

High-performance Analytic Appliance (HANA) TEC Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

SAMPLE

- 1. If three services in a single-host system are running, what is the effective allocation limit of service 3 given a total memory allocation?**
 - A. 15 GB**
 - B. 25 GB**
 - C. 30 GB**
 - D. 35 GB**
- 2. What is the main function of the backup.log file in SAP HANA?**
 - A. Monitor system performance**
 - B. Record backup and recovery activities**
 - C. Document user access**
 - D. Track database migrations**
- 3. Who is responsible for owning an SAP HANA Deployment Infrastructure (HDI) role?**
 - A. The container-specific technical user**
 - B. The system administrator**
 - C. The database administrator**
 - D. The application developer**
- 4. What type of algorithms does the Predictive Analysis Library (PAL) contain?**
 - A. Optimization algorithms**
 - B. Predictive algorithms**
 - C. Sorting algorithms**
 - D. Data mining algorithms**
- 5. Which encryption services can you activate in SAP HANA?**
 - A. Backup**
 - B. Data**
 - C. Log**
 - D. Configuration**

- 6. Which user credentials are required during the initial run of the configuration phase for the Database Migration Option to an SAP HANA 2.0 database tenant?**
- A. The SYSTEM user of the system database**
 - B. The DBACOCKPIT user of the tenant database**
 - C. The DDIC user in client 000 of the A BAP system**
 - D. All of the above**
- 7. Which information is provided by the backup catalog in SAP HANA?**
- A. Backup scheduling**
 - B. Backup configuration**
 - C. Hardware requirements**
 - D. Data restoration guidelines**
- 8. After performing an upgrade of the primary system, what manual action must be taken for SAP HANA system replication?**
- A. Reconfigure network settings**
 - B. Perform manual synchronization steps**
 - C. Restart the primary system**
 - D. Update user permissions**
- 9. Which privilege level is typically necessary for managing multiple databases in SAP HANA?**
- A. DATABASE ADMIN**
 - B. DATABASE USER**
 - C. DATABASE MANAGER**
 - D. DATABASE DEVELOPER**
- 10. Which option of the Manage Table Distribution tool do you use to evaluate if a partitioned table needs to be repartitioned?**
- A. Analyze data distribution**
 - B. Check the number of partitions**
 - C. Review partitioning strategy**
 - D. Assess table size**

Answers

SAMPLE

1. B
2. B
3. A
4. B
5. A
6. D
7. B
8. B
9. A
10. B

SAMPLE

Explanations

SAMPLE

1. If three services in a single-host system are running, what is the effective allocation limit of service 3 given a total memory allocation?

- A. 15 GB
- B. 25 GB**
- C. 30 GB
- D. 35 GB

To determine the effective allocation limit for service 3 in a single-host system where three services are running, it's essential to understand how memory allocation works in such an environment. Typically, in systems like HANA, memory allocation is preset with certain maximum limits based on total system resources and service configuration. If the total memory allocated for the host system is 30 GB and there are three services running, the effective allocation for each service would generally depend on the distribution of this total memory. In many configurations, including HANA, the memory might be divided based on given priorities or service requirements, while also considering a buffer or overhead for the system itself. Given that the answer is 25 GB, this suggests that the other services collectively require a total of 5 GB, meaning that service 3 can utilize up to 25 GB of the total 30 GB allocation. This makes sense in practical terms, as it allows enough memory for other system operations while still providing a substantial allocation for service 3, promoting effective performance without risking resource starvation for the other tasks. In summary, the 25 GB allocation reflects a realistic model of balancing resources among services, ensuring that while one service has a higher threshold, others can still operate effectively within system limits.

2. What is the main function of the backup.log file in SAP HANA?

- A. Monitor system performance
- B. Record backup and recovery activities**
- C. Document user access
- D. Track database migrations

The backup.log file in SAP HANA serves a critical role by recording backup and recovery activities. This file is essential for maintaining a log of all operations related to data backups, ensuring that database administrators can track each backup operation carried out on the system. By having a detailed record, administrators can review the history of backups taken, identify any issues that might have occurred during the backup process, and verify successful completion. This is particularly important for restoring the database to a specific point in time, as the log provides insights into what backups are available and when they were executed. The log also aids in compliance and auditing, as it provides a comprehensive trail of data protection measures that have been implemented. This functionality is integral to an effective disaster recovery plan, allowing organizations to confidently restore their data as needed based on the information encapsulated within the backup.log file.

3. Who is responsible for owning an SAP HANA Deployment Infrastructure (HDI) role?

- A. The container-specific technical user**
- B. The system administrator
- C. The database administrator
- D. The application developer

The responsibility for owning an SAP HANA Deployment Infrastructure (HDI) role lies with the container-specific technical user. In an HDI environment, this user is critical because they are specifically designed to create and manage the various database objects within a defined container. They have the required privileges and capabilities to perform tasks related to the deployment and management of the application artifacts, ensuring that the applications function correctly within the database. The container-specific technical user facilitates the connection between the application and the database, allowing for a more streamlined development process that adheres to the principles of application lifecycle management. Their role is essential for managing aspects such as authorization and the lifecycle of database artifacts, which are central to the successful use of SAP HANA's capabilities. Other roles, such as system administrators, database administrators, and application developers, perform their specific functions but do not own the HDI role that is specifically tied to container management. System administrators focus on overall system health and access control, while database administrators typically handle broader database management tasks, and application developers work on programming and developing applications rather than directly managing the deployment infrastructure responsibilities reserved for the container-specific technical user.

4. What type of algorithms does the Predictive Analysis Library (PAL) contain?

- A. Optimization algorithms
- B. Predictive algorithms**
- C. Sorting algorithms
- D. Data mining algorithms

The Predictive Analysis Library (PAL) within HANA is specifically designed to provide various predictive algorithms that are essential for analyzing data patterns and making forecasts based on historical datasets. These algorithms are particularly tailored for tasks such as regression, classification, time series analysis, and clustering, allowing users to derive insights and predictions from their data. The focus of PAL on predictive algorithms makes it a powerful tool for organizations looking to implement analytics that can inform decision-making and strategy. By leveraging these algorithms, users can gain a deeper understanding of trends and behaviors that may not be immediately apparent from raw data alone, thus enhancing the overall analytical capabilities of the HANA environment.

5. Which encryption services can you activate in SAP HANA?

- A. Backup**
- B. Data**
- C. Log**
- D. Configuration**

In SAP HANA, the capability to activate encryption services is vital for ensuring the security and integrity of data. The correct focus here is on backup encryption. Enabling backup encryption protects the data that is stored during the backup process by encrypting it, ensuring that unauthorized access to backup files is mitigated. This is a critical aspect of data management in SAP HANA, as backups often contain sensitive information that needs to be safeguarded. Encryption services for backup are especially important for compliance with various data protection regulations and standards, making it a key component of a robust data governance strategy. This measure adds an essential layer of security, ensuring that even if backup data is intercepted or accessed by unauthorized individuals, it remains unreadable without the proper decryption keys. Other options such as data and log encryption also serve significant roles within the SAP HANA environment, but they focus on protecting data at rest and in transit, rather than specifically addressing the security of backup content. Configuration does not pertain to encryption but refers to system settings and parameters. Therefore, backup encryption is the service that stands out as actively associated with managing the security of data backups within SAP HANA.

6. Which user credentials are required during the initial run of the configuration phase for the Database Migration Option to an SAP HANA 2.0 database tenant?

- A. The SYSTEM user of the system database**
- B. The DBACOCKPIT user of the tenant database**
- C. The DDIC user in client 000 of the A BAP system**
- D. All of the above**

During the initial run of the configuration phase for the Database Migration Option to an SAP HANA 2.0 database tenant, the required credentials include the SYSTEM user of the system database as well as the DBACOCKPIT user of the tenant database, and the DDIC user in client 000 of the ABAP system. This comprehensive requirement reflects the various roles needed for successful migration and configuration. The SYSTEM user of the system database is necessary because it holds the highest privileges required to create and manage database objects. This user is crucial during configuration, as it provides access to key system-level operations and settings. The DBACOCKPIT user, specific to the tenant database, is essential for database administration tasks and provides access to monitor and manage the database environment post-migration. This user ensures that migration activities can be conducted smoothly without any permission-related issues. The DDIC user in client 000 of the ABAP system is critical for integration between the database and the ABAP environment. This user is required to handle ABAP dictionary objects and ensure that data definitions are consistent during the migration. Overall, the inclusion of all these roles signifies the complexity of the migration process, ensuring that all necessary permissions and functionalities are covered to facilitate a successful configuration to

7. Which information is provided by the backup catalog in SAP HANA?

- A. Backup scheduling
- B. Backup configuration**
- C. Hardware requirements
- D. Data restoration guidelines

The backup catalog in SAP HANA primarily provides information about backup configuration. This includes details on how the backups are set up, such as types of backups (full, incremental, differential), backup locations, retention policies, and other relevant settings that specify how the system should behave when creating and managing backups. Understanding the configuration of backups is crucial because it ensures that data protection strategies are correctly implemented, allowing for efficient recovery and minimizing potential data loss. This capability helps administrators manage and optimize the backup process, ensuring that it aligns with business requirements and system performance considerations. The other options, while important in their own right, do not accurately represent the specific focus of the backup catalog. Backup scheduling relates to when backups are executed, hardware requirements are associated with the infrastructure needed for HANA, and data restoration guidelines pertain to procedures followed during recovery, rather than the configuration itself.

8. After performing an upgrade of the primary system, what manual action must be taken for SAP HANA system replication?

- A. Reconfigure network settings
- B. Perform manual synchronization steps**
- C. Restart the primary system
- D. Update user permissions

When upgrading the primary SAP HANA system, it is essential to perform manual synchronization steps for the system replication to ensure that the secondary systems are in sync with the upgraded primary system. During the upgrade process, important changes may occur within the database structure, configurations, or even system parameters. These changes necessitate manual synchronization actions, which typically involve confirming that the secondary system recognizes and adopts these updates. If these steps are neglected, the secondary system could fall out of sync, leading to inconsistencies in the data and potentially affecting the overall reliability and performance of the replicated system. The other options, while they may seem relevant in different contexts, do not directly pertain to the specific requirements following an upgrade in SAP HANA system replication. For instance, network settings might only need adjustments if there were changes to connectivity during the upgrade, and restarting the primary system is not typically required as part of post-upgrade steps specifically related to replication. User permissions may need updates over time, but they are not directly connected to the upgrade process of system replication. Thus, performing manual synchronization steps is the critical action required after a primary system upgrade to maintain system integrity and consistency in HANA system replication.

9. Which privilege level is typically necessary for managing multiple databases in SAP HANA?

- A. DATABASE ADMIN**
- B. DATABASE USER**
- C. DATABASE MANAGER**
- D. DATABASE DEVELOPER**

The privilege level typically necessary for managing multiple databases in SAP HANA is the DATABASE ADMIN. This role encompasses a broad range of administrative permissions, allowing users to perform various critical tasks, such as creating and deleting databases, managing user access, and configuring system settings. An administrator is responsible for overseeing the overall health and performance of the databases, which requires elevated privileges that are not available to other user roles. While other roles—such as DATABASE USER, DATABASE MANAGER, and DATABASE DEVELOPER—have their own distinct functions within the SAP HANA environment, they do not possess the comprehensive control required for full database management. DATABASE USERS generally have limited access to data and cannot perform administrative tasks. DATABASE MANAGER might suggest oversight but is not a defined role in SAP HANA terminology, which can lead to confusion. DATABASE DEVELOPER focuses on the development aspect, such as creating applications and content, rather than on administrative functions. Therefore, the DATABASE ADMIN role is essential for anyone looking to manage multiple databases effectively.

10. Which option of the Manage Table Distribution tool do you use to evaluate if a partitioned table needs to be repartitioned?

- A. Analyze data distribution**
- B. Check the number of partitions**
- C. Review partitioning strategy**
- D. Assess table size**

The option to evaluate if a partitioned table needs to be repartitioned is focused on understanding the distribution of data across existing partitions. The choice that pertains to this evaluation process is assessing the number of partitions. When determining if repartitioning is necessary, checking the number of partitions allows you to identify if they are balanced or imbalanced. An imbalanced distribution of data across partitions can lead to performance issues, as some partitions may end up being much larger than others, leading to inefficiencies during query execution or data processing. In practice, if the number of partitions isn't optimal for the volume and characteristics of the data, it would indicate a need for adjustment—either consolidating smaller partitions or splitting larger ones to optimize performance. Thus, analyzing the number of partitions directly informs whether the current partitioning strategy is effective or if it requires modification, establishing this choice as the right option for evaluating the need for repartitioning.