



Welcome to Cert007 - Your Ultimate IT Certification Partner



- Real Exam Questions
- Free Updates
- Expert Support
- Instant Access
- Money-Back Guarantee



Visit us at <https://www.cert007.com/> for more information

Exam : **TDVAN5**

Title : **Vantage Administration
Exam**

Version : **DEMO**

1. There is a call center application that repetitively sends tactical queries to Vantage. These queries use a few small tables that are joined on the primary index column.

The following maps are defined in the system:

* TD_GlobalMap

* TD_Map1

* TD_DataDictionaryMap

* TD_1AmpSparseMap_1Node

What can be done to optimize these queries?

A. Assign these tables to TD_1AmpSparseMap_1Node using different colocation names for each table.

B. Assign these tables to TD_1AmpSparseMap_1Node using the same colocation name for every table.

C. Assign these tables to TD_Map1 using the same colocation name for every table.

D. Assign these tables to TD_Map1 using different colocation names for each table.

Answer: B

Explanation:

Example:

TD_1AmpSparseMap_1Node is a sparse map that assigns the data to a single AMP (or a few AMPs), and using the same colocation name ensures that the tables are colocated on the same AMP. This helps in efficient joining because no data redistribution is required between AMPs when tables are joined on the primary index.

Using different colocation names for each table (options A and D) would place the tables on different AMPs, leading to less efficient joins since data would need to be shuffled between AMPs.

TD_Map1 is a predefined map in the system but does not specifically optimize small, frequently accessed tables in the same way that TD_1AmpSparseMap_1Node does, which is more suitable for these scenarios.

Thus, using the same colocation name within the TD_1AmpSparseMap_1Node ensures that the joins are AMP-local, optimizing the repetitive tactical queries.

2. Which workload management technique should the Administrator use to reject all QueryGrid queries during a pre-defined critical batch state?

A. An exception that includes the following Foreign Servers:

TD_SYSFNLIB.QGInitiatorExport

TD_SYSFNLIB.QGInitiatorImport

TD_SYSFNLIB.QGRemoteExport

TD_SYSFNLIB.QGRemoteImport

B. A filter that includes the following functions:

TD_SYSFNLIB.QGInitiatorExport

TD_SYSFNLIB.QGInitiatorImport

TD_SYSFNLIB.QGRemoteExport

TD_SYSFNLIB.QGRemoteImport

C. An exception that excludes the following Foreign Servers:

TD_SYSFNLIB.QGInitiatorExport

TD_SYSFNLIB.QGInitiatorImport

TD_SYSFNLIB.QGRemoteExport

TD_SYSFNLIB.QGRemoteImport

D. A filter that excludes the following functions:

TD_SYSFNLIB.QGInitiatorExport
TD_SYSFNLIB.QGInitiatorImport
TD_SYSFNLIB.QGRemoteExport
TD_SYSFNLIB.QGRemoteImport

Answer: A

Explanation:

In this case, using an exception with the Foreign Servers related to QueryGrid functions will prevent queries involving those foreign servers from running. This is a typical method to block or reject specific workloads during critical times, such as a batch processing window.

Exceptions in workload management are used to define specific conditions under which queries or workloads should be rejected or managed differently. By including these QueryGrid-related functions in an exception, the administrator ensures that any QueryGrid queries involving the listed foreign servers will be rejected during the critical batch state.

Options B and D, which mention filters, are not appropriate for rejecting queries. Filters are used for monitoring or routing purposes, not for outright rejection of workloads.

Option C, which suggests excluding the foreign servers in an exception, would not achieve the goal, as this would mean those queries are not affected by the exception. The goal here is to include these functions in an exception to actively reject them.

3.A customer is complaining that the creation of a large table in the lab environment is getting stuck in the merge step. The customer is using the following command to create the table: CREATE TABLE ... AS (SELECT * ...) WITH DATA

Which recommendation will help the merge step?

A. A column that provides good distribution should be specified as PRIMARY INDEX.

B. The MERGEBLOCKRATIO should be specified for the table to improve the data block merge operation.

C. The table should be created with "CREATE TABLE ... AS (SELECT *...) WITH NO DATA", and afterwards, "INSERT ... SELECT *" should be used to fill the table.

D. Users of the lab may have lower priority than other workloads, so the creation of the table should be moved to off-peak hours.

Answer: C

Explanation:

Merge step issues often occur when a large amount of data is being processed during the table creation, especially if the system is trying to simultaneously create the table and insert data.

By using "WITH NO DATA", the table structure is created first, without the actual data being inserted during the table creation process. The "INSERT ... SELECT " command can then be used afterwards to populate the table in a more controlled way, reducing the load on the system during the creation phase and potentially improving the efficiency of the merge step.

Specifying a good distribution for the primary index can help overall performance, but it doesn't directly address the issue with the merge step in this scenario.

Specifying the MERGEBLOCKRATIO isn't typically a solution for this specific problem; the merge block ratio is more about the optimization of data block merges rather than the creation of tables.

Moving the creation to off-peak hours may help if the environment is busy, but it doesn't directly address

the core issue of the merge step getting stuck.

4.Which benefit is achieved by creating profiles?

- A. Reducing the growth of the DBC.AccessRights table
- B. Managing user privileges on database objects
- C. Assigning permanent space for groups of users with similar needs
- D. Defining spool space for groups of users with similar needs

Answer: D

Explanation:

Profiles in Teradata are used to manage and enforce resource limits for groups of users. This includes defining spool space, temporary space, and permanent space for users who share similar needs. Spool space is important for query processing, and managing it by profile helps prevent individual users from consuming excessive resources and ensures better overall system performance. Reducing the growth of the DBC.AccessRights table is not directly related to profiles. Profiles do not reduce access rights or the growth of this table, which tracks individual user privileges. Managing user privileges on database objects is handled through roles and grants, not profiles.

Profiles focus on resource management, not privileges.

Assigning permanent space for groups of users is partially true since profiles can manage some space attributes, but the main focus of profiles is on managing spool space rather than permanent space.

5.What is a use case for Data Mover?

- A. Archiving data to a Hadoop system
- B. Copying data between Vantage systems for active-active replication
- C. Replicating data to a disaster recovery system
- D. Copying data between Hadoop systems

Answer: C

Explanation:

Teradata Data Mover is primarily designed to copy and replicate data between Teradata or Vantage systems. One of its common use cases is to move data to a disaster recovery system, ensuring that data is available in case of system failure or disaster, and making it a valuable tool for maintaining high availability and business continuity.

Archiving data to Hadoop and Copying data between Hadoop systems are more relevant to other tools such as Teradata QueryGrid, which integrates Vantage with Hadoop and other external systems.

Copying data between Vantage systems for active-active replication might involve Data Mover, but active-active replication typically involves more sophisticated real-time synchronization technologies like Teradata's Unity or QueryGrid.