


TDVAN5 Exam Questions Get Updated [2025 with Correct Answers [Q38-Q55]



TDVAN5 Exam Questions Get Updated [2025] with Correct Answers [Q38-Q55]

TDVAN5 Exam Questions Get Updated [2025] with Correct Answers

Practice TDVAN5 Questions With Certification guide Q&A from Training Expert Exams4sures

TeraData TDVAN5 Exam Syllabus Topics:

Topic 1- User Administration: This section evaluates the skills of system administrators in meeting user security requirements through roles and proxy access configurations. Candidates will also learn about the functionality and benefits of profiles, as well as the key attributes that should be set for a new user.
Topic 2- Security Management & Auditing: This section of the exam measures the skills of such database managers. It covers the features, functionality, and benefits of access logging and advanced security configurations.
Topic 3- Database Management: This section tests the skills of database administrators in managing access rights and advanced analytic functions. It covers logging benefits, privilege types, and the impact of changing global parameters on user sessions.
Topic 4- Monitoring Vantage: This section of the exam measures the skills of IT professionals such as system administrators. It covers how to effectively use Viewpoint to investigate system conditions, monitor external queries, and administer both Viewpoint and monitored systems.

Q38. The Administrator has received a request to add SELECT rights on the BusinessViews database to end users, developers, and batch accounts in the accounting unit. The following roles are set up for each group:

Accounting Group	Assigned Role
End Users	AcctUsers
Developers	AcctDev
Batch Accounts	AcctBatch

The Administrator created the AcctShared role and will use it in a role nesting strategy to provide the required access.

Which actions can the Administrator take to fulfill this request?

- * Grant SELECT on AcctShared to BusinessViews, then grant AcctUsers, AcctDev, and AcctBatch to AcctShared.
- * Grant SELECT on BusinessViews to AcctShared, then grant AcctUsers, AcctDev, and AcctBatch to AcctShared.
- * Grant SELECT on BusinessViews to AcctShared, then grant AcctShared to AcctUsers, AcctDev, and AcctBatch.
- * Grant SELECT on AcctShared to BusinessViews, then grant AcctShared to AcctUsers, AcctDev, and AcctBatch.

The AcctShared role should be granted SELECT access on the BusinessViews database. This ensures that the role itself has the necessary privileges.

Then, you can nest this role by granting AcctShared to the individual roles of AcctUsers, AcctDev, and AcctBatch. This role nesting strategy allows the users in these groups to inherit the permissions from AcctShared without having to directly grant the privileges to each individual role.

This approach maintains a clean and efficient permission structure using role nesting.

Q39. An update of a very large table was in progress when the system experienced an unplanned restart. After restarting, the system is available, but Recovery Manager shows it may be many hours before a large table is available for use. The Administrator needs to make the table available sooner.

Which option should be used to achieve this goal?

- * Use Workload Manager to elevate the rollback to the SLG tier.
- * Use MultiLoad to execute a DROP of the table, and restore it from backup.
- * Drop the transaction journal, and set RollbackPriority to TRUE.
- * Cancel the rollback, and restore the table from backup.

When a system experiences an unplanned restart during an update operation, the rollback process can take a significant amount of time, especially for very large tables. Canceling the rollback will stop the system from attempting to undo the incomplete transactions.

After canceling the rollback, the table can be restored from a backup, which is a much faster way to recover the table and make it available for use again.

The other options are less effective:

Use Workload Manager to elevate the rollback to the SLG tier: While this may prioritize the rollback, it doesn't significantly reduce the time required for a large table rollback.

Use MultiLoad to execute a DROP of the table, and restore it from backup: MultiLoad is not typically used for dropping tables, and this adds unnecessary complexity.

Drop the transaction journal, and set RollbackPriority to TRUE: Dropping the transaction journal could lead to data inconsistencies,

and setting RollbackPriority to TRUE does not directly make the table available sooner.

Q40. An Administrator has created a new view, ViewDB.SystemMetrics, that combines and aggregates several key system metrics. The data is non-secure in nature, and the intent is for all Vantage users to have SELECT access to it.

Which statement can the Administrator execute to meet this goal?

- * GRANT SELECT ON ViewDB.SystemMetrics TO ALL;
- * GRANT SELECT ON ViewDB.SystemMetrics TO PUBLIC;
- * GRANT SELECT ON ViewDB.SystemMetrics TO DBC;
- * GRANT SELECT ON ViewDB.SystemMetrics TO DEFAULT;

PUBLIC is a special keyword in Teradata that refers to all users on the system. Granting privileges to PUBLIC will automatically apply those privileges to every user, including future users, without the need to grant permissions individually.

Option A (GRANT SELECT ON ViewDB.SystemMetrics TO ALL) is not valid because ALL is not a recognized keyword for granting privileges to all users in Teradata.

Option C (GRANT SELECT ON ViewDB.SystemMetrics TO DBC) would grant the privilege to the DBC user (the system database user), but this would not provide access to all users.

Option D (GRANT SELECT ON ViewDB.SystemMetrics TO DEFAULT) is not a valid syntax in Teradata for granting privileges.

GRANT SELECT ON ViewDB.SystemMetrics TO PUBLIC; will allow all users to have SELECT access to the view.

Q41. An Administrator reports a power user running a query that is consuming significant CPU in a final step product join and impacting other users due to the high priority of the workload.

Upon contacting the user in the application team, they are advised that the query is a request from the CEO and needs to be completed. The team decides to remediate the situation by changing the workload to a lower priority.

Which Viewpoint portlet should be used to make this adjustment?

- * Workload Designer
- * Application Queries
- * Query Monitor
- * Workload Health

The Workload Designer portlet in Viewpoint is used to manage and adjust workloads, including setting priorities for different workloads. To address the situation where a high-priority query is consuming too many resources, the Workload Designer can be used to adjust the priority of the workload that the query belongs to, ensuring that the impact on other users is reduced while still allowing the query to complete.

The other options are less suited for adjusting workload priorities:

Application Queries is used to monitor and manage queries related to specific applications but does not provide direct options to change workload priorities.

Query Monitor allows the monitoring of active queries and possibly aborting problematic queries but does not handle workload priority changes.

Workload Health is used for monitoring the health and performance of workloads, but it is not the tool used to modify priorities.

Q42. Which QueryGrid connector can only be a target?

- * Hive
- * Oracle
- * Presto
- * Teradata

In QueryGrid, the Oracle connector can only be used as a target, meaning data can be sent to an Oracle database but not sourced from it for querying in Teradata Vantage.

Hive, Presto, and Teradata connectors can act as both source and target, allowing data to be retrieved from or sent to these systems as part of a QueryGrid query.

Thus, Oracle is the connector that can only be a target in QueryGrid.

Q43. Which description accurately characterizes the use of external authentication for Vantage?

- * The directory username must match a database username.
- * Single Sign-On is available with LDAP authentication.
- * User authorization roles can also be supplied by the directory.
- * External authentication is not permitted for mainframe clients.

This is because external authentication systems like LDAP or Active Directory can supply both authentication (verifying the user's identity) and authorization (defining what the user is allowed to do) roles for users.

Q44. An Administrator manages a Vantage system that is continually updated. The system is critical to the business and must be available as much as possible. The Administrator decides to use a backup strategy that will allow changes to tables while a backup is in progress.

Which backup strategy should be used?

- * Offline
- * Skip statistics
- * Online
- * Dictionary Only

An Online backup strategy allows the database to remain accessible and operational while the backup is taking place. This means that users can continue to modify data, and the system can remain available without requiring downtime. It is a common approach for mission-critical systems that need to maximize uptime.

Offline would require the system or certain tables to be unavailable during the backup process, which is not suitable for a system that needs to remain available.

Skip statistics refers to skipping the backup of statistics on database objects and does not relate to whether the system can be updated during the backup.

Dictionary Only involves backing up only the system catalog or dictionary data, not the actual table data, and would not fulfill the requirement of a comprehensive backup while allowing updates.

Q45. An Administrator has been presented with these performance metrics from the DBQLOGTBL table for four queries:

QRYRESPTIME	AMPCPUTIME	TOTALIOCOUNT	CPU SKEW	IO SKEW	PJ
00:20:31.6	15,876	406,587	1.04	1.08	39.
00:18:26.2	2,472	8,616,362	1.01	2.76	0.
00:06:23.2	334	6,101,075	1.03	0.99	0.
00:01:26.9	231	1,119,840	1.06	1.11	0.

The Administrator needs to identify which query is suspected to be causing a full table scan or large redistribution.

Which query is producing this result?

- * QRY_C because of UII
- * QRY_D because of CPU SKEW
- * QRY_B because of IO SKEW
- * QRY_A because of PJI

Q46. The Administrator has just created a new database object on a Vantage system.

Which privilege is NOT automatically granted to that database on itself?

- * CREATE AUTHORIZATION
- * CREATE MACRO
- * CREATE PROCEDURE
- * CREATE TRIGGER

When a new database object is created in Teradata Vantage, certain object creation privileges (such as CREATE MACRO, CREATE PROCEDURE, and CREATE TRIGGER) are typically granted automatically to the database or user who created the object.

However, the CREATE AUTHORIZATION privilege, which allows a user to create an authorization object (used for external authentication to external systems, such as a cloud storage service), is not automatically granted. This privilege requires an explicit grant.

Q47. Given a user creation request on a 10-AMP system:

```
CREATE USER hr_user AS PERM = 100e9 SKEW = 10 PERCENT,
```

```
SPOOL = 100e9 SKEW = 20 PERCENT;
```

How does the SKEW factor affect the user's Perm space, assuming the total space consumed is under the specified perm space limit?

- * The per AMP limit of any AMP can reach 20 GB.
- * The per AMP limit of 100 GB can be breached up to 10 percent.
- * The per AMP limit of 10 GB can be breached to any percent.
- * The per AMP limit of any AMP can reach 11 GB

In the given CREATE USER statement, the SKEW = 10 PERCENT parameter applies to Perm space and allows some AMPs (Access Module Processors) to use up to 10% more space than the average allocation across the AMPs.

The user is allocated 100 GB of Perm space across a 10-AMP system, meaning the average space per AMP is 10 GB.

With a 10% skew allowed, this means that an AMP can use up to 10% more than the average allocation, which is 10 GB + 1 GB = 11 GB.

Q48. An Administrator notices that a system appears to be near capacity and needs to access a dictionary to assess that AMPs are entering into flow control.

Which dictionary should be accessed for this purpose?

- * DBC.SessionInfoV

- * DBC.ResUsageSawt
- * DBC.AMPUsage
- * DBC.ResUsageSldv

The DBC.ResUsageSawt view provides detailed information about the resource usage of AMP Worker Tasks (AWTs), including whether AMPs are experiencing flow control. Flow control occurs when AMPs are overwhelmed and need to throttle the workload, and this view tracks metrics related to AWT usage and system resource contention, which would indicate when AMPs are under strain.

Option A (DBC.SessionInfoV) provides information about current user sessions but does not provide insights into AMP-level flow control or resource usage.

Option C (DBC.AMPUsage) provides general statistics about AMP usage but doesn't give detailed information about flow control or AWT usage.

Option D (DBC.ResUsageSldv) tracks statistics related to logical disk usage but isn't focused on AMP flow control.

Q49. What is a use case for Data Mover?

- * Archiving data to a Hadoop system
- * Copying data between Vantage systems for active-active replication
- * Replicating data to a disaster recovery system
- * Copying data between Hadoop systems

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.

Q50. A middle-tier application server logs on to the database as TrustedUser and submits requests on behalf of application end users. The server is shared by Finance and Human Resources groups and uses ProxyUser query band to identify end users to the database. Each group needs access to its own sensitive data, so the Administrator has created two separate roles with the appropriate permissions.

What is the best way to control access to each group's sensitive data?

- * Define the roles as external and use the ProxyRole query band to specify one role.
- * Grant both roles to TrustedUser, and add the ProxyRole query band to specify one role.
- * Specify the appropriate role for each end user in a grant connect through statement.
- * Include both roles in the grant connect through statement, and use ProxyRole in the query band to select the appropriate role.

The GRANT CONNECT THROUGH statement allows the TrustedUser to act on behalf of multiple end users while securely connecting to the database. By granting both roles (Finance and Human Resources) in this statement, you allow the ProxyUser to switch between roles depending on the query band's ProxyRole value.

Using the ProxyRole query band, the application can specify which role (Finance or Human Resources) should be used for each specific request. This approach provides flexibility, as the application can dynamically assign the appropriate role to the user based on the query context.

Option A (Defining roles as external and using ProxyRole) wouldn't fully address the need to manage multiple roles

dynamically for a shared server.

Option B (Granting both roles to TrustedUser) doesn't allow for flexible role switching on a per-request basis without the use of GRANT CONNECT THROUGH and could lead to over-granting of permissions.

Option C (Specifying a role for each end user in GRANT CONNECT THROUGH) isn't as flexible as allowing both roles to be used and dynamically selected through the query band.

Thus, Option D is the most appropriate solution, as it provides both security and flexibility, enabling the application to use the correct role based on the ProxyRole query band for each query submitted.

Q51. A client has a healthy system but often sees delays in some queries because of workload concurrency limits. These limits have been carefully chosen, so the client needs a solution that will not modify them.

What should the Administrator use to assist this client?

- * Use a system throttle.
- * Use Query Group Viewpoint portlet to change the throttle limit temporarily.
- * Use Flex Throttle option.

Use Query Monitor Viewpoint portlet to change query workloads.

Explanation:

The Flex Throttle option allows the system to temporarily adjust workload concurrency limits based on system conditions. This provides more flexibility when handling spikes in query concurrency, without permanently modifying the established workload limits. Flex Throttle is ideal for handling temporary peaks in activity, helping to smooth out delays while keeping the core concurrency limits intact.

Option A (Use a system throttle) would enforce strict concurrency limits but doesn't provide the flexibility needed in this scenario, where the client is trying to avoid modifying existing limits.

Option B (Use Query Group Viewpoint portlet to change the throttle limit temporarily) suggests manually adjusting the throttle limit, which is not desirable in this case as the limits have been carefully chosen.

Option D (Use Query Monitor Viewpoint portlet to change query workloads) would involve changing the way queries are handled or prioritized but does not address the need to keep concurrency limits unchanged while still dealing with temporary delays.

Thus, Flex Throttle (Option C) provides the best solution to assist the client without altering the concurrency limits permanently.

Q52. A system in Viewpoint is regularly reported as being in a critical state due to a lack of available AWT. No flow control is observed on the system. The Administrator identified that this is due to a recently completed cloud migration for the system that increased the number of available AWT from 80 to 120.

Which process task is required to set up the system in Viewpoint to address this problem?

- * Configure the AWT Info data collector with the updated setting of 120 maximum AWT.
- * Increase by 40 the degraded and critical thresholds for the AWT in the system health setup portlet.
- * Update the performance data collection portlet job that collects resource usage data with the 120 maximum AWT value.
- * Adjust the system alert that has been configured for AWT to the recommended critical threshold of 92.

After the cloud migration increased the number of available AWTs from 80 to 120, the thresholds for critical and degraded states in Viewpoint are likely still based on the old maximum of 80 AWTs. Since the system is now falsely reporting critical states due to this change, the thresholds need to be updated to reflect the new maximum of 120 AWTs. Increasing the degraded and critical thresholds

by 40 (to account for the additional AWTs) will prevent unnecessary critical alerts.

Q53. On a Vantage system, the data load process has recently become much slower than normal and is now running for two hours in low concurrency. During the time the process is running, the Administrator measured the following average values from ResUsage:

System CPU busy = 75%

WIO = 25%

Inter-AMP parallelism = 90%

Average BYNET usage = 10%

Which resource is most constrained?

- * CPU
- * I/O
- * BYNET
- * Parallelism

WIO (Work in I/O) is at 25%, which indicates that the system is spending a significant amount of time waiting for I/O operations to complete. This suggests that I/O is the most constrained resource in this scenario. A high WIO typically points to I/O bottlenecks, causing delays in processes like data loading.

The other resources appear to be less constrained:

CPU: The system is busy with 75% CPU usage, which indicates the CPU has available capacity (not fully utilized).

BYNET: With 10% BYNET usage, there is no indication of network congestion.

Parallelism: Inter-AMP parallelism is at 90%, which suggests that parallel processing is functioning well and not the bottleneck.

Thus, the data load slowdown is likely due to I/O constraints, making I/O the most constrained resource in this case.

Q54. An Administrator is working through a set of tasks to set up a new Vantage system. The current task is to set up alerts to trigger BTEQ scripts to run when certain conditions are met.

Which action should an Administrator take?

- * Define a delivery setting to run SQL.
- * Ensure the SMTP host is using TLS.
- * Configure an action set specifying BTEQ.
- * Specify BTEQ in a SNMP Trap.

Action sets in Teradata Vantage allow the administrator to define a set of actions that should occur when certain conditions (alerts) are triggered. In this case, the action set should be configured to run a BTEQ script when the conditions are met. This approach ensures that specific automated responses (such as running BTEQ scripts) are executed in response to alerts.

Option A (Define a delivery setting to run SQL) refers to running SQL commands but does not involve triggering BTEQ scripts, which is what the question requires.

Option B (Ensure the SMTP host is using TLS) is related to secure email notifications and is unrelated to running BTEQ scripts.

Option D (Specify BTEQ in a SNMP Trap) refers to Simple Network Management Protocol traps, which are used for network

monitoring and not directly for triggering BTEQ scripts.

Q55. An analytics team uses a multibillion row table that is relevant to a great number of queries with different filters and joins. The Administrator needs to identify an effective strategy to collect statistics on this table.

Which statistics should be collected?

- * Full-table
- * [Dynamic AMP Sample
- * Sampled
- * Summary

Dynamic AMP Sample is an efficient method for collecting statistics on large tables. It collects sample statistics from a subset of AMPs (Access Module Processors) in the system, making it much faster and less resource-intensive than collecting full-table statistics, while still providing sufficiently accurate information for the optimizer.

Full-table statistics collection would be too resource-intensive for a multibillion-row table, potentially causing performance issues due to the size of the data.

Sampled statistics might be an option, but Dynamic AMP Sample is generally preferred because it provides a more efficient and balanced approach in large distributed systems like Teradata.

Summary statistics typically apply to aggregate data rather than large, detailed tables, and would not be sufficient for query optimization across different filters and joins.

Hence, Dynamic AMP Sample is the most effective strategy in this scenario.

Prepare Top TeraData TDVAN5 Exam Audio Study Guide Practice Questions Edition:

<https://www.exams4sures.com/TeraData/TDVAN5-practice-exam-dumps.html>