

Real ARA-C01 Exams & ARA-C01 Latest Test Dumps



ARA-C01
Snowflake SnowPro
Advanced
Architect



Certification Questions
& Exams Dumps

www.edurely.com

P.S. Free & New ARA-C01 dumps are available on Google Drive shared by ExamPrepAway: https://drive.google.com/open?id=1uz9RS82mitSCE0YAa_7aGVx3bQN7P58U

Before you choose to end your practices of the ARA-C01 study materials, the screen will display the questions you have done, which help you check again to ensure all questions of ARA-C01 practice prep are well finished. The report includes your scores of the ARA-C01 learning guide. Also, it will display how many questions of the ARA-C01 exam questions you do correctly and mistakenly. In a word, you can compensate for your weakness and change a correct review plan of the study materials.

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.

>> Real ARA-C01 Exams <<

ARA-C01 Latest Test Dumps - ARA-C01 Valid Dumps Ppt

In today's era, knowledge is becoming more and more important, and talents are becoming increasingly saturated. In such a tough situation, how can we highlight our advantages? It may be a good way to get the test ARA-C01 certification. In fact, we always will unconsciously score of high and low to measure a person's level of strength, believe that we have experienced as a child by elders inquire achievement feeling, now, we still need to face the fact. Our society needs all kinds of comprehensive talents, the ARA-C01 Study Materials can give you what you want, but not just some boring book knowledge, but flexible use of combination with the social practice.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q129-Q134):

NEW QUESTION # 129

A Snowflake Architect is designing a multiple-account design strategy. This strategy will be MOST cost-effective with which scenarios? (Select TWO).

- A. The company must use a specific network policy for certain users to allow and block given IP addresses.
- B. The company security policy mandates the use of different Active Directory instances for the development, test, and production environments.
- C. The company needs to support different role-based access control features for the development, test, and production environments.
- D. The company wants to clone a production database that resides on AWS to a development database that resides on Azure.
- E. The company needs to share data between two databases, where one must support Payment Card Industry Data Security Standard (PCI DSS) compliance but the other one does not.

Answer: C,E

Explanation:

A multiple-account design strategy is a way of organizing Snowflake accounts into logical groups based on different criteria, such as cloud provider, region, environment, or business unit. A multiple-account design strategy can help achieve various goals, such as cost optimization, performance isolation, security compliance, and data sharing¹. In this question, the scenarios that would be most cost-effective with a multiple-account design strategy are:

* The company wants to clone a production database that resides on AWS to a development database that resides on Azure. This scenario would benefit from a multiple-account design strategy because it would allow the company to leverage the cross-cloud replication feature of Snowflake, which enables replicating databases across different cloud platforms and regions. This feature can help reduce the data transfer costs and latency, as well as provide high availability and disaster recovery².

* The company security policy mandates the use of different Active Directory instances for the development, test, and production environments. This scenario would benefit from a multiple-account design strategy because it would allow the company to use different federated authentication methods for each environment, and integrate them with different Active Directory instances. This can help improve the security and governance of the access to the Snowflake accounts, as well as simplify the user management and provisioning³.

The other scenarios would not be most cost-effective with a multiple-account design strategy, because:

* The company needs to share data between two databases, where one must support Payment Card Industry Data Security Standard (PCI DSS) compliance but the other one does not. This scenario can be handled within a single Snowflake account, by using secure views and secure UDFs to mask or filter the sensitive data, and applying the appropriate roles and privileges to the users who access the data. This can help achieve the PCI DSS compliance without incurring the additional costs of managing multiple accounts⁴.

* The company needs to support different role-based access control features for the development, test, and production environments. This scenario can also be handled within a single Snowflake account, by using the native role-based access control (RBAC) features of Snowflake, such as roles, grants, and privileges, to define different access levels and permissions for each environment. This can help ensure the security and integrity of the data and the objects, as well as the separation of duties and responsibilities among the users.

* The company must use a specific network policy for certain users to allow and block given IP addresses.

This scenario can also be handled within a single Snowflake account, by using the network policy

* feature of Snowflake, which enables creating and applying network policies to restrict the IP addresses that can access the Snowflake account. This can help prevent unauthorized access and protect the data from malicious attacks.

References:

* Designing Your Snowflake Topology

* Cross-Cloud Replication

* Configuring Federated Authentication and SSO

* Using Secure Views and Secure UDFs to Comply with PCI DSS

* [Understanding Access Control in Snowflake]

* [Network Policies]

NEW QUESTION # 130

When loading data into a table that captures the load time in a column with a default value of either CURRENT_TIME () or CURRENT_TIMESTAMP() what will occur?

- A. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were read from the source.
- B. All rows loaded using a specific COPY statement will have the same timestamp value.
- C. All rows loaded using a specific COPY statement will have varying timestamps based on when the rows were inserted.
- D. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were created in the source.

Answer: B

Explanation:

According to the Snowflake documentation, when loading data into a table that captures the load time in a column with a default value of either CURRENT_TIME () or CURRENT_TIMESTAMP(), the default value is evaluated once per COPY statement, not once per row. Therefore, all rows loaded using a specific COPY statement will have the same timestamp value. This behavior ensures that the timestamp value reflects the time when the data was loaded into the table, not when the data was read from the source or created in the source. References:

* Snowflake Documentation: Loading Data into Tables with Default Values

* Snowflake Documentation: COPY INTO table

NEW QUESTION # 131

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. Production and development run in different databases in the same account, and Developers need to see production-like data but with specific columns masked.
- B. 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.
- C. Developers create their own datasets to work against transformed versions of the live data.
- D. 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.
- 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: A,E

Explanation:

<https://docs.snowflake.com/en/user-guide/tag-based-masking-policies#considerations>

NEW QUESTION # 132

A new table and streams are created with the following commands:

```
CREATE OR REPLACE TABLE LETTERS (ID INT, LETTER STRING);
```

```
CREATE OR REPLACE STREAM STREAM_1 ON TABLE LETTERS;
```

```
CREATE OR REPLACE STREAM STREAM_2 ON TABLE LETTERS APPEND_ONLY = TRUE;
```

The following operations are processed on the newly created table:

```
INSERT INTO LETTERS VALUES (1, 'A');
```

```
INSERT INTO LETTERS VALUES (2, 'B');
```

```
INSERT INTO LETTERS VALUES (3, 'C');
```

```
TRUNCATE TABLE LETTERS;
```

```
INSERT INTO LETTERS VALUES (4, 'D');
```

```
INSERT INTO LETTERS VALUES (5, 'E');
```

```
INSERT INTO LETTERS VALUES (6, 'F');
```

```
DELETE FROM LETTERS WHERE ID = 6;
```

What would be the output of the following SQL commands, in order?

```
SELECT COUNT (*) FROM STREAM_1;
```

```
SELECT COUNT (*) FROM STREAM_2;
```

- A. 4 & 6
- B. 2 & 3
- C. 4 & 3
- D. 2 & 6

Answer: C

Explanation:

In Snowflake, a stream records data manipulation language (DML) changes to its base table since the stream was created or last consumed. STREAM_1 will show all changes including the TRUNCATE operation, while STREAM_2, being APPEND_ONLY, will not show deletions like TRUNCATE. Therefore, STREAM_1 will count the three inserts, the TRUNCATE (counted as a single

