#)

38. True or False

The major components of DB2 are system services component, Locking services component, Database services component and Distributed data facility component.

True

False

[Ans](#)

39. What are the functions of the System services component?

System operation

Operator communication

Logging

All of the above

[Ans](#)

40. True or False

The Locking services component provides the necessary controls for managing concurrent access to data.

True

False

[Ans](#)

41. What are the functions of the Database services component, which supports the definition, retrieval and update of user and system data.

Data definition

Data retrieval

Data manipulation

All of the above

[Ans](#)

42. True or False

The Distributed data facility component provides DB2's distributed database support

True

False

[Ans](#)

43. True or False

DB2 directory can be accessed by SQL statements

True

False

[Ans](#)

44. True or False

The DB2 catalog consists of regular tables, and is accessible by means of SQL data retrieval statements.

True

False

[Ans](#)

45. What is the function of the DB2 pre-compiler?

Removes all the SQL statements in the source code

Replaces all the SQL statements with host language CALL statements

Converts SQL statements into host language statements

None of the above

[Ans](#)

46. True or False

DB2 Pre-compiler uses the SQL statements to build a Database Request Module (DBRM) for the program.

True

False

[Ans](#)

47. What does the DBRM contain?

Edited form of the SQL statement with some additional info

Compiled version of the SQL statements

Compiled source code

None of the above

[Ans](#)

48. True or False

The function of the Bind is to convert the high-level DBRMs into an optimised internal form.

True

False

[Ans](#)

49. True or False

Bind performs two major jobs syntax checking and optimisation.

True

False

[Ans](#)

50. What are the various environments in which DB2 operates?

IMS/DB and IMS/DC

CICS

TSO batch and on-line

All of the above

[Ans](#)

51. What is CAF?

Call According to function

Call Attachment Facility

Continuous Application Facility

None of the above

[Ans](#)

SQL PART - I

52. What is the expansion of SQL?

- Sequential Query Language
- Simplified Query Language
- Structured Query Language
- None of the above

[Ans](#)

53. What is SQL2 or SQL/92

- The second version of SQL.
- The second release of SQL
- The ANSI SQL standard introduced in 1992
- None of the above

[Ans](#)

54. True or False

- The primary function of the SQL is to support the definition, manipulation and control of data in relational databases
- True
- False

[Ans](#)

55. What are the different types of SQL statements?

- Data Definition Language
- Data Manipulation Language
- Data Control Language
- All of the above

[Ans](#)

56. True or False

- The Data Definition statements of SQL are GRANT and REVOKE.
- True
- False

[Ans](#)

57. What are the DDL statements of SQL?

- CREATE
- ALTER
- DROP
- All of the above

[Ans](#)

58. What are the Data Control statements of SQL?

- GRANT
- REVOKE

All of the above
None of the above

[Ans](#)

59. True or False

The data manipulation SQL statements are SELECT, INSERT, UPDATE and DELETE.

True

False

[Ans](#)

60. What are the DB2 objects that can be altered using the ALTER statement?

Views

Tables

Indexes

All of the above

[Ans](#)

61. True or False

You can create a base table which has the same structure as some existing table using the CREATE TABLE LIKE..... command

True

False

[Ans](#)

62. True or False

When creating a table using the LIKE command the primary, alternate and foreign key definitions are not inherited.

True

False

[Ans](#)

63. True or False

Alternate key specifications cannot be changed using the ALTER TABLE statement.

True

False

[Ans](#)

64. True or False

ALTER TABLE statement does not support any kind of change to the width or data type of an existing column neither does it support the deletion of an existing column.

True

False

[Ans](#)

65. What are the numeric data types that DB2 supports?

INTEGER, SMALLINT, DECIMAL, FLOAT

INTEGER, DECIMAL, FLOAT

INTEGER, FLOAT

None of the above

[Ans](#)

66. What are the string data types that DB2 supports?

CHARACTER, GRAPHIC

CHARACTER, GRAPHIC, VARCHAR, VARGRAPHIC

CHARACTER, STRING

None of the above

[Ans](#)

67. What are the date/time data types that DB2 supports?

DATE
TIME
TIMESTAMP
All of the above

[Ans](#)

68. If NOT NULL WITH DEFAULT is specified during a CREATE TABLE command what will be the default values for the DATE, TIME and TIMESTAMP?

SPACES
ZEROES
CURRENT DATE, CURRENT TIME and CURRENT TIMESTAMP
None of the above

[Ans](#)

69. If NOT NULL WITH DEFAULT is specified during an ALTER TABLE command what will be the default values for the DATE, TIME and TIMESTAMP?

SPACES
CURRENT DATE, CURRENT TIME and CURRENT TIMESTAMP
01/01/0001, 00:00 AM, 0001-01-01-00.00.00.000000
None of the above

[Ans](#)

70. True or False

The expression DATE('5/31/1977') + 1 MONTH will yield the result '6/30/1977'.
True
False

[Ans](#)

71. True or False

The expression DATE('6/30/1977') - 1 MONTH yields the result '5/31/1977'.
True
False

[Ans](#)

72. True or False

The expression DATE('5/31/1988') + 1MONTH -1MONTH does not yield '5/31/1988', but the expression DATE('5/31/1988')+30 DAYS-30DAYS will yield '5/31/1988'.
True
False

[Ans](#)

73. Which is the function used for converting a date, time, timestamp or decimal number to its character string representation

CHAR.
DATE

DATE TO CHAR

None of the above

[Ans](#)

74. Which is the function used for converting a scalar value to date
CHAR.

DATE

DATE TO CHAR

None of the above

[Ans](#)

75. Which is the function used for extracting the hour portion of the time or timestamp
TIME
HOUR
TIMESTAMP TO HOUR
None of the above

[Ans](#)

76. Which is the function used for extracting the year portion form a date or timestamp
TIME
YEAR
TIMESTAMP TO YEAR
None of the above

[Ans](#)

77. Which are the special registers of DB2
CURRENT SQLID
CURRENT DATE, CURRENT TIME
CURRENT TIMESTAMP
All of the above

[Ans](#)

78. If NOT NULL WITH DEFAULT is specified what will be the default values for the numeric fields?
SPACES
NULLS
ZEROES
None of the above

[Ans](#)

79. If NOT NULL WITH DEFAULT is specified what will be the default values for the character fields?
SPACES
NULLS
ZEROES
None of the above

[Ans](#)

SQL PART II

80. True or False

The 'UNIQUE' option in the CREATE INDEX statement specifies that no two rows in the indexed base table will be allowed to take the same value for the indexed column or column combination at the same time, or in other words no duplicates will be allowed.

True

False

[Ans](#)

81. True or False

If a given column has nulls allowed, a UNIQUE index on that column will permit at the most one null to appear at any given time.

True

False

[Ans](#)

82. True or False

Whenever the base table is dropped the indexes for that table are not automatically dropped

True

False

[Ans](#)

83. True or False

The SELECT * FROM..... will fetch all the columns of the table meeting the search criteria.

True

False

[Ans](#)

84. What is the command that is used to eliminate duplicates in a SELECT statement?

WITH NO DUPLICATES

NO DUPLICATES

DISTINCT

None of the above

[Ans](#)

85. True or False

The default value in a SELECT statement is SELECT DISTINCT.

True

False

[Ans](#)

86. True or False

SQL statements cannot be used for retrieving computed values?

True
False

[Ans](#)

87. What is the clause used for getting the result in a particular order?

ORDER BY
GROUP BY
ORDER
None of the above

[Ans](#)

88. True or False

The column(s) on which the ORDER BY is specified need not be part of the result table

True

False

[Ans](#)

89. True or False

It is possible to identify columns in the ORDER BY clause by column number instead of column name, that is by the ordinal, left-to-right, position of the column in question within the result table.

True

False

[Ans](#)

90. What is the clause used along with the SELECT statement to get those items that fall within a range?

AND

OR

BETWEEN

None of the above

[Ans](#)

91. True or False

```
SELECT PART_NAME
  FROM PARTS_TABLE
 WHERE WEIGHT IN (12, 17, 20)
```

will get the name of all the parts whose weight is 12, 17 or 20.

True

False

[Ans](#)

92. True or False

```
SELECT PNAME
  FROM P
 WHERE PNAME LIKE 'C%'
```

will get the name of all the parts whose name starts with letter C.

True

False

[Ans](#)

93. True or False

```
SELECT PNAME
  FROM P
 WHERE PNAME LIKE 'C_'
```

will get the name of all the parts whose name starts with letter C.

True

False

[Ans](#)

94. True or False

SELECT PNAME

FROM P

WHERE PNAME LIKE '%C%'

will get the name of all the parts whose name has a character C in it.

True

False

[Ans](#)

95. Is the statement

SELECT *

FROM S

WHERE STATUS = NULL; correct

Yes

No

Can't say

[Ans](#)

96. True or False

If the WHERE clause in the UPDATE or DELETE includes a subquery, then the FROM clause of that subquery must not refer to the table that is being updated or deleted.

True

False

[Ans](#)

97. True or False

In the subquery that is used in the INSERT statement the FROM clause can refer to the table that is the target of the INSERT.

True

False

[Ans](#)

98. What are the aggregate functions supported by DB2?

SUM and AVG

SUM, MAX and MIN

COUNT, SUM, AVG, MAX, MIN

None of the above

[Ans](#)

99. What is the function which is used to all rows without any duplicate elimination?

COUNT

COUNT ALL

COUNT (&)

COUNT (*)

[Ans](#)

100. True or False

The key word DISTINCT is not allowed for COUNT(*) function.

True

False

[Ans](#)

101. True or False

The SQL ' SELECT AVG(MIN(QTY)) AS AVERAGE' is legal.

True

False

[Ans](#)

102. Which is the aggregate function that returns zero when there are no rows satisfying the search criteria?

SUM

COUNT

AVG

None of the above

[Ans](#)

103. What will the aggregate function SUM return if there are no rows satisfying the search criteria?

ZERO

-1

NULL

1

[Ans](#)

104. True or False

If HAVING is specified, GROUP BY clause also must be specified and if GROUP BY clause is omitted then the entire table will be considered as a single group.

True

False

[Ans](#)

105. True or False

SELECT name, salary FROM employee-table

WHERE EXISTS (SELECT * FROM department-table

WHERE dept LIKE 'C%')

The above SQL will fetch the names and salaries of all the employees whose department name starts with 'C'.

True

False

[Ans](#)

106. True or False

SELECT name FROM employee-table

WHERE NOT EXISTS

[SELECT * FROM department

WHERE COUNT(*) >25]

The above SQL the sub-query will get all the departments whose strength is more than 25 and the NOT EXISTS clause will get the employee names who does not belong to the those departments

True

False

[Ans](#)

SQL PART III

107. True or False

Subqueries are nested SELECT statements.

True

False

[Ans](#)

108. What will happen if the subquery in the following SQL returns more than one value?

```
SELECT name, salary
```

```
FROM employee-table
```

```
WHERE deptno = (SELECT deptno FROM department-table WHERE dept  
='C%')
```

Query will fail

Query will run

Query will be ignored

Nothing will happen

[Ans](#)

109. True or False

A non-correlated subquery is evaluated in a top-to-bottom fashion and a correlated subquery is resolved in a bottom-to-top manner.

True

False

[Ans](#)

110. If the comparison operator in a join is equality then that join is called

Outer Join

Equality Join

Equijoin

None of the above

[Ans](#)

111. The equijoin must produce a result containing two identical columns. If one of these columns is eliminated then that join is called

Equijoin

Natural join

Outer join

None of the above

[Ans](#)

112. True or False

More than 2 tables cannot be joined.

True

False

[Ans](#)

113. True or False

The maximum number of tables that can be joined using an SQL in DB2 is 15.

True

False

[Ans](#)

114. True or false

The union operation combines two sets of rows into a single set composed of all the rows in either or both of the two original sets whereas a join combines two sets of columns into a single set.

True

False

[Ans](#)

115. What are the conditions of two tables to be union compatibility:

The 2 tables must have the same number of columns .

Each col of the 1st set must be same data type as corresponding col of the 2nd set

Each col of the 1st set must be convertible to same data type as corresponding col of the 2nd

All of the above

[Ans](#)

116. True or False

The UNION ALL verb eliminates duplicates whereas the UNION verb retains them.

True

False

[Ans](#)

117. True or False

Any ORDER BY clause in the query must appear as part of the final SELECT only and must identify ordering columns by their ordinal position or number and not by name.

True

False

[Ans](#)

118. Does DB2 support outer joins?

Yes

No

Can't say

[Ans](#)

119. Which are the situations in which column names must be specified explicitly for all columns of the view?

Any column of the view is derived from a function

Any column of the view is derived from an operational expression

Two or more columns of the view would otherwise have the same name

All of the above

[Ans](#)

120. True or False

For a view to be updatable it must be derived from a single base-table.

True

False

[Ans](#)

121. True or False

If a column of the view is derived from an expression involving a scalar operator or a scalar function or a literal, then INSERT and UPDATE operations are allowed, but DELETE operations are not allowed.

True

False

[Ans](#)

122. True or False

If a column of a view is derived from an aggregate function, then the view is not updatable.

True

False

[Ans](#)

123. True or False

If the definition of the view involves either a GROUP BY or a HAVING clause at the outermost level, then the view is updatable.

True

False

[Ans](#)

124. True or False

If the definition of the view involves DISTINCT at the outermost level, then that view is not updatable.

True

False

[Ans](#)

125. If the definition of the view includes a nested subquery that refers to the base table on which the view is defined, then it is not updatable.

True

False

[Ans](#)

126. True or False

If the FROM clause in the view definition involves multiple range variables, then it is not updatable.

True

False

[Ans](#)

127. A view defined on a non-updatable view is updatable.

True

False

[Ans](#)

128. Fill in the blanks to make the SQL complete and to ensure that all the inserts and updates made on the view can specify values 'M' or 'F' only:

```
CREATE VIEW emp (emp-no, name, sex, age, dept, salary)
AS   SELECT emp-no, name, sex, age, dept, salary
      FROM emp-table WHERE sex IN ('M', 'F')
      WITH ----- OPTION;
EXAMINE
CHECK
ISSUE
```

None of the above

[Ans](#)

SQL - TIPS, TRICKS AND TOOLS

129. True or False

EXPLAIN facility provided by DB2 is used to get knowledge into the performance potential of each SQL in an application.

True

False

[Ans](#)

130. Name the DB2 table where the EXPLAIN analysis results can be got.

PLAN_TABLE

EXPLAIN_TABLE

PACKAGE_TABLE

None of the above

[Ans](#)

131. What will you use when you want to retrieve more than one row in a program?

Singleton SELECT

SELECT

Cursor

None of the above

[Ans](#)

132. What is the error handling SQL statement?

SELECT

INSERT

UPDATE

WHENEVER

[Ans](#)

133. What are the statements used for using a Cursor?

DECLARE

OPEN and FETCH

CLOSE

All of the above

[Ans](#)

134. When you issue a SELECT and now matching rows are found which is the SQL code that is returned?

-911

0

-100

100

[Ans](#)

135. True or False

The BETWEEN predicate is more efficient than the 'greater/less than or equal to' predicates.

True

False

[Ans](#)

136. True or False

When you are coding a subquery using negation logic, use NOT EXISTS instead of NOT IN to increase the efficiency of the SQL.

True

False

[Ans](#)

DATA SECURITY AND ACCESS

137. What are the Data control statements of SQL?

- GRANT and REVOKE
- COMMIT and ROLLBACK
- GETMAIN and FREEMAIN
- None of the above

[Ans](#)

138. True or False

SYSADM or System Administration authority allows the holder to execute any operation that the system supports.

- True
- False

[Ans](#)

139. True or False

CURRENT SQLID is a special register like CURRENT TIMESTAMP whose value is an authorisation ID.

- True
- False

[Ans](#)

140. True or False

DBADM or Database Administration authority on a specific database allows the holder to execute any operation that the system supports on that database.

- True
- False

[Ans](#)

141. True or False

Database Control authority (DBCTRL) on specific database allows the holder to execute any operation that the system supports on that database (like RECOVER DATABASE) except for operations that access the data content of that database.

- True
- False

[Ans](#)

142. True or False

Database Maintenance authority on a specific database allows the holder to execute read-only maintenance functions such as 'IMAGE COPY' on that database.

- True
- False

[Ans](#)

143. True or False

System Operator authority allows the holder to carry out console operator functions on the system such as starting and stopping system trace activities.

True

False

[Ans](#)

- 144.** What are the table privileges?
ALTER, DELETE
INDEX, INSERT, SELECT, UPDATE
ALL
All of the above

[Ans](#)

- 145.** What are the plan privileges?
ALTER, DELETE
ALL
BIND, EXECUTE
None of the above

[Ans](#)

- 146.** What are the database privileges?
CREATETAB, CREATETS, DISPLAYDB
DROP, IMAGECOPY, STARTDB
LOAD, RECOVERDB, REORG
All of the above

[Ans](#)

- 147.** What are the use privileges?
USE OF BUFFERPOOL
USE OF STOGROUP
USE OF TABLESPACE
All of the above

[Ans](#)

- 148.** True or False
DB2 is a hierarchical database
True
False

[Ans](#)

- 149.** True or False
CICS and DB2 can exist in the same region under the Operating system
True
False

[Ans](#)

- 150.** What is the facility that connects CICS and DB2?
SQL
VSAM
CICS attachment Facility
DB2 Attachment Facility

[Ans](#)

151. What is the syntax of the DB2 access commands?

EXEC DLI function
EXEC CICS function
EXEC DB2 function
EXEC SQL function

[Ans](#)

152. True or False

The application programs that contain the SQL statements must be Pre-compiled for converting the SQL statements into equivalent COBOL statements

True

False

[Ans](#)

INTEGRITY

153. What are the different types of data integrity?

- Entity integrity
- Domain integrity
- Referential integrity
- All of the above

[Ans](#)

154. True or False

Domain integrity is making sure that every field value is a member of a domain.

- True
- False

[Ans](#)

155. True or False

Entity integrity makes sure that each row is unique.

- True
- False

[Ans](#)

156. True or False

The value of the primary key can be a null.

- True
- False

[Ans](#)

157. True or False

A table can have more than one primary key.

- True
- False

[Ans](#)

158. True or False

A table can have more than one unique key.

- True
- False

[Ans](#)

159. True or False

Referential Integrity ensures that the related tables stay in sync.

- True
- False

[Ans](#)

160. Primary keys

Are unique identifiers of a table

Cannot contain nulls
Not more than one for a table
All of the above

[Ans](#)

- 161.** When the primary key is a combination of multiple columns it is called
Complex key
Composite key
Alternate key
None of the above

[Ans](#)

- 162.** True or False
Every column participating in the primary key must be explicitly declared NOT NULL
True
False

[Ans](#)

- 163.** True or False
When a primary key is dropped, the primary index will also be dropped.
True
False

[Ans](#)

- 164.** A column or combination of columns in one table whose values are required to match with the values of the primary key in some other table is called
Referential key
Foreign key
Alternate key
None of the above

[Ans](#)

- 165.** True or False
The ON DELETE CASCADE option means that the delete operation is cascaded, that is if a row in the parent table is deleted then all the matching rows in the child table is automatically deleted.
True
False

[Ans](#)

DB2 APPLICATION DEVELOPMENT

166. What are the host languages supported by DB2?

APL2, Assembler H, IBM BASIC
OS/VS COBOL, VS COBOL II
C/370, FORTRAN, LISP, PL/
All of the above

[Ans](#)

167. What are the delimiters for embedded SQLs in COBOL?

EXEC and END EXEC
EXEC SQL and END-EXEC
SQL and END-SQL
None of the above

[Ans](#)

168. True or False

The SQLCA contains fields which are used for communicating information describing the success or failure of the execution of an embedded SQL.

True
False

[Ans](#)

169. True or False

A host variable is an area of storage allocated by the host language and referenced in the SQL statement.

True
False

[Ans](#)

170. What is the DB2 utility used to generate host variables ?

QMF
SPUFI
DCLGEN
None of the above

[Ans](#)

171. When there are no more rows to be fetched in a cursor DB2 returns an SQL code of

- 911
100
-912
None of the above

[Ans](#)

172. True or False

If you want to read a row and depending upon the values in the row, you want to modify, delete or do nothing, you can do that with a cursor, this is accomplished with a cursor and a special clause of UPDATE and DELETE statements usable only by embedded SQLs, namely WHERE CURRENT OF.

True

False

[Ans](#)

- 173.** When WHERE CURRENT OF clause is used the cursor should be declared with.
UPDATE clause
FOR UPDATE OF clause
MODIFY clause
None of the above

[Ans](#)

- 174.** What are the two types of SQLs?
Simple and complex
Static and dynamic
re-entrant and non-re-entrant
None of the above

[Ans](#)

- 175.** True or False
Dynamic SQLs are characterised by the capability to change columns, tables and predicates during a program's execution.
True
False

[Ans](#)

- 176.** What are the dynamic SQL statements?
DECLARE
EXECUTE
PREPARE and EXECUTE IMMEDIATE
All of the above

[Ans](#)

- 177.** True or False
EXECUTE IMMEDIATE dynamic SQL supports SELECT statement.
True
False

[Ans](#)

- 178.** What is the parameter marker in DB2?
*
?
/
None of the above

[Ans](#)

- 179.** True or False
There are two types of BINDs - BIND PLAN and BIND PACKAGE.
True
False

[Ans](#)

180. What is the input for the Bind operation?

Source code

SQLs

DBRMs

None of the above

[Ans](#)

181. What is the output of the Bind operation?

- Load module
- Application Plan
- DBRMs
- None of the above

[Ans](#)

182. True or False

The output of the BIND PLAN is an application plan containing the executable logic representing optimised access paths to DB2 data.

- True
- False

[Ans](#)

183. True or False

A package is a single, bound DBRM with optimised access paths.

- True
- False

[Ans](#)

184. True or False

A collection is a user defined name for every package

- True
- False

[Ans](#)

185. What is the TSO parameter which controls the number of users that can access DB2 simultaneously from the TSO foreground

- IDFOR
- IDNO
- IDFORE
- None of the above

[Ans](#)

186. What is the TSO parameter controls the number of concurrent DB2 batch connections

- IDFORE
- IDNO
- IDBACK
- None of the above

[Ans](#)

187. True or False

DB2 batch programs are executed in the background under the control of the TSO Terminal Monitor Program, IKJEFT01.

- True
- False

[Ans](#)

- 188.** What is the expansion of SPUFI?
- Sequential Programming Using Functional Input
 - SQL Processor Using File Input
 - SQL Program Using Fast Interaction
 - None of the above

[Ans](#)

189. What is the expansion of QMF?

Question Management Function

Query Management Function

Query Management Facility

None of the above

[Ans](#)

190. True or False

A mechanism called CICS Attachment Facility connects CICS with DB2.

True

False

[Ans](#)

191. True or False

The SQL commands that can be used in a CICS application program are SELECT, FETCH, UPDATE, DELETE and INSERT.

True

False

[Ans](#)

DB2 INTERNALS

192. True or False

The total collection of stored data is divided into a number of user databases and system databases, each of which is divided into a number of tablespaces and index spaces.

True

False

[Ans](#)

193. True or False

A 'space' is a dynamically extendible collection of pages, where a 'page' is a block of physical storage.

True

False

[Ans](#)

194. What is the page size of an index space?

32KB

4KB

16KB

32KB and 4KB

[Ans](#)

195. What is the page size of a table space?

32KB

4KB

16KB

32KB and 4KB

[Ans](#)

196. True or False

A table space can be considered as a logical address space on secondary storage that is used to hold one or more stored tables.

True

False

[Ans](#)

197. True or False

The table space is the storage unit for recovery and reorganisation purposes, that is, it is the unit that can be recovered via the RECOVER utility or reorganised using the REORG utility.

True

False

[Ans](#)

198. What are the different types of table spaces?

Simple
Partitioned
Segmented
All of the above

[Ans](#)

199. True or False

Partitioned table spaces are intended for tables that are sufficiently large that is operationally difficult to deal with the entire table as a single unit.

True
False

[Ans](#)

200. True or False

The DDL for creating an index space is CREATE INDEXSPACE.

True
False

[Ans](#)

201. True or False

Indexes in DB2 are based on a structure known as B-tree.

True
False

[Ans](#)

202. What is the name of the DB2 subsystem which analyses the SQL statements and determines the most efficient access path available for satisfying the statement?

EXPLAIN
RUNSTATS
OPTIMISER
None of the above

[Ans](#)

203. True or False

The DB2 Catalog is a single database composed of 11 tablespaces and 43 tables.

True
False

[Ans](#)

204. What is the name of the DB2 catalog database?

DSNDB
DSNDB01
DSNDB06
None of the above

[Ans](#)

205. What is the name of the DB2 utility that is used for making the DB2 catalog up-to-date?

REORG.
REPAIR
RUNSTATS
None of the above

[Ans](#)

206. True or False

The DB2 Directory is for DB2's internal use only and the data in the DB2 directory cannot be accessed by the SQL statements.

True

False

[Ans](#)

207. True or False
DB2 has record level locking.
True
False

[Ans](#)

208. What is the system that DB2 uses for locking?
Lock Manager
Resource Manager
IRLM
None of the above

[Ans](#)

209. When DB2 tries to do locking without the help of IRLM that locks are called
Locks
Latches
Links
None of the above

[Ans](#)

210. True or False
Latches are more efficient than locks.
True
False

[Ans](#)

211. What is the situation in which two or more transactions are in simultaneous wait state, each waiting for one of the others to release a lock before it can proceed called?
Time-out
Deadlock
Rollback
None of the above

[Ans](#)

PERFORMANCE MONITORING

212. What are the different types of traces provided by DB2?

- Accounting Trace and Audit Trace
- Global Trace and Monitor Trace
- Performance Trace and Statistics Trace
- All of the above

[Ans](#)

213. What are the different destinations to which DB2 traces are written ?

- GTF and RES
- SMF and SRV
- OPn and OPX
- All of the above

[Ans](#)

214. Which is the most widely used batch performance monitor for DB2.

- DB2I
- EXPLAIN
- DB2PM
- None of the above

[Ans](#)

215. True or False

EXPLAIN allows the user to obtain information regarding the optimiser's choice of access strategy for a specified SQL statement.

- True
- False

[Ans](#)

216. What is the default destination for the Statistics trace?

- GTF
- RES
- SMF
- None of the above

[Ans](#)

217. What is the default destination for the Accounting trace?

- GTF
- RES
- SMF
- None of the above

[Ans](#)

218. What is the default destination for the Audit trace?

- GTF
- RES

SMF
None of the above

[Ans](#)

219. What is the default destination for the Performance trace?

GTF

RES

SMF

None of the above

[Ans](#)

220. What is the default destination for the Monitor trace?

GTF

RES

OPX

SMF

[Ans](#)

221. What is the default destination for the Global trace?

GTF

RES

SMF

None of the above

[Ans](#)

DB2 UTILITIES AND COMMANDS

222. True or False

There are three data consistency utilities CHECK, REPAIR and REPORT.

True

False

[Ans](#)

223. What are the functions of the CHECK utility?

Checks the integrity of data structures

Checks the referential integrity between 2 tables

Checks the indexes for consistency

All of the above

[Ans](#)

224. True or False

The CHECK utility has two options CHECK DATA and CHECK INDEX .

True

False

[Ans](#)

225. True or False

The REPAIR utility is designed to modify DB2 data and associated data structures when there is an error or problem.

True

False

[Ans](#)

226. What are the functions of the REPAIR utility?

Synchronise DB2 catalog info with DB2 directory & DBD definition

Physically change specific locations in a dataset

Reset the pending flags that are erroneously set

All of the above

[Ans](#)

227. True or False

The REPORT TABLESPACESET option generates a report detailing all tables and tablespaces in a referential tablespace set.

True

False

[Ans](#)

228. True or False

The input to the utility is a single tablespace and the output is a report of all related tablespaces and tables.

True

False

[Ans](#)

229. True or False

The REPORT RECOVERY can be used to generate a report on tablespace recovery information. The report contains information from the DB2 directory, the DB2 catalog, and the BSDS.

True

False

[Ans](#)

230. What is the input to the REPORT RECOVERY utility?

A tablespace

single partition of a partitioned tablespace

All of the above

[Ans](#)

231. What are the main backup utilities of DB2?

COPY and MERGE COPY

QUIESEC

RECOVER

All of the above

[Ans](#)

232. True or False

The COPY utility is used to create an image copy back up data set for a complex tablespace or a single partition of a tablespace.

True

False

[Ans](#)

233. True or False

The MERGECOPY utility combines multiple incremental image copy data sets into a new full or incremental image copy data set.

True

False

[Ans](#)

234. True or False

QUIESEC utility ensures that all tablespaces in the scope of QUIESEC are referentially intact.

True

False

[Ans](#)

235. What are the data organisation utilities of DB2?

LOAD and REORG

LOAD and RECOVER

REORG and REPAIR

None of the above

[Ans](#)

236. True or False

The LOAD utility is used to accomplish bulk inserts to DB2 tables.

True

False

[Ans](#)

237. True or False

The catalog manipulation utilities of DB2 are CATMAINT, MODIFY, RUNSTATS and STOSPACE.

True

False

[Ans](#)

238. True or False

The MODIFY utility is used to delete rows from the DB2 catalog and directory tables.

True

False

[Ans](#)

239. True or False

The RUNSTATS utility collects statistical information for DB2 tables, tablespaces, partitions, indexes, and columns.

True

False

[Ans](#)

DB2 RECOVERY AND RESTART SCENARIOS

240. True or False

The COMMIT and ROLLBACK statements are instruction to the DBMS.

True

False

[Ans](#)

241. True or False

A syncpoint represents a boundary point between two transactions and it corresponds to the end of a logical unit of work and therefore a point in which the database is in a state of consistency.

True

False

[Ans](#)

242. True or False

COMMIT and COMMIT WORK does the same function.

True

False

[Ans](#)

243. True or False

The DB2 log registers data changes and significant events as they occur into a DASD dataset called the active log.

True

False

[Ans](#)

244. What is the process of copying the contents of the active log to archive log called?

Load-shedding

Off-loading

Archiving

None of the above

[Ans](#)

ANSWERS

- | | |
|--|-------------------|
| 1. Database | Q |
| 2. International Business Machines Corporation | Q |
| 3. Relational Database Management System | Q |
| 4. True | Q |
| 5. All of the above | Q |
| 6. True | Q |
| 7. Tuples | Q |
| 8. Attributes | Q |
| 9. Domains | Q |
| 10. True | Q |
| 11. True | Q |
| 12. All of the above | Q |
| 13. IMS | Q |
| 14. IDMS | Q |
| 15. E.F. Codd | Q |
| 16. E.F. Codd | Q |
| 17. Answer is not available | Q |
| 18. Answer is not available | Q |
| 19. DB2 | Q |
| 20. True | Q |
| 21. Degree | Q |
| 22. Cardinality | Q |
| 23. True | Q |

24. False	Q	
25. False	Q	
26. Forgein Key		Q
27. True	Q	
28. Referential Integrity		Q
29. Composite Key	Q	
30. True	Q	
31. MVS	Q	
32. False	Q	
33. DB2I	Q	
34. True	Q	
35. True	Q	
36. Base Tables and Views	Q	
37. True	Q	
38. True	Q	
39. Logging	Q	
40. True	Q	
41. All of the above	Q	
42. True	Q	
43. False	Q	
44. True	Q	
45. Replaces all the SQL statements with host language CALL statements	Q	
46. True	Q	
47. Compiled version of the SQL statements	Q	
48. True	Q	

49. True	Q	
50. All of the above	Q	
51. Call attachment facility	Q	
52. Structured Query Language		Q
53. The ANSI SQL standard introduced in 1992		Q
54. True	Q	
55. All of the above	Q	
56. False	Q	
57. All of the above	Q	
58. All of the above	Q	
59. True	Q	
60. Tables	Q	
61. True	Q	
62. True	Q	
63. True	Q	
64. True	Q	
65. INTEGER, SMALLINT, DECIMAL, FLOAT		Q
66. CHARACTER, GRAPHIC, VARCHAR, VARGRAPHIC		Q
67. All of the above	Q	
68. CURRENT DATE, CURRENT TIME and CURRENT TIMESTAMP	Q	
69. 01/01/0001, 00:00 AM, 0001-01-01-00.00.00.000000		Q
70. True	Q	
71. False	Q	
72. True	Q	
73. CHAR	Q	

74. DATE	Q	
75. HOUR	Q	
76. YEAR	Q	
77. All of the above	Q	
78. ZEROES	Q	
79. SPACES	Q	
80. True	Q	
81. True	Q	
82. False	Q	
83. True	Q	
84. DISTINCT		Q
85. False	Q	
86. True	Q	
87. ORDER BY		Q
88. False	Q	
89. True	Q	
90. BETWEEN		Q
91. True	Q	
92. True	Q	
93. Flase	Q	
94. True	Q	
95. No	Q	
96. True	Q	
97. True	Q	
98. COUNT, SUM, AVG, MAX, MIN		Q

99.	COUNT(*)		Q
100.	True	Q	
101.	False	Q	
102.	COUNT		Q
103.	NULL	Q	
104.	True	Q	
105.	True	Q	
106.	False	Q	
107.	True	Q	
108.	Query will fail	Q	
109.	False	Q	
110.	Equijoin		Q
111.	Natural Join		Q
112.	False	Q	
113.	True	Q	
114.	True	Q	
115.	All of the above		Q
116.	True	Q	
117.	False	Q	
118.	Yes		Q
119.	All of the above		Q
120.	True	Q	
121.	False	Q	
122.	True	Q	
123.	False	Q	

124.	True	Q	
125.	True	Q	
126.	True	Q	
127.	False	Q	
128.	CHECK		Q
129.	True	Q	
130.	PLAN_TABLE		Q
131.	Cursor	Q	
132.	WHENEVER	Q	
133.	All of the above		Q
134.	+100	Q	
135.	True	Q	
136.	True	Q	
137.	GRANT and REVOKE		Q
138.	True	Q	
139.	True	Q	
140.	True	Q	
141.	True	Q	
142.	True	Q	
143.	True	Q	
144.	All of the above		Q
145.	BIND, EXECUTE		Q
146.	CREATETB, CREATETS, DISPLAYDB		Q
147.	USE OF TABLESPACE		Q
148.	False	Q	

149.	False	Q	
150.	CICS Attachment Facility		Q
151.	EXEC SQL function		Q
152.	False	Q	
153.	All of the above		Q
154.	True	Q	
155.	True	Q	
156.	False	Q	
157.	False	Q	
158.	False	Q	
159.	True	Q	
160.	All of the above		Q
161.	Composite key		Q
162.	True	Q	
163.	True	Q	
164.	Foreign key		Q
165.	True	Q	
166.	All of the above		Q
167.	EXEC-SQL and END-EXEC		Q
168.	True	Q	
169.	True	Q	
170.	DCLGEN		Q
171.	+100	Q	
172.	True	Q	
173.	FOR UPDATE OF clause		Q

174.	Static and Dynamic		Q
175.	True	Q	
176.	All of the above		Q
177.	False	Q	
178.	?		Q
179.	True	Q	
180.	DBRMs		Q
181.	Application Plan		Q
182.	True	Q	
183.	True	Q	
184.	True	Q	
185.	<i>Answer not available</i>		Q
186.	<i>Answer not available</i>		Q
187.	True	Q	
188.	SQL Processor Using File Input		Q
189.	Query Management Facility		Q
190.	True	Q	
191.	True	Q	
192.	True	Q	
193.	True	Q	
194.	4KB	Q	
195.	32 KB and 4KB		Q
196.	True	Q	
197.	True	Q	
198.	All of the above		Q

199.	True	Q	
200.	True	Q	
201.	True	Q	
202.	EXPLAIN		Q
203.	True	Q	
204.	<i>Answer not available</i>		Q
205.	RUNSTATS		Q
206.	True	Q	
207.	True	Q	
208.	IRLM	Q	
209.	Latches		Q
210.	True	Q	
211.	Deadlock		Q
212.	All of the above		Q
213.	All of the above		Q
214.	DB2PM		Q
215.	True	Q	
216.	SMF	Q	
217.	SMF	Q	
218.	SMF	Q	
219.	SMF	Q	
220.	OPX	Q	
221.	RES	Q	
222.	True	Q	
223.	All of the above		Q

224.	True	Q	
225.	True	Q	
226.	All of the above		Q
227.	True	Q	
228.	True	Q	
229.	True	Q	
230.	All of the above		Q
231.	COPY and MERGE COPY		Q
232.	True	Q	
233.	True	Q	
234.	True	Q	
235.	LOAD and REORG		Q
236.	True	Q	
237.	True	Q	
238.	True	Q	
239.	True	Q	
240.	True	Q	
241.	True	Q	
242.	<i>Answer is not available</i>		Q
243.	<i>Answer is not available</i>		Q
244.	<i>Answer is not available</i>		Q

1. How would you find out the total number of rows in a table? -

2. How do you eliminate duplicate values in SELECT? -
3. How do you select a row using indexes? -
4. What are aggregate functions?
5. How do you find the maximum value in a column? -
6. Can you use MAX on a CHAR column?
7. My SQL statement `SELECT AVG(SALARY) FROM EMP` yields inaccurate results. Why?
8. How do you retrieve the first 5 characters of FIRSTNAME column of EMP table?
9. How do you concatenate the FIRSTNAME and LASTNAME from EMP table to give a complete name?
10. What is the use of VALUE function?
11. What is UNION, UNION ALL? -
12. Suppose I have five SQL SELECT statements connected by UNION/UNION ALL, how many times should I specify UNION to eliminate the duplicate rows? -
13. What is the restriction on using UNION in embedded SQL?
14. In the WHERE clause what is BETWEEN and IN? -
15. Is BETWEEN inclusive of the range values specified? -
16. What is 'LIKE' used for in WHERE clause? What are the wildcard characters? -
17. When do you use a LIKE statement?
18. What is the meaning of underscore ('_') in the LIKE statement? -
19. What do you accomplish by GROUP BY ... HAVING clause? -
20. Consider the employee table with column PROJECT nullable. How can you get a list of employees who are not assigned to any project?

21. What is the result of this query if no rows are selected:
22. Why SELECT * is not preferred in embedded SQL programs?
23. What are correlated subqueries? -
24. What are the issues related with correlated subqueries? -
25. What is a cursor? why should it be used? -
26. How would you retrieve rows from a DB2 table in embedded SQL? -
27. Apart from cursor, what other ways are available to you to retrieve a row from a table in embedded SQL? -
28. Where would you specify the DECLARE CURSOR statement? -
29. How do you specify and use a cursor in a COBOL program? -
30. What happens when you say OPEN CURSOR?
31. Is DECLARE CURSOR executable?
32. Can you have more than one cursor open at any one time in a program ? -
33. When you COMMIT, is the cursor closed?
34. How do you leave the cursor open after issuing a COMMIT? (for DB2 2.3 or above only)
35. Give the COBOL definition of a VARCHAR field.
36. What is the physical storage length of each of the following DB2 data types:
DATE, TIME, TIMESTAMP?
37. What is the COBOL picture clause of the following DB2 data types:
DATE, TIME, TIMESTAMP?
38. What is the COBOL picture clause for a DB2 column defined as
DECIMAL(11,2)? -
39. What is DCLGEN ? -

40. What are the contents of a DCLGEN? -
41. Is it mandatory to use DCLGEN? If not, why would you use it at all? -
42. Is DECLARE TABLE in DCLGEN necessary? Why it used?
43. Will precompile of an DB2-COBOL program bomb, if DB2 is down?
44. How is a typical DB2 batch pgm executed ?
45. Assuming that a site's standard is that pgm name = plan name, what is the easiest way to find out which pgms are affected by change in a table's structure ?
46. Name some fields from SQLCA.
47. How can you quickly find out the # of rows updated after an update statement?
48. What is EXPLAIN? -
49. What do you need to do before you do EXPLAIN?
50. Where is the output of EXPLAIN stored? -
51. EXPLAIN has output with MATCHCOLS = 0. What does it mean? -
52. How do you do the EXPLAIN of a dynamic SQL statement?
53. How do you simulate the EXPLAIN of an embedded SQL statement in SPUFI/QMF? Give an example with a host variable in WHERE clause.)
54. What are the isolation levels possible ? -
55. What is the difference between CS and RR isolation levels?
56. Where do you specify them ?
57. When do you specify the isolation level? How?
58. I use CS and update a page. Will the lock be released after I am done with that page?
59. What are the various locking levels available?
60. How does DB2 determine what lock-size to use?
61. What are the disadvantages of PAGE level lock?
62. What is lock escalation?

63. What are the various locks available?
64. Can I use LOCK TABLE on a view?
65. What is ALTER ? -
66. What is a DBRM, PLAN ?
67. What is ACQUIRE/RELEASE in BIND?
68. What else is there in the PLAN apart from the access path? -
69. What happens to the PLAN if index used by it is dropped?
70. What are PACKAGES ? -
71. What are the advantages of using a PACKAGE?
72. What is a collection?
73. In SPUFI suppose you want to select max. of 1000 rows , but the select returns only 200 rows. What are the 2 sqlcodes that are returned? –
74. How would you print the output of an SQL statement from SPUFI? -
75. How do you pull up a query which was previously saved in QMF ? -
76. Lot of updates have been done on a table due to which indexes have gone haywire. What do you do? -
77. What is dynamic SQL? -
78. When is the access path determined for dynamic SQL? -
79. Suppose I have a program which uses a dynamic SQL and it has been performing well till now. Off late, I find that the performance has deteriorated. What happened? -
80. How does DB2 store NULL physically?
81. How do you retrieve the data from a nullable column? -
82. What is the picture clause of the null indicator variable? -
83. What does it mean if the null indicator has -1, 0, -2? -
84. How do you insert a record with a nullable column?

85. What is RUNSTATS? -
86. When will you chose to run RUNSTATS?
87. Give some example of statistics collected during RUNSTATS?
88. What is REORG? When is it used?
89. What is IMAGECOPY ? -
90. When do you use the IMAGECOPY? -
91. What is COPY PENDING status?
92. What is CHECK PENDING ?
93. What is QUIESCE?
94. What is a clustering index ? -
95. How many clustering indexes can be defined for a table?
96. What is the difference between primary key & unique index ?
97. What is sqlcode -922 ?
98. What is sqlcode -811?
99. What does the sqlcode of -818 pertain to? -
100. Are views updatable ?
101. If I have a view which is a join of two or more tables, can this view be updatable?
-
102. What are the 4 environments which can access DB2 ?
103. What is an inner join, and an outer join ?
104. What is FREEPAGE and PCTFREE in TABLESPACE creation?
105. What are simple, segmented and partitioned table spaces ?
106. What is filter factor?
107. What is index cardinality? -
108. What is a synonym ?

109. What is the difference between SYNONYM and ALIAS?
110. What do you mean by NOT NULL WITH DEFAULT? When will you use it?
111. What do you mean by NOT NULL? When will you use it?
112. When would you prefer to use VARCHAR?
113. What are the disadvantages of using VARCHAR?
114. How do I create a table MANAGER (EMP#, MANAGER) where MANAGER is a foreign key which references to EMP# in the same table? Give the exact DDL.
115. When is the authorization check on DB2 objects done - at BIND time or run time?
116. What is auditing?
117. What is a DB2 bind?
118. What is a DB2 access path?
119. What is a DB2 plan?
120. What is normalization and what are the five normal forms?
121. What are foreign keys?
122. Describe the elements of the SELECT query syntax.
123. Explain the use of the WHERE clause.
124. What techniques are used to retrieve data from more than one table in a single SQL statement?
125. What do the initials DDL and DML stand for and what is their meaning?
126. What is a view? Why use it?
127. Explain an outer join.
128. What is a subselect? Is it different from a nested select?
129. What is the difference between group by and order by?
130. Explain the EXPLAIN statement.
131. What is tablespace?

- 132.What is a cursor and what is its function?
- 133.What is referential integrity?
- 134.Usually, which is more important for DB2 system performance - CPU processing or I/O access?
- 135.Is there any advantage to denormalizing DB2 tables?
- 136.What is the database descriptor?
- 137.What is lock contention?
- 138.What is SPUFI?
- 139.What is the significance of DB2 free space and what parameters control it?
- 140.what is a NULL value? What are the pros and cons of using NULLS?
- 141.What is a synonym? How is it used?
- 142.What is an alias and how does it differ from a synonym?
- 143.What is a LIKE table and how is it created?
- 144.If the base table underlying a view is restructured, eg. attributes are added, does the application code accessing the view need to be redone?
- 145.Under what circumstances will DB2 allow an SQL statement to update more than one primary key value at a time?
- 146.What is the cascade rule and how does it relate to deletions made with a subselect.
- 147.What is the self-referencing constraint?
- 148.What are delete-connected tables?
- 149.When can an insert of a new primary key value threaten referential integrity?
- 150.In terms of DB2 indexing, what is the root page?
- 151.How does Db2 use multiple table indexes?
- 152.What are some characteristics of columns that benefit from indexes?
- 153.What is a composite index and how does it differ from a multiple index?
- 154.What is meant by index cardinality?

155. What is a clustered index?
156. What keyword does an SQL SELECT statement use for a string search?
157. What are some sql aggregates and other built-in functions?
158. How is the SUBSTR keyword used in sql?
159. What are the three DB2 date and time data types and their associated functions?
160. Explain transactions, commits and rollbacks in DB2.
161. What is deadlock?
162. What are the four lockable units for DB2?
163. What are the three lock types?
164. What is isolation level?
165. What is an intent lock?
166. What is the difference between static and dynamic sql?
167. What is cursor stability?
168. What is the significance of the CURSOR WITH HOLD clause in a cursor declaration?
169. What is the SQL Communications Area and what are some of its key fields?
170. What is the purpose of the WHENEVER statement?
171. What is DCLGEN?
172. What is the FREE command?
173. DB2 can implement a join in three ways using a merge join, a nested join or a hybrid join. Explain the differences.
174. Compare a subselect to a join.
175. What is the difference between IN subselects and EXISTS subselect?
176. What is a Cartesian product?
177. DB2 What is the difference between a package and a plan? How does one bind 2 versions of a CICS transaction with the same module name in two different CICS regions that share the same DB2 subsystem?

178.What is an asychronous write?

179.What is a lock?

180.What is meant by isolation level?

181.What are leaf pages?

182.What is a precompiler?

183.What is a root page?

184.What is a thread?

DB2 Answers.

- A1 Use SELECT COUNT(*) ...
- A2 Use SELECT DISTINCT ...
- A3 Specify the indexed columns in the WHERE clause.
- A4 Bulit-in mathematical functions to use in a SELECT clause.
- A5 Use SELECT MAX(...
- A6 YES.
- A7 Because SALARY is not declared to have NULLs and the employees for whom the salary is not known are also counted.
- A8 SELECT SUBSTR(FIRSTNAME,1,5) FROM EMP;
- A9 SELECT FIRSTNAME || ' ' || LASTNAME FROM EMP;
- A10 1. Avoid -ve SQLCODEs by handling nulls and zeroes in computations
2. Substitute a numeric value for any nulls used in computation
- A11 UNION : eliminates duplicates
UNION ALL: retains duplicates
Both these are used to combine the results of different SELECT statements.
- A12 Once.
- A13 It has to be in a CURSOR.
- A14 BETWEEN supplies a range of values while IN supplies a list of values.
- A15 Yes.
- A16 LIKE is used for partial string matches. '%' (for a string of any character) and '_' (for any single character) are the two wild card characters.
- A17 To do partial search e.g. to search employee by name, you need not specify the complete name; using LIKE, you can search for partial string matches.
- A18 Match for any single character.
- A19 GROUP BY partitions the selected rows on the distinct values of the column on which you group by.
HAVING selects GROUPs which match the criteria specified

- A20 SELECT EMPNO
 FROM EMP
 WHERE PROJECT IS NULL;
- A21 SELECT SUM(SALARY)
 FROM EMP
 WHERE QUAL='MSC';

 NULL
- A22 For three reasons:
 ☐ If the table structure is changed (a field is added), the program will have to be modified
 ☐ Program might retrieve the columns which it might not use, leading on I/O over head.
 ☐ The chance of an index only scan is lost.
- A23 A subquery in which the inner (nested) query refers back to the table in the outer query. Correlated subqueries must be evaluated for each qualified row of the outer query that is referred to.
- A24 ANSWER NOT KNOWN
- A25 Cursor is a programming device that allows the SELECT to find a set of rows but return them one at a time.
 Cursor should be used because the host language can deal with only one row at a time.
- A26 Either by using the single row SELECT statements, or by using the CURSOR.
- A27 Single row SELECTs.
- A28 See answer to next question.
- A29 Use DECLARE CURSOR statement either in working storage or in procedure division (before open cursor), to specify the SELECT statement. Then use OPEN, FETCH rows in a loop and finally CLOSE.
- A30 If there is an ORDER BY clause, rows are fetched, sorted and made available for the FETCH statement. Other wise simply the cursor is placed on the first row.
- A31 No.
- A32 Yes.
- A33 Yes
- A34 Use WITH HOLD option in DECLARE CURSOR statement. But, it has not effect in psuedo-conversational CICS programs.

- A35 A VARCHAR column REMARKS would be defined as follows:
...
10 REMARKS.
49 REMARKS-LEN PIC S9(4) USAGE COMP.
49 REMARKS-TEXT PIC X(1920).
- A36 DATE: 4bytes
TIME: 3bytes
TIMESTAMP: 10bytes
- A37 DATE: PIC X(10)
TIME : PIC X(08)
TIMESTAMP: PIC X(26)
- A38 PIC S9(9)V99 COMP-3.
Note: In DECIMAL(11,2), 11 indicates the size of the data type and 2 indicates the precision.
- A39 DeCLarations GENerator: used to create the host language copy books for the table definitions. Also creates the DECLARE table.
- A40 1. EXEC SQL DECLARE TABLE statement which gives the layout of the table/view in terms of DB2 datatypes.
2. A host language copy book that gives the host variable definitions for the column names.
- A41 It is not mandatory to use DCLGEN.
Using DCLGEN, helps detect wrongly spelt column names etc. during the pre-compile stage itself (because of the DECLARE TABLE). DCLGEN being a tool, would generate accurate host variable definitions for the table reducing chances of error.
- A42 It not necessary to have DECLARE TABLE statement in DCLGEN. This is used by the pre-compiler to validate the table-name, view-name, column name etc., during pre-compile.
- A43 No. Because the precompiler does not refer to the DB2 catalogue tables.
- A44 1. Use DSN utility to run a DB2 batch program from native TSO. An example is shown:

DSN SYSTEM(DSP3)
RUN PROGRAM(EDD470BD) PLAN(EDD470BD) LIB('ED01T.OBJ.LOADLIB')
END

2. Use IKJEFT01 utility program to run the above DSN command in a JCL.
- A45 Query the catalogue tables SYSPLANDEP and SYSPACKDEP.
- A46 SQLCODE, SQLERRM, SQLERRD
- A47 Check the value stored in SQLERRD(3).

- A48 EXPLAIN is used to display the access path as determined by the optimizer for a SQL statement. It can be used in SPUFI (for single SQL statement) or in BIND step (for embedded SQL).
- A49 Make sure that the PLAN_TABLE is created under the AUTHID.
- A50 In userid.PLAN_TABLE
- A51 a non matching index scan if ACESSTYPE = I.
- A52 1. Use SPUFI or QMF to EXPLAIN the dynamic SQL statement
2. Include EXPLAIN command in the embedded dynamic SQL statements
- A53 Use a question mark in place of a host variable (or an unknown value). e.g.
SELECT EMP_NAME
FROM EMP
WHERE EMP_SALARY > ?
- A54 CS: Cursor Stability
RR: Repeatable Read
- A55 CS: Releases the lock on a page after use
RR: Retains all locks acquired till end of transaction
- A56 ISOLATION LEVEL is a parameter for the bind process.
- A57 During the BIND process. ISOLATION (CS/RR)...
- A58 No.
- A59 PAGE, TABLE, TABLESPACE
- A60 1. Based on the lock-size given while creating the tablespace
2. Programmer can direct the DB2 what lock-size to use
3. If lock-size ANY is specified, DB2 usually choses a lock-size of PAGE
- A61 Hig resource utilization if large updates are to be done
- A62 Promoting a PAGE lock-size to table or tablespace lock-size when a transaction has aquired more locks than specified in NUMLKTS. Locks should be taken on objects in single tablespace for escalation to occur.
- A63 SHARE, EXCLUSIVE, UPDATE
- A64 No. To lock a view, take lock on the underlying tables.
- A65 SQL command used to change the definition of DB2 objects.

- A66 DBRM: DataBase Request Module, has the SQL statements extracted from the host language program by the pre-compiler.
PLAN: A result of the BIND process. It has the executable code for the SQL statements in the DBRM.
- A67 Determine the point at which DB2 acquires or releases locks against table and tablespaces, including intent locks.
- A68 PLAN has the executable code for the SQL statements in the host program
- A69 Plan is marked as invalid. The next time the plan is accessed, it is rebound.
- A70 They contain executable code for SQL statements for one DBRM.
- A71
1. Avoid having to bind a large number of DBRM members into a plan
 2. Avoid cost of a large bind
 3. Avoid the entire transaction being unavailable during bind and automatic rebind of a plan
 4. Minimize fallback complexities if changes result in an error.
- A72 a user defined name that is the anchor for packages. It has not physical existence. Main usage is to group packages.
- A73
- 100 (for successful completion of the query), 0 (for successful COMMIT if AUTOCOMMIT is set to Yes).
- A74 Print the output dataset.
- A75 ANSWER NOT KNOWN
- A76 Looks like index page split has ocured. DO a REORG of the indexes.
- A77 Dynamic SQL is a SQL statement created at program execution time.
- A78 At run time, when the PREPARE statement is issued.
- A79 Probably RUN STATS is not done and the program is using a wrong index due to incorrect stats.
Probably RUNSTATS is done and optimizer has chosen a wrong access path based on the latest statistics.
- A80 as an extra-byte prefix to the column value. physically, the nul prefix is Hex '00' if the value is present and Hex 'FF' if it is not.
- A81 Use null indicators. Syntax ... INTO :HOSTVAR:NULLIND
- A82 S9(4) COMP.

A83 -1 : the field is null
 0 : the field is not null
 -2 : the field value is truncated

A84 To insert a NULL, move -1 to the null indicator
 To insert a valid value, move 0 to the null indicator

A85 A DB2 utility used to collect statistics about the data values in tables which can be used by the optimizer to decide the access path. It also collects statistics used for space management. These statistics are stored in DB2 catalog tables.

A86 After a load, or after mass updates, inserts, deletes, or after REORG.

A87 # of rows in the table
 Percent of rows in clustering sequence
 # of distinct values of indexed column
 # of rows moved to a nearby/farway page due to row length increase

A88 REORG reorganizes data on physical storage to recluster rows, positioning overflowed rows in their proper sequence, to reclaim space, to restore free space. It is used after heavy updates, inserts and delete activity and after segments of a segmented tablespace have become fragmented.

A89 It is full backup of a DB2 table which can be used in recovery.

A90 To take routine backup of tables
 After a LOAD with LOG NO
 After REORG with LOG NO

A91 A state in which, an image copy on a table needs to be taken. In this status, the table is available only for queries. You cannot update this table. To remove the COPY PENDING status, you take an image copy or use REPAIR utility.

A92 When a table is LOADED with ENFORCE NO option, then the table is left in CHECK PENDING status. It means that the LOAD utility did not perform constraint checking.

A93 A QUIESCE flushes all DB2 buffers on to the disk. This gives a correct snapshot of the database and should be used before and after any IMAGECOPY to maintain consistency.

A94 Causes the data rows to be stored in the order specified in the index. A mandatory index defined on a partitioned table space.

A95 Only one.

A96 Primary : a relational database constraint. Primary key consists of one or more columns that uniquely identify a row in the table. For a normalized relation, there is one designated primary key.
 Unique index: a physical object that stores only unique values. There can be one or more unique indexes on a table.

A97 Authorization failure

A98 SELECT statement has resulted in retrieval of more than one row.

A99 This is generated when the consistency tokens in the DBRM and the load module are different.

A100 Not all of them. Some views are updatable e.g. single table view with all the fields or mandatory fields. Examples of non-updatable views are views which are joins, views that contain aggregate functions(such as MIN), and views that have GROUP BY clause.

A101 No.

A102 TSO, CICS, IMS and BATCH

A103 Inner Join: combine information from two or more tables by comparing all values that meet the search criteria in the designated column or columns of one table with all the values in corresponding columns of the other table or tables. This kind of join which involves a match in both columns are called inner joins.

Outer join is one in which you want both matching and non matching rows to be returned. DB2 has no specific operator for outer joins, it can be simulated by combining a join and a correlated sub query with a UNION.

A104 PCTFREE: percentage of each page to be left free
FREEPAGE: Number of pages to be loaded with data between each free page

A105 Simple Tablespace:

Can contain one or more tables

Rows from multiple tables can be interleaved on a page under the DBAs control and maintenance

Segmented Tablespace:

Can contain one or more tables

Tablespace is divided into segments of 4 to 64 pages in increments of 4 pages. Each segment is dedicated to single table. A table can occupy multiple segments

Partitioned Tablespace:

Can contain one table

Tablespace is divided into parts and each part is put in a separate VSAM dataset.

A106 one divided by the number of distinct values of a column.

A107 The number of distinct values a column or columns contain.

- A108 Synonym is an alternate name for a table or view used mainly to hide the leading qualifier of a table or view.. A synonym is accessible only by the creator.
- A109 SYNONYM: is dropped when the table or tablespace is dropped. Synonym is available only to the creator.
ALIAS: is retained even if table or tablespace is dropped. ALIAS can be created even if the table does not exist. It is used mainly in distributed environment to hide the location info from programs. Alias is a global object & is available to all.
- A110 This column cannot have nulls and while insertion, if no value is supplied then it wil have zeroes, spaces or date/time depending on whether it is numeric, character or date/time.
Use it when you do not want to have nulls but at the same time cannot give values all the time you insert this row.
- A111 The column cannot have nulls. Use it for key fields.
- A112 When a column which contains long text, e.g. remarks, notes, may have in most cases less than 50% of the maximum length.
- A113 1. Can lead to high space utilization if most of the values are close to maimum.
2. Positioning of VARCHAR column has to be done carefully as it has performance implications.
3. Relocation of rows to different pages can lead to more I/Os on retrieval.
- A114 First CREATE MANAGER table with EMP# as the primary key. Then ALTER it to define the foreign key.
- A115 At run time.
- A116 Recording SQL statements that access a table. Specified at table creation time or thru alter.
- A117 A DB2 bind is a process that builds an access path to DB2 tables.
- A118 An access path is the method used to access data specified in DB2 sql statements.
- A119 An application plan or package is generated by the bind to define an access path.
- A120 Normalization is a design procedure for representing data in tabular format. The five normal forms are progressive rules to represent the data with minimal redundancy.
- A121 These are attributes of one table that have matching values in a primary key in another table, allowing for relationships between tables
- A122 SELECT element FROM table WHERE conditional statement.
- A123 WHERE is used with a relational statement to isolate the object element or row.
- A124 Joins, unions and nested selects are used to retrieve data.

- A125 DDL is data definition language and DML is data manipulation language. DDL statements are CREATE, ALTER, TRUNCATE. DML statements are SELECT, INSERT, DELETE and UPDATE.
- A126 A view is a virtual table made up of data from base tables and other views, but not stored separately.
- A127 An outer join includes rows from tables when there are no matching values in the tables.
- A128 A subselect is a select which works in conjunction with another select. A nested select is a kind of subselect where the inner select passes to the where criteria for the outer select.
- A129 Group by controls the presentation of the rows, order by controls the presentation of the columns for the results of the SELECT statement.
- A130 The explain statement provides information about the optimizer's choice of access path of the sql.
- A131 Tables are stored in tablespaces (hence the name)! There are three types of tablespaces: simple, segmented and partitioned.
- A132 An embedded sql statement may return a number of rows while the programming language can only access one row at a time. The programming device called a cursor controls the position of the row.
- A133 Referential integrity refers to the consistency that must be maintained between primary and foreign keys, ie every foreign key value must have a corresponding primary key value.
- A134 I/O operations are usually most critical for DB2 performance (or any other database for that matter).
- A135 Denormalizing DB2 tables reduces the need for processing intensive relational joins and reduces the
Number of foreign keys.
- A136 The database descriptor, DBD is the DB2 component that limits access to the database whenever objects are created, altered or dropped.
- A137 To maintain the integrity of DB2 objects the DBD permits access to only on object at a time. Lock contention happens if several objects are required by contending application processes simultaneously.
- A138 SPUFI stands for SQL processing using file input. It is the DB2 interactive menu-driven tool used by developers to create database objects.

- A139 The two parameters used in the CREATE statement are the PCTFREE which specifies the percentage of free space for each page and FREEPAGE which indicates the number of pages to be loaded with data between each free page. Free space allows room for the insertion of new rows.
- A140 A NULL value takes up one byte of storage and indicates that a value is not present as opposed to a space or zero value. It's the DB2 equivalent of TBD on an organizational chart and often correctly portrays a business situation. Unfortunately, it requires extra coding for an application program to handle this situation.
- A141 A synonym is used to reference a table or view by another name. The other name can then be written in the application code pointing to test tables in the development stage and to production entities when the code is migrated. The synonym is linked to the AUTHID that created it.
- A142 An alias is an alternative to a synonym, designed for a distributed environment to avoid having to use the location qualifier of a table or view. The alias is not dropped when the table is dropped.
- A143 A LIKE table is created by using the LIKE parameter in a CREATE table statement. LIKE tables are typically created for a test environment from the production environment.
- A144 No. The table and its view are created anew, but the programs accessing the view do not need to be changed if the view and attributes accessed remain the same.
- A145 Never. Such processing could produce duplicate values violating entity integrity. Primary keys must be updated one at a time.
- A146 The cascade rule will not allow deletions based on a subselect that references the same table from which the deletions are being made.
- A147 The self-referencing constraint limits in a single table the changes to a primary key that the related foreign key defines. The foreign key in a self referencing table must specify the DELETE CASCADE rule.
- A148 Tables related with a foreign key are called delete-connected because a deletion in the primary key table can affect the contents of the foreign key table.
- A149 Never. New primary key values are not a problem. However, the values of foreign key inserts must have corresponding primary key values in their related tables. And updates of primary key values may require changes in foreign key values to maintain referential integrity.
- A150 The simplest DB2 index is the B-tree and the B-tree's top page is called the root page. The root page entries represent the upper range limits of the index and are referenced first in a search.
- A151 DB2 use the multiple indexes to satisfy multiple predicates in a SELECT statement that are joined by an AND or OR.

- A152 Primary key and foreign key columns; columns that have unique values; columns that have aggregates computed frequently and columns used to test the existence of a value.
- A153 A multiple index is not one index but two indexes for two different columns of a table. A composite index is one index made up of combined values from two columns in a table. If two columns in a table will often be accessed together a composite index will be efficient.
- A154 The number of distinct values for a column is called index cardinality. DB2's RUNSTATS utility analyzes column value redundancy to determine whether to use a tablespace or index scan to search for data.
- A155 For a clustered index DB2 maintains rows in the same sequence as the columns in the index for as long as there is free space. DB2 can then process that table in that order efficiently.
- A156 The LIKE keyword allows for string searches. The % sign is used as a wildcard.
- A157 The common aggregate, built-in functions are AVG, SUM, MIN, MAX, COUNT and DISTINCT.
- A158 SUBSTR is used for string manipulation with column name, first position and string length used as arguments. Eg. SUBSTR (NAME, 1 3) refers to the first three characters in the column NAME.
- A159 The three data types are DATE, TIME and TIMESTAMP. CHAR can be used to specify the format of each type. The DAYS function calculates the number of days between two dates. (It's Y2K compliant).
- A160 In DB2 a transaction typically requires a series of updates, insertions and deletions that represent a logical unit of work. A transaction puts an implicit lock on the DB2 data. Programmers can use the COMMIT WORK statement to terminate the transaction creating smaller units for recovery. If the transaction fails DB2 uses the log to roll back values to the start of the transaction or to the preceding commit point.
- A161 Deadlock occurs when transactions executing at the same time lock each other out of data that they need to complete their logical units of work.
- A162 DB2 imposes locks of four differing sizes: pages, tables, tablespace and for indexes subpage.
- A163 The three types are shared, update and exclusive. Shared locks allow two or more programs to read simultaneously but not change the locked space. An exclusive lock bars all other users from accessing the space. An update lock is less restrictive; it allows other transactions to read or acquire shared locks on the space.
- A164 SQL statements may return any number of rows, but most host languages deal with one row at a time by declaring a cursor that presents each row at a unique isolation level.
- A165 An intent lock is at the table level for a segmented tablespace or at the tablespace level for a nonsegmented tablespace. They indicate at the table or tablespace level the kinds of locks at lower levels.

- A166 Static sql is hard-coded in a program when the programmer knows the statements to be executed. For dynamic sql the program must dynamically allocate memory to receive the query results.
- A167 Cursor stability means that DB2 takes a lock on the page the cursor is accessing and releases the lock when the cursor moves to another page.
- A168 The clause avoids closing the cursor and repositioning it to the last row processed when the cursor is reopened.
- A169 It is a data structure that must be included in any host-language program using SQL. It is used to pass feedback about the sql operations to the program. Fields are return codes, error messages, handling codes and warnings.
- A170 The WHENEVER statement is coded once in the host program to control program actions depending on the SQL-CODE returned by each sql statement within the program.
- A171 DCLGEN stands for declarations generator; it is a facility to generate DB2 sql data structures in COBOL or PL/I programs.
- A172 The FREE command can be used to delete plans and/or packages no longer needed.
- A173 A merge join requires that the tables being joined be in a sequence; the rows are retrieved with a high cluster ratio index or are sorted by DB2. A nested join does not require a sequence and works best on joining a small number of rows. DB2 reads the outer table values and each time scans the inner table for matches. The hybrid join is a nested join that requires the outer table be in sequence.
- A174 Any subselect can be rewritten as a join, but not vice versa. Joins are usually more efficient as join rows can be returned immediately, subselects require a temporary work area for inner selects results while processing the outer select.
- A175 If there is an index on the attributes tested an IN is more efficient since DB2 uses the index for the IN. (IN for index is the mnemonic).
- A176 A Cartesian product results from a faulty query. It is a row in the results for every combination in the join tables.
- A177 Package and plan are usually used synonymously, as in this site. Both contain optimized code for SQL statements - a package for a single program, module or subroutine contained in the database request module (DBRM) library. A plan may contain multiple packages and pointers to packages. The one CICS module would then exist in a package that could be referenced in two different plans.
- A178 It is a write to disk that may occur before or long after a commit. The write is controlled by the buffer manager.
- A179 A lock is the mechanism that controls access to data pages and tablespaces.

- A180 This is a key concept for any relational database. Isolation level is the manner in which locks are applied and released during a transaction. For DB2 a 'repeatable read' holds all locks until the transaction completes or a syncpoint is issued. For transactions using 'cursor stability' the page lock releases are issued as the cursor 'moves', i.e. as the transaction releases addressability to the records.
- A181 They are the opposite of root pages. Leaf pages are the lowest level index pages - the pages that contain index entries and information to the corresponding table rows.
- A182 It is a DB2 facility for static SQL statements - it replaces these statements with calls to the DB2 language interface module.
- A183 The opposite of a leaf page; it is the highest level index page. An index can contain only the one root page; all other index pages are associated to the root.
- A184 A thread is the connection between DB2 and some other subsystem, such as CICS or IMS/DC.