

ARA-C01 Exam New Exam Fee & Useful ARA-C01 Exam Guide Pass Success



DOWNLOAD the newest Pass4sures ARA-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1I04QG-3GSSmEBKsuKCuEJbNRPp56K_Je

If without a quick purchase process, users of our ARA-C01 quiz guide will not be able to quickly start their own review program. So, our company employs many experts to design a fast sourcing channel for our ARA-C01 exam prep. All users can implement fast purchase and use our ARA-C01 learning materials. We have specialized software to optimize the user's purchase channels, if you decide to purchase our ARA-C01 prepare questions, you can achieve the ARA-C01 exam questions content even if the update service and efficient and convenient user experience and you will pass the exam for sure.

The Snowflake ARA-C01 Exam consists of 90 multiple-choice questions that must be completed within two hours. The questions are designed to test an individual's knowledge of Snowflake's architecture, including multi-cluster warehouses, virtual warehouses, and resource management. It also covers topics such as data modeling, security, performance optimization, and data integration.

[**>> New ARA-C01 Exam Fee <<**](#)

Snowflake ARA-C01 Exam Guide | Latest ARA-C01 Exam Online

Will you feel nervous in facing the real exam? ARA-C01 Soft test engine can stimulate the real exam environment, so that you can know what the real exam is like, you're your nerves will be reduced, at the same time, your confidence will be strengthened. In addition, ARA-C01 Soft test engine can install in more than 200 computers, and it supports MS operating system, and it has two modes for practicing. We offer you free demo for ARA-C01 Soft test engine, you can have a try before buying, so that you can have a better understanding of what you are going to buy. You can enjoy free update for 365 days, and the update version for ARA-C01 exam materials will be sent to you automatically.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q225-Q230):

NEW QUESTION # 225

A DevOps team has a requirement for recovery of staging tables used in a complex set of data pipelines. The staging tables are all located in the same staging schema. One of the requirements is to have online recovery of data on a rolling 7-day basis.

After setting up the DATA_RETENTION_TIME_IN_DAYS at the database level, certain tables remain unrecoverable past 1 day. What would cause this to occur? (Choose two.)

- A. The staging tables are of the TRANSIENT type.
- B. The tables exceed the 1 TB limit for data recovery.
- C. The staging schema has not been setup for MANAGED ACCESS.
- D. The DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day.

- E. The DevOps role should be granted ALLOW_RECOVERY privilege on the staging schema.

Answer: A,D

Explanation:

- * The DATA_RETENTION_TIME_IN_DAYS parameter controls the Time Travel retention period for an object (database, schema, or table) in Snowflake. This parameter specifies the number of days for which historical data is preserved and can be accessed using Time Travel operations (SELECT, CREATE ... CLONE, UNDROP)1.
- * The requirement for recovery of staging tables on a rolling 7-day basis means that the DATA_RETENTION_TIME_IN_DAYS parameter should be set to 7 at the database level. However, this parameter can be overridden at the lower levels (schema or table) if they have a different value1.
- * Therefore, one possible cause for certain tables to remain unrecoverable past 1 day is that the DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day. This would override the database level setting and limit the Time Travel retention period for all the tables in the schema to 1 day. To fix this, the parameter should be unset or set to 7 at the schema level1. Therefore, option B is correct.
- * Another possible cause for certain tables to remain unrecoverable past 1 day is that the staging tables are of the TRANSIENT type. Transient tables are tables that do not have a Fail-safe period and can have a Time Travel retention period of either 0 or 1 day. Transient tables are suitable for temporary or intermediate data that can be easily reproduced or replicated2. To fix this, the tables should be created as permanent tables, which can have a Time Travel retention period of up to 90 days1. Therefore, option D is correct.
- * Option A is incorrect because the MANAGED ACCESS feature is not related to the data recovery requirement. MANAGED ACCESS is a feature that allows granting access privileges to objects without explicitly granting the privileges to roles. It does not affect the Time Travel retention period or the data availability3.
- * Option C is incorrect because there is no 1 TB limit for data recovery in Snowflake. The data storage size does not affect the Time Travel retention period or the data availability4.
- * Option E is incorrect because there is no ALLOW_RECOVERY privilege in Snowflake. The privilege required to perform Time Travel operations is SELECT, which allows querying historical data in tables5.

References: : Understanding & Using Time Travel : Transient Tables : Managed Access : Understanding Storage Cost : Table Privileges

NEW QUESTION # 226

Which of the below commands will use warehouse credits?

- A. SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY FLAKE_ID;
- B. SELECT MAX(FLAKE_ID) FROM SNOWFLAKE;
- C. SELECT COUNT(*) FROM SNOWFLAKE;
- D. SHOW TABLES LIKE 'SNOWFL%';

Answer: A,B,C

Explanation:

- * Warehouse credits are used to pay for the processing time used by each virtual warehouse in Snowflake. A virtual warehouse is a cluster of compute resources that enables executing queries, loading data, and performing other DML operations. Warehouse credits are charged based on the number of virtual warehouses you use, how long they run, and their size1.
- * Among the commands listed in the question, the following ones will use warehouse credits:
 - * SELECT MAX(FLAKE_ID) FROM SNOWFLAKE: This command will use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the SNOWFLAKE table and return the maximum value of the FLAKE_ID column2. Therefore, option B is correct.
 - * SELECT COUNT(*) FROM SNOWFLAKE: This command will also use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the SNOWFLAKE table and return the number of rows in the table3. Therefore, option C is correct.
 - * SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY FLAKE_ID: This command will also use warehouse credits because it is a query that requires a virtual warehouse to execute. The query will scan the SNOWFLAKE table and return the number of rows for each distinct value of the FLAKE_ID column4. Therefore, option D is correct.
- * The command that will not use warehouse credits is:
 - * SHOW TABLES LIKE 'SNOWFL%': This command will not use warehouse credits because it is a metadata operation that does not require a virtual warehouse to execute. The command will return the names of the tables that match the pattern 'SNOWFL%' in the current database and schema5. Therefore, option A is incorrect.

References: : Understanding Compute Cost : MAX Function : COUNT Function : GROUP BY Clause : SHOW TABLES

NEW QUESTION # 227

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 needs to support different role-based access control features for the development, test, and production environments.
- B. The company security policy mandates the use of different Active Directory instances for the development, test, and production environments.
- C. The company must use a specific network policy for certain users to allow and block given IP addresses.
- 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: A,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.

Reference:

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 # 228

A data share exists between a provider and a consumer account. Five tables are already shared, and the consumer role has been granted IMPORTED PRIVILEGES.

What happens if a new table is added to the provider schema?

- A. The consumer sees the table after granting IMPORTED PRIVILEGES again on the consumer side.

- B. The consumer sees the table after USAGE is granted on the database and schema and SELECT is granted on the table to the consumer database.
- **C. The consumer sees the table only after SELECT is granted on the new table to the share on the provider side.**
- D. The consumer automatically sees the new table.

Answer: C

Explanation:

In Snowflake Secure Data Sharing, adding a new table to an existing schema does not automatically make it visible to consumers. The provider must explicitly grant SELECT on the new table to the share (Answer C).

This ensures intentional and controlled data exposure.

Once the grant is applied on the provider side, consumer roles that already have IMPORTED PRIVILEGES will automatically see the new table without requiring additional grants on the consumer account. This separation of responsibilities is a core Snowflake governance principle.

This question reinforces SnowPro Architect knowledge of provider-versus-consumer responsibilities in Secure Data Sharing.

NEW QUESTION # 229

A Snowflake Architect is designing a multi-tenant application strategy for an organization in the Snowflake Data Cloud and is considering using an Account Per Tenant strategy.

Which requirements will be addressed with this approach? (Choose two.)

- A. Compute costs must be optimized.
- B. Security and Role-Based Access Control (RBAC) policies must be simple to configure.
- **C. Tenant data shape may be unique per tenant.**
- **D. Storage costs must be optimized.**
- E. There needs to be fewer objects per tenant.

Answer: C,D

Explanation:

* An Account Per Tenant strategy means creating a separate Snowflake account for each tenant (customer or business unit) of the multi-tenant application.

* This approach has some advantages and disadvantages compared to other strategies, such as Database Per Tenant or Schema Per Tenant.

* One advantage is that each tenant can have a unique data shape, meaning they can define their own tables, views, and other objects without affecting other tenants. This allows for more flexibility and customization for each tenant. Therefore, option D is correct.

* Another advantage is that storage costs can be optimized, because each tenant can use their own storage credits and manage their own data retention policies. This also reduces the risk of data spillover or cross-tenant access. Therefore, option E is correct.

* However, this approach also has some drawbacks, such as:

* It requires more administrative overhead and complexity to manage multiple accounts and their resources.

* It may not optimize compute costs, because each tenant has to provision their own warehouses and pay for their own compute credits. This may result in underutilization or overprovisioning of compute resources. Therefore, option C is incorrect.

* It may not simplify security and RBAC policies, because each account has to define its own roles, users, and privileges. This may increase the risk of human errors or inconsistencies in security configurations. Therefore, option B is incorrect.

* It may not reduce the number of objects per tenant, because each tenant still has to create their own databases, schemas, and other objects within their account. This may affect the performance and scalability of the application. Therefore, option A is incorrect.

References: : Multi-Tenant Application Strategies

NEW QUESTION # 230

.....

SnowPro Advanced Architect Certification Practice exams of Pass4sures i.e. desktop software and web-based are customizable and you can attempt them for multiple times. These practice exam save progress report of each attempt so you can assess it to find and overcome mistakes. As mentioned earlier, these SnowPro Advanced Architect Certification (ARA-C01) practice exams can be customized according to your requirements. You can change their time and numbers of SnowPro Advanced Architect Certification (ARA-C01) dumps questions as you want.

ARA-C01 Exam Guide: <https://www.pass4sures.top/SnowPro-Advanced-Certification/ARA-C01-testking-braindumps.html>

What's more, part of that Pass4sures ARA-C01 dumps now are free: <https://drive.google.com/open?id=1I04QG-3GSSmEBKsuKCuEJbNRPp56K> Je