

# PASSTCERT

QUESTION & ANSWER

Higher Quality  
Better Service!

We offer free update service for one year  
[HTTP://WWW.PASSTCERT.COM](http://www.passtcert.com)

**Exam** : **NR0-012**

**Title** : Teradata Physical  
Implementation v2r5 Exam

**Version** : DEMO

**1.Which join plan could make use of a primary index on each of two tables with matching Partitioned Primary Indexes?**

- A.Merge
- B.Bitmap
- C.Product
- D.Exclusion Merge

**Correct:A**

**2.What are three functions performed by the Index Wizard? (Choose three.)**

- A.automatically drops existing indexes based on the workload
- B.recommends adding indexes based on specified workloads
- C.recommends dropping the indexes that may not be necessary
- D.automatically creates DDL statements for recommended indexes
- E.recommends Partitioned Primary Indexes based on specified workloads

**Correct:B C D**

**3.Which three statements are true concerning join plans? (Choose three.)**

- A.A Nested Join is only used when the user specifies a constant value in a WHERE clause.
- B.A Hash Join requires that the data from both tables be sorted before performing the join operation.
- C.A Merge Join requires that the data from the tables be sorted before performing the join operation.
- D.A Merge Join retrieves rows from two tables and requires them to be on the same AMP based on the row hash of the columns involved in the join.

**Correct:A C D**

**4.What defines whether parallel steps have been generated in an EXPLAIN output?**

- A.the phrase "we execute parallel steps on each AMP"
- B.the phrase "we execute the following steps in unison"
- C.the phrase "we execute the following steps in parallel"
- D.the phrase "we execute the following steps as an all-AMPs RETRIEVE"

**Correct:C**

**5.At query execution time, which two statements about nested views are true? (Choose two.)**

- A.Nested views are not updatable.
- B.Maximum size of the fully expanded text is 2 MB.
- C.Teradata validates nested view privileges of the executing user at execution time.
- D.If the view contains WITH CHECK OPTION, the WHERE clause used in defining the view is ignored.

**Correct:B C**

**6.Which two statements about dynamic AMP Sampling are true? (Choose two.)**

- A.An AMP is selected to represent an estimation of population demographics.
- B.A random data block from a random cylinder is selected and the number of rows it contains are counted.
- C.The Optimizer understands the reliability of the information it collects using a random AMP and it can be more aggressive in its pursuit of optimal access and join strategies.
- D.The potential for error introduced from a large table is much higher than the estimates gained by the full-table statistics gathered using the Optimizer form of COLLECT STATISTICS.

**Correct:A B**

**7.An EXPLAIN output contains the following phrase: "by way of a traversal of index #n extracting row ids only" Which statement is true?**

- A.A spool file is built from the rowhash found in the cylinder index.

- B.Rowid values are accessed in memory without the need for a spool file.
- C.A spool file is built from rowid values found in a secondary index (Index #n).
- D.A temporary file is accessed with rowid values found in a secondary index (Index #n).

**Correct:C**

**8.When using Visual EXPLAIN, which three menu options are available? (Choose three.)**

- A.Show Connectors
- B.Replace Execution Plans
- C.Open Plan from Database
- D.Explain Textual Differences

**Correct:A C D**

**9.Collected statistics on a Teradata table allows the Optimizer to predict which two factors? (Choose two.)**

- A.the size of spool files
- B.the AMP to send data to
- C.the number of occurrences of a value
- D.the PE that will be assigned to the process

**Correct:A C**

**10.Which three statements are true about Teradata Statistics? (Choose three.)**

- A.Statistics can be collected on a global temporary table.
- B.Table statistics are refreshed automatically by Teradata.
- C.Statistics are eliminated when all rows from the table are deleted.
- D.Table statistics can be viewed using the STATS module of Teradata Manager.
- E.User must have at least one appropriate privilege on the object to COLLECT STATISTICS.

**Correct:A D E**

**11.An EXPLAIN output contains one of the following phrases: "a single partition of" or "# partitions of" Which three statements are true? (Choose three.)**

- A.The query has a range constraint on the partitioning columns.
- B.The query has an equality constraint on the partitioning column.
- C.Indicates that the table is being accessed using a sparse index.
- D.Indicates AMPs will only access a single partition or # partitions of a table.

**Correct:A B D**

**12.Which three statements about Visual EXPLAIN are true? (Choose three.)**

- A.identifies what the Optimizer chooses and why
- B.identifies what the Optimizer chooses, not why
- C.translates Optimizer output (the parse tree) into icons
- D.may be used on any SQL statement, except EXPLAIN
- E.may be used on any SQL statement, including EXPLAIN

**Correct:B C D**

**13.In an EXPLAIN output, what indicates that a query contains a join or a subquery?**

- A.the subquery step that is identified
- B.the type of join step that is identified
- C.the number of retrieve steps that are generated
- D.the number of parallel steps that are generated

**Correct:B**

**14.Which three are characteristics of a Volatile table? (Choose three.)**

- A.Rows are deleted at session end.
- B.Spool space is used to materialize the table.
- C.Rows are lost following a system restart.
- D.Definitions are stored in the data dictionary.
- E.Temporary space is used to materialize the table.

**Correct:A B C**

**15.Which two definitions are required on the referenced columns when implementing Referential Integrity? (Choose two.)**

- A.Unique
- B.Nullable
- C.Indexed
- D.Not Null

**Correct:A D**

**16.Partitioned Primary Indexes are best suited for which two types of tables? (Choose two.)**

- A.large tables that are not referenced using range constraints
- B.large tables that are frequently accessed using restrictive conditions
- C.large tables that are maintained by deleting and adding ranges of rows
- D.large tables that are frequently accessed for Referential Integrity checks

**Correct:B C**

**17.What are three valid DATABLOCKSIZE options in the CREATE TABLE statement? (Choose three.)**

- A.MINIMUM DATABLOCKSIZE
- B.MAXIMUM DATABLOCKSIZE
- C.OWNERS DATABLOCKSIZE
- D.DATABLOCKSIZE = MAXIMUM
- E.DATABLOCKSIZE = 32 KILOBYTES

**Correct:A B E**

**18.Which three statements about aggregations in views are correct? (Choose three.)**

- A.They simplify end-user coding requirements.
- B.Alias names are not required in views with aggregates.
- C.They simplify the maintenance of complicated calculations.
- D.They are affected if new columns are added to underlying tables.
- E.They are unaffected if new columns are added to underlying tables.

**Correct:A C E**

**19.Which three are valid for the ALTER TABLE command? (Choose three.)**

- A.ADD columns up to 2048 columns
- B.ADD compression to an existing column
- C.ADD a new column with multi-value compression
- D.MODIFY CONSTRAINT for a named CHECK constraint
- E.ADD a compressible value to a column with multi-value compression

**Correct:A C D**

**20.To equi-join a Partitioned Primary Index (PPI) table to a Non-Partitioned Primary Index (NPPI) table, or to another PPI table defined with a different partitioning expression, what does the**

**Optimizer evaluate? (Choose three.)**

- A.using a "sliding window" join of the two tables without re-partitioning either table
- B.spooling one or both PPIs into a NPPI spool and joining these using a conventional merge join
- C.re-partitioning and redistributing both tables to have matching partition definitions in preparation for a rowkey-based Merge Join
- D.re-partitioning, if possible, one of the two tables into a PPI spool with the same partitioning as the PPI table in preparation for a rowkey-based Merge Join

**Correct:A B D**