

HOT ARA-C01 Learning Engine: SnowPro Advanced Architect Certification - High Pass-Rate Snowflake Latest ARA-C01 Dumps Free



BTW, DOWNLOAD part of ActualCollection ARA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1Zh6YdS6bCRLNExgFsdguV9jLDBKGCLHw>

Our website focus on helping candidates pass Snowflake certification exams with our Valid ARA-C01 Practice Questions and detailed test answers. The most reliable ARA-C01 dumps pdf are written by our professional IT experts who have rich experience in actual test. And you will be enjoyed one-year free updating after you make payment.

The service of ARA-C01 test guide is very prominent. It always considers the needs of customers in the development process. There are three versions of our ARA-C01 learning question, PDF, PC and APP. Each version has its own advantages. You can choose according to your needs. Of course, you can use the trial version of ARA-C01 Exam Training in advance. After you use it, you will have a more profound experience. You can choose your favorite our study materials version according to your feelings. When you use ARA-C01 test guide, you can also get our services at any time.

>> ARA-C01 Learning Engine <<

100% Pass 2026 ARA-C01 Learning Engine - SnowPro Advanced Architect Certification Latest Dumps Free

Overall, we can say that with the Snowflake ARA-C01 exam you can gain a competitive edge in your job search and advance your career in the tech industry. However, to pass the SnowPro Advanced Architect Certification (ARA-C01) exam you have to prepare well. For the quick ARA-C01 exam preparation the ARA-C01 Questions is the right choice.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q95-Q100):

NEW QUESTION # 95

Assuming all Snowflake accounts are using an Enterprise edition or higher, in which development and testing scenarios would be copying of data be required, and zero-copy cloning not be suitable? (Select TWO).

- A. Data is in a production Snowflake account that needs to be provided to Developers in a separate development/testing Snowflake account in the same cloud region.
- B. Developers create their own datasets to work against transformed versions of the live data.
- C. Production and development run in different databases in the same account, and Developers need to see production-like data but with specific columns masked.
- D. The release process requires pre-production testing of changes with data of production scale and complexity. For security reasons, pre-production also runs in the production account.
- E. Developers create their own copies of a standard test database previously created for them in the development account, for their initial development and unit testing.

Answer: A,B

Explanation:

Zero-copy cloning is a feature that allows creating a clone of a table, schema, or database without physically copying the data. Zero-copy cloning is suitable for scenarios where the cloned object needs to have the same data and metadata as the original object, and where the cloned object does not need to be modified or updated frequently. Zero-copy cloning is also suitable for scenarios where the cloned object needs to be shared within the same Snowflake account or across different accounts in the same cloud region² However, zero-copy cloning is not suitable for scenarios where the cloned object needs to have different data or metadata than the original object, or where the cloned object needs to be modified or updated frequently. Zero-copy cloning is also not suitable for scenarios where the cloned object needs to be shared across different accounts in different cloud regions. In these scenarios, copying of data would be required, either by using the COPY INTO command or by using data sharing with secure views³ The following are examples of development and testing scenarios where copying of data would be required, and zero-copy cloning would not be suitable:

Developers create their own datasets to work against transformed versions of the live data. This scenario requires copying of data because the developers need to modify the data or metadata of the cloned object to perform transformations, such as adding, deleting, or updating columns, rows, or values. Zero-copy cloning would not be suitable because it would create a read-only clone that shares the same data and metadata as the original object, and any changes made to the clone would affect the original object as well⁴ Data is in a production Snowflake account that needs to be provided to Developers in a separate development/testing Snowflake account in the same cloud region. This scenario requires copying of data because the data needs to be shared across different accounts in the same cloud region. Zero-copy cloning would not be suitable because it would create a clone within the same account as the original object, and it would not allow sharing the clone with another account. To share data across different accounts in the same cloud region, data sharing with secure views or COPY INTO command can be used⁵ The following are examples of development and testing scenarios where zero-copy cloning would be suitable, and copying of data would not be required: Production and development run in different databases in the same account, and Developers need to see production-like data but with specific columns masked. This scenario can use zero-copy cloning because the data needs to be shared within the same account, and the cloned object does not need to have different data or metadata than the original object. Zero-copy cloning can create a clone of the production database in the development database, and the clone can have the same data and metadata as the original database. To mask specific columns, secure views can be created on top of the clone, and the developers can access the secure views instead of the clone directly⁶ Developers create their own copies of a standard test database previously created for them in the development account, for their initial development and unit testing. This scenario can use zero-copy cloning because the data needs to be shared within the same account, and the cloned object does not need to have different data or metadata than the original object. Zero-copy cloning can create a clone of the standard test database for each developer, and the clone can have the same data and metadata as the original database. The developers can use the clone for their initial development and unit testing, and any changes made to the clone would not affect the original database or other clones⁷ The release process requires pre-production testing of changes with data of production scale and complexity. For security reasons, pre-production also runs in the production account. This scenario can use zero-copy cloning because the data needs to be shared within the same account, and the cloned object does not need to have different data or metadata than the original object. Zero-copy cloning can create a clone of the production database in the pre-production database, and the clone can have the same data and metadata as the original database. The pre-production testing can use the clone to test the changes with data of production scale and complexity, and any changes made to the clone would not affect the original database or the production environment⁸ Reference:

- 1: SnowPro Advanced: Architect | Study Guide ⁹
- 2: Snowflake Documentation | Cloning Overview
- 3: Snowflake Documentation | Loading Data Using COPY into a Table
- 4: Snowflake Documentation | Transforming Data During a Load
- 5: Snowflake Documentation | Data Sharing Overview
- 6: Snowflake Documentation | Secure Views
- 7: Snowflake Documentation | Cloning Databases, Schemas, and Tables
- 8: Snowflake Documentation | Cloning for Testing and Development
- 9: SnowPro Advanced: Architect | Study Guide
- 10: Cloning Overview
- 11: Loading Data Using COPY into a Table
- 12: Transforming Data During a Load
- 13: Data Sharing Overview
- 14: Secure Views
- 15: Cloning Databases, Schemas, and Tables
- 16: Cloning for Testing and Development

NEW QUESTION # 96

A Snowflake Architect Is working with Data Modelers and Table Designers to draft an ELT framework specifically for data loading

using Snowpipe. The Table Designers will add a timestamp column that Inserts the current timestamp as the default value as records are loaded into a table. The Intent is to capture the time when each record gets loaded into the table; however, when tested the timestamps are earlier than the load_time column values returned by the copy_history function or the Copy_HISTORY view (Account Usage).

Why Is this occurring?

- A. The timestamps are different because there are parameter setup mismatches. The parameters need to be realigned
- B. The Table Designer team has not used the localtimeStamp or systimestamp functions in the Snowflake copy statement.
- **C. The CURRENT_TIME is evaluated when the load operation is compiled in cloud services rather than when the record is inserted into the table.**
- D. The Snowflake timezone parameter Is different from the cloud provider's parameters causing the mismatch.

Answer: C

Explanation:

* The correct answer is D because the CURRENT_TIME function returns the current timestamp at the start of the statement execution, not at the time of the record insertion. Therefore, if the load operation takes some time to complete, the CURRENT_TIME value may be earlier than the actual load time.

* Option A is incorrect because the parameter setup mismatches do not affect the timestamp values. The parameters are used to control the behavior and performance of the load operation, such as the file format, the error handling, the purge option, etc.

* Option B is incorrect because the Snowflake timezone parameter and the cloud provider's parameters are independent of each other. The Snowflake timezone parameter determines the session timezone for displaying and converting timestamp values, while the cloud provider's parameters determine the

* physical location and configuration of the storage and compute resources.

* Option C is incorrect because the localtimeStamp and systimestamp functions are not relevant for the Snowpipe load operation. The localtimeStamp function returns the current timestamp in the session timezone, while the systimestamp function returns the current timestamp in the system timezone.

Neither of them reflect the actual load time of the records. References:

* Snowflake Documentation: Loading Data Using Snowpipe: This document explains how to use Snowpipe to continuously load data from external sources into Snowflake tables. It also describes the syntax and usage of the COPY INTO command, which supports various options and parameters to control the loading behavior.

* Snowflake Documentation: Date and Time Data Types and Functions: This document explains the different data types and functions for working with date and time values in Snowflake. It also describes how to set and change the session timezone and the system timezone.

* Snowflake Documentation: Querying Metadata: This document explains how to query the metadata of the objects and operations in Snowflake using various functions, views, and tables. It also describes how to access the copy history information using the COPY_HISTORY function or the COPY_HISTORY view.

NEW QUESTION # 97

A company is designing its serving layer for data that is in cloud storage. Multiple terabytes of the data will be used for reporting. Some data does not have a clear use case but could be useful for experimental analysis.

This experimentation data changes frequently and is sometimes wiped out and replaced completely in a few days.

The company wants to centralize access control, provide a single point of connection for the end-users, and maintain data governance.

What solution meets these requirements while MINIMIZING costs, administrative effort, and development overhead?

- A. Import the data used for reporting into a Snowflake schema with native tables. Then create views that have SELECT commands pointing to the cloud storage files for the experimentation data. Then create two different roles to match the different user personas, and grant these roles to the corresponding users.
- B. Import all the data in cloud storage to be used for reporting into a Snowflake schema with native tables. Then create two different roles with grants to the different datasets to match the different user personas, and grant these roles to the corresponding users.
- C. Import all the data in cloud storage to be used for reporting into a Snowflake schema with native tables. Then create a role that has access to this schema and manage access to the data through that role.
- **D. Import the data used for reporting into a Snowflake schema with native tables. Then create external tables pointing to the cloud storage folders used for the experimentation data. Then create two different roles with grants to the different datasets to match the different user personas, and grant these roles to the corresponding users.**

Answer: D

Explanation:

The most cost-effective and administratively efficient solution is to use a combination of native and external tables. Native tables for reporting data ensure performance and governance, while external tables allow for flexibility with frequently changing experimental data. Creating roles with specific grants to datasets aligns with the principle of least privilege, centralizing access control and simplifying user management¹².

References

*Snowflake Documentation on Optimizing Cost¹.

*Snowflake Documentation on Controlling Cost².

NEW QUESTION # 98

A group of order_admin users must delete records older than 5 years from an orders table without having DELETE privileges. The order_manager role has DELETE privileges.

How can this be achieved?

- A. Create a stored procedure that runs with owner's rights and grant usage to order_admin.
- B. Create a stored procedure with selectable caller/owner rights.
- C. Create a stored procedure with caller's rights.
- D. This is not possible without DELETE privileges.

Answer: A

Explanation:

Snowflake stored procedures can execute using either caller's rights or owner's rights. When a procedure runs with owner's rights, it executes using the privileges of the role that owns the procedure-not the role invoking it.

By creating a stored procedure owned by the order_manager role (which has DELETE privileges) and defining it to run with owner's rights, the Architect can encapsulate the business logic (delete records older than 5 years) securely (Answer C). Granting only USAGE on the procedure to order_admin allows them to execute the cleanup without directly holding DELETE privileges on the table.

Caller's rights would fail because order_admin lacks DELETE privileges. Snowflake does not support runtime selection of execution rights. This is a classic SnowPro Architect security pattern: least privilege access combined with controlled procedural execution.

NEW QUESTION # 99

What are purposes for creating a storage integration? (Choose three.)

- A. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider's key management service.
- B. Store a generated identity and access management (IAM) entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account.
- C. Avoid supplying credentials when creating a stage or when loading or unloading data.
- D. Support multiple external stages using one single Snowflake object.
- E. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.
- F. Manage credentials from multiple cloud providers in one single Snowflake object.

Answer: B,C,D

Explanation:

The purpose of creating a storage integration in Snowflake includes:

B: Store a generated identity and access management (IAM) entity for an external cloud provider - This helps in managing authentication and authorization with external cloud storage without embedding credentials in Snowflake. It supports various cloud providers like AWS, Azure, or GCP, ensuring that the identity management is streamlined across platforms.

C: Support multiple external stages using one single Snowflake object - Storage integrations allow you to set up access configurations that can be reused across multiple external stages, simplifying the management of external data integrations.

D: Avoid supplying credentials when creating a stage or when loading or unloading data - By using a storage integration, Snowflake can interact with external storage without the need to continuously manage or expose sensitive credentials, enhancing security and ease of operations.

Reference: Snowflake documentation on storage integrations, found within the SnowPro Advanced: Architect course materials.

DOWNLOAD the newest ActualCollection ARA-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1Zh6YdS6bCRLNExgFsdguV9jLDBKGCLHw>