

ARA-C01 Valid Test & ARA-C01 Cert Material & ARA-C01 Sure Pass Exam



P.S. Free & New ARA-C01 dumps are available on Google Drive shared by Pass4sureCert: <https://drive.google.com/open?id=14AR80fVcYKeIw2B3fgUFcPuQCTL676rz>

Pass4sureCert provides you with actual Snowflake ARA-C01 in PDF format, Desktop-Based Practice tests, and Web-based Practice exams. These 3 formats of Snowflake ARA-C01 exam preparation are easy to use. This is a Printable ARA-C01 PDF dumps file. The Snowflake ARA-C01 PDF dumps enables you to study without any device, as it is a portable and easily shareable format.

Snowflake ARA-C01 exam covers a wide range of topics, including data modeling, security, performance, and scalability. ARA-C01 exam also tests the candidate's knowledge of Snowflake's data warehousing capabilities, including how to design and implement a data warehouse solution using Snowflake's features. ARA-C01 Exam is designed to test a candidate's ability to design and implement complex data solutions on the Snowflake platform.

>> Trustworthy ARA-C01 Pdf <<

ARA-C01 Examcollection Dumps | ARA-C01 Valid Exam Sims

Our company is a professional certificate exam materials provider, and we have worked on this industry for years, therefore we have rich experiences. ARA-C01 exam dumps of us have questions and answers, and it will be easier for you to check the right answers

after practicing. ARA-C01 Exam Braindumps are famous for high quality, we use the skilled professionals to compile them, and the quality is guaranteed. Furthermore, our professional technicians will check the safety of our website, and we will provide you with a safe shopping environment.

Snowflake ARA-C01 Exam is a rigorous assessment of an individual's technical abilities and understanding of Snowflake's architecture. ARA-C01 exam consists of multiple-choice questions and covers a broad range of topics, including data modeling, security, performance tuning, and integration with other systems. ARA-C01 exam also tests an individual's ability to design and implement advanced Snowflake solutions that meet specific business requirements.

The SnowPro Advanced Architect Certification Exam covers a wide range of topics related to Snowflake's advanced architecture, including advanced Snowflake concepts, Snowflake security and access control, Snowflake data modeling, Snowflake data warehousing, and Snowflake performance tuning. SnowPro Advanced Architect Certification certification exam is designed to test your ability to apply this knowledge to real-world scenarios, and to demonstrate your expertise in designing and implementing Snowflake solutions that meet business needs.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q148-Q153):

NEW QUESTION # 148

What is a valid object hierarchy when building a Snowflake environment?

- A. Account --> Schema > Table --> Stage
- B. Organization --> Account --> Stage --> Table --> View
- C. Account --> Database --> Schema --> Warehouse
- D. Organization --> Account --> Database --> Schema --> Stage

Answer: D

Explanation:

Explanation

This is the valid object hierarchy when building a Snowflake environment, according to the Snowflake documentation and the web search results. Snowflake is a cloud data platform that supports various types of objects, such as databases, schemas, tables, views, stages, warehouses, and more. These objects are organized in a hierarchical structure, as follows:

* Organization: An organization is the top-level entity that represents a group of Snowflake accounts that are related by business needs or ownership. An organization can have one or more accounts, and can enable features such as cross-account data sharing, billing and usage reporting, and single sign-on across accounts12.

* Account: An account is the primary entity that represents a Snowflake customer. An account can have one or more databases, schemas, stages, warehouses, and other objects. An account can also have one or more users, roles, and security integrations. An account is associated with a specific cloud platform, region, and Snowflake edition34.

* Database: A database is a logical grouping of schemas. A database can have one or more schemas, and can store structured, semi-structured, or unstructured data. A database can also have properties such as retention time, encryption, and ownership56.

* Schema: A schema is a logical grouping of tables, views, stages, and other objects. A schema can have one or more objects, and can define the namespace and access control for the objects. A schema can also have properties such as ownership and default warehouse .

* Stage: A stage is a named location that references the files in external or internal storage. A stage can be used to load data into Snowflake tables using the COPY INTO command, or to unload data from Snowflake tables using the COPY INTO LOCATION command. A stage can be created at the account, database, or schema level, and can have properties such as file format, encryption, and credentials .

The other options listed are not valid object hierarchies, because they either omit or misplace some objects in the structure. For example, option A omits the organization level and places the warehouse under the schema level, which is incorrect. Option C omits the organization, account, and stage levels, and places the table under the schema level, which is incorrect. Option D omits the database level and places the stage and table under the account level, which is incorrect.

References:

- * Snowflake Documentation: Organizations
- * Snowflake Blog: Introducing Organizations in Snowflake
- * Snowflake Documentation: Accounts
- * Snowflake Blog: Understanding Snowflake Account Structures
- * Snowflake Documentation: Databases
- * Snowflake Blog: How to Create a Database in Snowflake
- * [Snowflake Documentation: Schemas]
- * [Snowflake Blog: How to Create a Schema in Snowflake]

* [Snowflake Documentation: Stages]

* [Snowflake Blog: How to Use Stages in Snowflake]

NEW QUESTION # 149

What are some of the characteristics of result set caches? (Choose three.)

- A. The result set cache is not shared between warehouses.
- B. **Snowflake persists the data results for 24 hours.**
- C. Time Travel queries can be executed against the result set cache.
- D. **Each time persisted results for a query are used, a 24-hour retention period is reset.**
- E. The data stored in the result cache will contribute to storage costs.
- F. **The retention period can be reset for a maximum of 31 days.**

Answer: B,D,F

Explanation:

Explanation

Comprehensive and Detailed Explanation: According to the SnowPro Advanced: Architect documents and learning resources, some of the characteristics of result set caches are:

- * Snowflake persists the data results for 24 hours. This means that the result set cache holds the results of every query executed in the past 24 hours, and can be reused if the same query is submitted again and the underlying data has not changed1.
- * Each time persisted results for a query are used, a 24-hour retention period is reset. This means that the result set cache extends the lifetime of the results every time they are reused, up to a maximum of 31 days from the date and time that the query was first executed1.
- * The retention period can be reset for a maximum of 31 days. This means that the result set cache will purge the results after 31 days, regardless of whether they are reused or not. After 31 days, the next time the query is submitted, a new result is generated and persisted1.

The other options are incorrect because they are not characteristics of result set caches. Option A is incorrect because Time Travel queries cannot be executed against the result set cache. Time Travel queries use the AS OF clause to access historical data that is stored in the storage layer, not the result set cache2. Option D is incorrect because the data stored in the result set cache does not contribute to storage costs. The result set cache is maintained by the service layer, and does not incur any additional charges1. Option F is incorrect because the result set cache is shared between warehouses. The result set cache is available across virtual warehouses, so query results returned to one user are available to any other user on the system who executes the same query, provided the underlying data has not changed1. References: Using Persisted Query Results | Snowflake Documentation, Time Travel | Snowflake Documentation

NEW QUESTION # 150

Role A has the following permissions:

- . USAGE on db1
- . USAGE and CREATE VIEW on schema1 in db1
- . SELECT on table1 in schema1

Role B has the following permissions:

- . USAGE on db2
- . USAGE and CREATE VIEW on schema2 in db2
- . SELECT on table2 in schema2

A user has Role A set as the primary role and Role B as a secondary role.

What command will fail for this user?

- A. **use database db2;**
use schema schema2;
create view v2 as select * from db1.schema1.table1;
- B. **use database db1;**
use schema schema1;
create view v1 as select * from db2.schema2.table2;
- C. **use database db1;**
use schema schema1;
select * from db2.schema2.table2;

- D. use database db2;
use schema schema2;
select * from db1.schema1.table1 union select * from table2;

Answer: A

Explanation:

This command will fail because while the user has USAGE permission on db2 and schema2 through Role B, and can create a view in schema2, they do not have SELECT permission on db1.schema1.table1 with Role B.

Since Role A, which has SELECT permission on db1.schema1.table1, is not the currently active role when the view v2 is being created in db2.schema2, the user does not have the necessary permissions to read from db1.schema1.table1 to create the view. Snowflake's security model requires that the active role have all necessary permissions to execute the command.

NEW QUESTION # 151

When using the Snowflake Connector for Kafka, what data formats are supported for the messages? (Choose two.)

- A. Parquet
- B. XML
- C. Avro
- D. JSON
- E. CSV

Answer: C,D

Explanation:

The data formats that are supported for the messages when using the Snowflake Connector for Kafka are Avro and JSON. These are the two formats that the connector can parse and convert into Snowflake table rows. The connector supports both schemaless and schematized JSON, as well as Avro with or without a schema registry1. The other options are incorrect because they are not supported data formats for the messages. CSV, XML, and Parquet are not formats that the connector can parse and convert into Snowflake table rows. If the messages are in these formats, the connector will load them as VARIANT data type and store them as raw strings in the table2. References: Snowflake Connector for Kafka | Snowflake Documentation, Loading Protobuf Data using the Snowflake Connector for Kafka | Snowflake Documentation

NEW QUESTION # 152

A company is following the Data Mesh principles, including domain separation, and chose one Snowflake account for its data platform.

An Architect created two data domains to produce two data products. The Architect needs a third data domain that will use both of the data products to create an aggregate data product. The read access to the data products will be granted through a separate role. Based on the Data Mesh principles, how should the third domain be configured to create the aggregate product if it has been granted the two read roles?

- A. Request that the two data domains share data using the Data Exchange.
- B. Use secondary roles for all users.
- C. Request a technical ETL user with the sysadmin role.
- D. Create a hierarchy between the two read roles.

Answer: A

Explanation:

In the scenario described, where a third data domain needs access to two existing data products in a Snowflake account structured according to Data Mesh principles, the best approach is to utilize Snowflake's Data Exchange functionality. Option D is correct as it facilitates the sharing and governance of data across different domains efficiently and securely. Data Exchange allows domains to publish and subscribe to live data products, enabling real-time data collaboration and access management in a governed manner. This approach is in line with Data Mesh principles, which advocate for decentralized data ownership and architecture, enhancing agility and scalability across the organization. References:

* Snowflake Documentation on Data Exchange

* Articles on Data Mesh Principles in Data Management

NEW QUESTION # 153

• • • • •

ARA-C01 Examcollection Dumps: <https://www.pass4surecert.com/Snowflake/ARA-C01-practice-exam-dumps.html>

BTW, DOWNLOAD part of Pass4sureCert ARA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=14AR80fVcYKelw2B3fgUFcPuQCTL676rz>