

IBM 000-553

Netezza Performance Software v6

Practice Test

Version: 14.20

QUESTION NO: 1

Which CREATE DATABASE attributes are required?

- A. The database name.
- B. The database name and the redo log file name.
- C. The database name and the table space name.
- D. The database name and the temporary table space name.

Answer: A

Explanation:

QUESTION NO: 2

Which statement is true for database users and groups?

- A. All users must belong to admin group.
- B. Creation of users and groups is not allowed.
- C. Users and groups are local and tied to a particular database.
- D. Users and groups are global and not tied to a particular database.

Answer: D

Explanation:

QUESTION NO: 3

What is the maximum number of columns you can choose as organizing keys for clustered base table?

- A. 1
- B. 2
- C. 4
- D. 8

Answer: C

Explanation:

QUESTION NO: 4

Which two characteristics describe the materialized view on TwinFin. (Choose two.)

- A. Materialized views can contain aggregates.
- B. Materialized views can contain a HAVING clause.
- C. Materialized views can contain a WHERE clause.
- D. Materialized views can contain an ORDER BY clause.
- E. Materialized views can reference only one base table in the FROM clause.

Answer: D,E

Explanation:

QUESTION NO: 5

Which statement is true regarding any database user?

- A. The user has a table space defined.
- B. The user can be assigned to more than one group.
- C. The user must have a corresponding host account.
- D. The user can be removed from the default PUBLIC group.

Answer: B

Explanation:

QUESTION NO: 6

When working with sequences, which of the following statements will always be true?

- A. One-up numbers are generated.
- B. Integers (whole numbers) are generated.
- C. The first value generated is the number 1.
- D. The sequence number assigned will be unique across all tables and databases.

Answer: B

Explanation:

QUESTION NO: 7

When are zonemaps initially created for a table?

- A. nzbackup
- B. nzreclaim
- C. CREATE TABLE
- D. ALTER TABLE ADD ZONEMAPS

Answer: C

Explanation:

QUESTION NO: 8

Columns in a table may be zonemapped. How are those columns chosen?

- A. They are automatically selected by the system.
- B. They are based on the same columns as the table's distribution key.
- C. They are added via the ALTER TABLE ADD INDEX(column_name) statement.
- D. They are added via the ALTER TABLE ADD ZONEMAP(column_name) statement.

Answer: A

Explanation:

QUESTION NO: 9

Which clause CANNOT be used with a CREATE VIEW statement?

- A. SUM ()
- B. GROUP BY
- C. ORDER BY
- D. OUTER JOIN

Answer: C

Explanation:

QUESTION NO: 10

What is a characteristic of a column that is based on the TIMESTAMP datatype?

- A. It includes a date portion.
- B. It includes a timezone offset.

- C. The value is reevaluated each time the row is selected.
- D. The value is initially assigned when the row is first created.

Answer: A

Explanation:

QUESTION NO: 11

When creating a table, you can specify a column constraint for each individual column of either "NULL" or "NOT NULL". When column values are later evaluated within a query statement, which of these is true?

- A. 0 is equal to null, the empty string (") is equal to null.
- B. 0 is equal to null, the empty string (") is not equal to null.
- C. 0 is not equal to null, the empty string (") is equal to null.
- D. 0 is not equal to null, the empty string (") is not equal to null.

Answer: D

Explanation:

QUESTION NO: 12

After performing a GROOM on a clustered base table, which statement is true?

- A. The table is redistributed on the ORGANIZE ON clause.
- B. The table is reorganized on the ORGANIZE ON clause.
- C. The table is reorganized on the DISTRIBUTE ON clause.
- D. The table is redistributed on the DISTRIBUTE ON clause.

Answer: B

Explanation:

QUESTION NO: 13

Which two statistical data points are automatically updated when performing an INSERT on a table? (Choose two.)

- A. The number of rows in the table.

- B. The column histogram statistics.
- C. The number of NULLs in each column.
- D. The number of unique values in each column.
- E. The minimum and maximum column value information (excluding text columns).

Answer: A,E

Explanation:

QUESTION NO: 14

Using a date column as a distribution key may distribute rows evenly across dataslices, why is this a poor choice of a distribution key?

- A. Queries on the table will invoke compression.
- B. Queries on date range may involve processing skew.
- C. Queries on date range will utilize only host processing.
- D. Queries on the table will invoke compression on the date columns.

Answer: B

Explanation:

QUESTION NO: 15

Why should you avoid using Boolean data types as a distribution key for a table ?

- A. Table will not join.
- B. Table will not broadcast.
- C. Table is likely to be skewed.
- D. Table is likely to be evenly spread across data slices.

Answer: C

Explanation:

QUESTION NO: 16

Which statement about the TwinFin Groom feature is true?

- A. GROOM TABLE requires space to make a complete copy of the table.

- B. GROOM TABLE only operates on tables that include an ORGANIZE ON clause.
- C. GROOM TABLE prevents nzload from adding data while GROOM is in progress.
- D. GROOM TABLE allows UPDATE, DELETE, and INSERT operations to occur while GROOM is in progress.

Answer: D

Explanation:

QUESTION NO: 17

What should be considered when you are asked to select the distribution key columns for a very large fact table?

- A. As many columns as possible.
- B. Columns that contain many nulls.
- C. Columns that contain few unique values.
- D. Columns used to join to other large tables.

Answer: D

Explanation:

QUESTION NO: 18

What is the raw user data space available on a TwinFin 24?

- A. 24TB
- B. 48TB
- C. 64TB
- D. 100TB

Answer: C

Explanation:

QUESTION NO: 19

What is the TwinFin host Linux file that must be modified to adjust optimizer settings permanently?

- A. /nz/data/postgresql.conf

- B. /nz/data/config/system.cfg
- C. /nz/kit/sys/initTopology.cfg
- D. /nz/kit/log/sysmgr/sysmgr.log

Answer: A

Explanation:

QUESTION NO: 20

Click on the exhibit.

CREATE Table TableA

```
( c_custkey      integer,  
  c_name        varchar(25),  
  c_address     VARCHAR(40),  
  c_nationkey   integer,  
  c_phone       character(15),  
  c_acctbal     numeric(15,2),  
  c_mktsegment  character(10),  
  c_comment     VARCHAR(117)  
) DISTRIBUTE ON ( c_custkey );
```

CREATE TABLE TableB

```
( o_orderkey     integer,  
  o_custkey     integer,  
  o_orderstatus  character(1),  
  o_totalprice  numeric(15,2),  
  o_orderdate   date,  
  o_orderpriority character(15),  
  o_clerk       character(15),  
  o_shippriority integer,  
  o_comment     VARCHAR(79)  
) DISTRIBUTE ON ( o_custkey );
```

Given the following query:

```
SELECT * FROM TableA, TableB WHERE TableA.c_custkey = TableB.o_custkey;
```

What type of join processing will occur when the query is executed?

- A. Collocation hash join processing
- B. Distributed hash join processing
- C. Broadcast sort merge join processing