

ARA-C01 New Real Exam | Reliable ARA-C01 Test Questions



2026 Latest PrepAwayETE ARA-C01 PDF Dumps and ARA-C01 Exam Engine Free Share: <https://drive.google.com/open?id=1BLGKFIQK0iKeVIEzqlhkANMW7MVV9C2R7>

New SnowPro Advanced Architect Certification ARA-C01 study guide and latest learning materials and practice materials have been provided for customers. PrepAwayETE is a good platform that has been providing reliable, true, updated, and free SnowPro Advanced Architect Certification ARA-C01 Exam Questions. The SnowPro Advanced Architect Certification ARA-C01 exam fee is affordable, in order to succeed in your career, you need to pass SnowPro Advanced Architect Certification exam.

Snowflake ARA-C01 certification exam consists of 100 multiple-choice questions that are designed to test the knowledge and skills of the candidates across various domains such as data modeling, data loading, performance tuning, security and access control, and data processing. ARA-C01 Exam is conducted in a timed format and candidates are expected to complete it within two hours. SnowPro Advanced Architect Certification certification exam is proctored and can be taken online from anywhere in the world.

>> ARA-C01 New Real Exam <<

ARA-C01 Dump with the Help of PrepAwayETE Exam Questions

The passing rate of our ARA-C01 training braindump is 99% which means that you almost can pass the ARA-C01 test with no doubts. The reasons why our ARA-C01 test guide's passing rate is so high are varied. That is because our test bank includes two forms and they are the PDF test questions which are selected by the senior lecturer, published authors and professional experts and the practice test software which can test your mastery degree of our ARA-C01 study question at any time. The two forms cover the syllabus of the entire ARA-C01 test. You will pass the ARA-C01 exam with it.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q26-Q31):

NEW QUESTION # 26

A user can change object parameters using which of the following roles?

- A. SYSADMIN, SECURITYADMIN
- B. ACCOUNTADMIN, USER with PRIVILEGE
- C. SECURITYADMIN, USER with PRIVILEGE
- D. ACCOUNTADMIN, SECURITYADMIN

Answer: D

NEW QUESTION # 27

An Architect entered the following commands in sequence:

□

USER1 cannot find the table.

Which of the following commands does the Architect need to run for USER1 to find the tables using the Principle of Least Privilege? (Choose two.)

- A. GRANT USAGE ON DATABASE SANDBOX TO ROLE INTERN;
- B. GRANT OWNERSHIP ON DATABASE SANDBOX TO USER INTERN;
- C. GRANT ALL PRIVILEGES ON DATABASE SANDBOX TO ROLE INTERN;
- D. GRANT USAGE ON SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;
- E. GRANT ROLE PUBLIC TO ROLE INTERN;

Answer: A,D

Explanation:

* According to the Principle of Least Privilege, the Architect should grant the minimum privileges necessary for the USER1 to find the tables in the SANDBOX database.

* The USER1 needs to have USAGE privilege on the SANDBOX database and the SANDBOX.PUBLIC schema to be able to access the tables in the PUBLIC schema. Therefore, the commands B and C are the correct ones to run.

* The command A is not correct because the PUBLIC role is automatically granted to every user and role in the account, and it does not have any privileges on the SANDBOX database by default.

* The command D is not correct because it would transfer the ownership of the SANDBOX database from the Architect to the USER1, which is not necessary and violates the Principle of Least Privilege.

* The command E is not correct because it would grant all the possible privileges on the SANDBOX database to the USER1, which is also not necessary and violates the Principle of Least Privilege.

References: : Snowflake - Principle of Least Privilege : Snowflake - Access Control Privileges : Snowflake - Public Role : Snowflake - Ownership and Grants

NEW QUESTION # 28

USERADMIN and Security administrators (i.e. users with the SECURITYADMIN role) or higher can create roles.

- A. FALSE
- B. TRUE

Answer: B

NEW QUESTION # 29

The diagram shows the process flow for Snowpipe auto-ingest with Amazon Simple Notification Service (SNS) with the following steps:

Step 1: Data files are loaded in a stage.

Step 2: An Amazon S3 event notification, published by SNS, informs Snowpipe - by way of Amazon Simple Queue Service (SQS) - that files are ready to load. Snowpipe copies the files into a queue.

Step 3: A Snowflake-provided virtual warehouse loads data from the queued files into the target table based on parameters defined in the specified pipe.

□ If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, what will happen to the pipe that references the topic to receive event messages from Amazon S3?

- A. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription to the same SNS topic and will recreate the pipe by specifying the same SNS topic name in the pipe definition.
- B. The pipe will no longer be able to receive the messages and the user must wait for 24 hours from the time when the SNS topic subscription was deleted. Pipe recreation is not required as the pipe will reuse the same subscription to the existing SNS topic after 24 hours.
- C. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription by creating a new SNS topic. Snowflake will then recreate the pipe by specifying the new SNS topic name in the pipe definition.
- D. The pipe will no longer be able to receive the messages. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition.

Answer: D

Explanation:

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, the pipe that references the topic to receive event messages from Amazon S3 will no longer be able to receive the messages. This is because the SQS subscription is the link between the SNS topic and the Snowpipe notification channel. Without the subscription, the SNS topic will not be able to send notifications to the Snowpipe queue, and the pipe will not be triggered to load the new files. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition. This will create a new notification channel and a new SQS subscription for the pipe. Alternatively, the user can also recreate the SQS subscription to the existing SNS topic and then alter the pipe to use the same SNS topic name in the pipe definition. This will also restore the notification channel and the pipe functionality. Reference:

Automating Snowpipe for Amazon S3

Enabling Snowpipe Error Notifications for Amazon SNS

HowTo: Configuration steps for Snowpipe Auto-Ingest with AWS S3 Stages

NEW QUESTION # 30

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. 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.
- 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. 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.

Answer: B,E

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
- : SnowPro Advanced: Architect | Study Guide
- : Cloning Overview
- : Loading Data Using COPY into a Table
- : Transforming Data During a Load
- : Data Sharing Overview
- : Secure Views
- : Cloning Databases, Schemas, and Tables
- : Cloning for Testing and Development

NEW QUESTION # 31

If you do not know how to pass the exam more effectively, I'll give you a suggestion is to choose a good training site. This can play a multiplier effect. PrepAwayETE site has always been committed to provide candidates with a real Snowflake ARA-C01 Certification Exam training materials. The PrepAwayETE Snowflake ARA-C01 Certification Exam software are authorized products by vendors, it is wide coverage, and can save you a lot of time and effort.

Reliable ARA-C01 Test Questions: <https://www.prepawayte.com/Snowflake/ARA-C01-practice-exam-dumps.html>

BONUS!!! Download part of PrepAwayETE ARA-C01 dumps for free: <https://drive.google.com/open?id=1BLGKFIQK0iKeVIEzqhkANMW7MVV9C2R7>