

New Snowflake ARA-C01 Practice Materials & New ARA-C01 Practice Questions

Snowflake ARA-C01 Practice Questions

SnowPro Advanced Architect Certification

Order our ARA-C01 Practice Questions Today and Get Ready to Pass with Flying Colors!



ARA-C01 Practice Exam Features | QuestionsTube

- Latest & Updated Exam Questions
- Subscribe to FREE Updates
- Both PDF & Exam Engine
- Download Directly Without Waiting

<https://www.questiontube.com/exam/ara-c01/>

At QuestionsTube, you can read ARA-C01 free demo questions in pdf file, so you can check the questions and answers before deciding to download the Snowflake ARA-C01 practice questions. These free demo questions are parts of the ARA-C01 exam questions. Download and read them carefully, you will find that the ARA-C01 test questions of QuestionsTube will be your great learning materials online. Share some ARA-C01 exam online questions below.

1. An Architect needs to allow a user to create a database from an inbound share.

P.S. Free 2026 Snowflake ARA-C01 dumps are available on Google Drive shared by ITexamReview:
https://drive.google.com/open?id=1MhMIj49eR1PBpCy--xPsGgc_a9Y4o49C

We will offer the preparation for the ARA-C01 training materials, we will also provide you the guide in the process of using. The materials of the exam dumps offer you enough practice for the ARA-C01 as well as the knowledge points of the ARA-C01 exam, the exam will become easier. If you are interested in the ARA-C01 training materials, free demo is offered, you can have a try. And the downloading link will send to you within ten minutes, so you can start your preparation as quickly as possible. In fact, the outcome of the ARA-C01 Exam most depends on the preparation for the ARA-C01 training materials. With the training materials, you can make it.

To pass the Snowflake ARA-C01 Exam, candidates need to demonstrate proficiency in various domains, including Snowflake architecture, security, data integration, performance tuning, and optimization. ARA-C01 exam is designed to evaluate the candidate's ability to design, implement and manage complex Snowflake environments that meet the business requirements of organizations. SnowPro Advanced Architect Certification certification is beneficial for architects, data engineers, and data scientists who work with Snowflake Data Cloud.

>> **New Snowflake ARA-C01 Practice Materials** <<

100% Pass Quiz Snowflake - ARA-C01 - Pass-Sure New SnowPro Advanced

Architect Certification Practice Materials

Our customer service is available all day, and your problems can be solved efficiently at any time. Last but not least, we can guarantee the security of the purchase process of ARA-C01 test questions and the absolute confidentiality of customer information. You do not have to worry about these issues, because we know that this is a basic condition for us to establish a good business model. At the same time, if you want to continue learning, ARA-C01 Test Torrent will provide you with the benefits of free updates within one year and a discount of more than one year.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q164-Q169):

NEW QUESTION # 164

A company has several sites in different regions from which the company wants to ingest data. Which of the following will enable this type of data ingestion?

- A. The company must have a Snowflake account in each cloud region to be able to ingest data to that account.
- B. The company must replicate data between Snowflake accounts.
- C. The company should provision a reader account to each site and ingest the data through the reader accounts.
- **D. The company should use a storage integration for the external stage.**

Answer: D

Explanation:

This is the correct answer because it allows the company to ingest data from different regions using a storage integration for the external stage. A storage integration is a feature that enables secure and easy access to files in external cloud storage from Snowflake. A storage integration can be used to create an external stage, which is a named location that references the files in the external storage. An external 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 storage integration can support multiple regions and cloud platforms, as long as the external storage service is compatible with Snowflake12.

Reference:

Snowflake Documentation: Storage Integrations

Snowflake Documentation: External Stages

NEW QUESTION # 165

Which columns can be included in an external table schema? (Select THREE).

- **A. METADATA\$FILENAME**
- **B. METADATA\$FILE_ROW_NUMBER**
- C. METADATA\$EXTERNAL_TABLE_PARTITION
- D. METADATA\$ISUPDATE
- **E. VALUE**
- F. METADATA\$ROW_ID

Answer: A,B,E

Explanation:

External tables in Snowflake expose a combination of user-defined columns and system-generated metadata columns. The VALUE column is commonly used to store semi-structured data (such as JSON or Avro records) read directly from external storage (Answer A).

Snowflake also provides metadata columns that describe the source file. METADATA\$FILENAME identifies the name of the file from which a given row was read (Answer D), and METADATA\$FILE_ROW_NUMBER indicates the row number within that file (Answer E). These columns are frequently used for auditing, debugging, and data lineage tracking.

METADATA\$ROW_ID and METADATA\$ISUPDATE are associated with streams and change tracking, not external tables.

METADATA\$EXTERNAL_TABLE_PARTITION is not a valid selectable column in the external table schema definition. This question reinforces SnowPro Architect knowledge of how Snowflake represents external data and exposes file-level metadata for data lake architectures.

NEW QUESTION # 166

An Architect is designing a solution that will be used to process changed records in an orders table.

Newly-inserted orders must be loaded into the `f_orders` fact table, which will aggregate all the orders by multiple dimensions (time, region, channel, etc.). Existing orders can be updated by the sales department within 30 days after the order creation. In case of an order update, the solution must perform two actions:

1. Update the order in the `f_ORDERS` fact table.
2. Load the changed order data into the special table `ORDER_REPAIRS`.

This table is used by the Accounting department once a month. If the order has been changed, the Accounting team needs to know the latest details and perform the necessary actions based on the data in the `order_repairs` table.

What data processing logic design will be the MOST performant?

- A. Use two streams and one task.
- B. Use one stream and one task.
- C. Use one stream and two tasks.
- D. Use two streams and two tasks.

Answer: C

Explanation:

The most performant design for processing changed records, considering the need to both update records in the `f_orders` fact table and load changes into the `order_repairs` table, is to use one stream and two tasks. The stream will monitor changes in the orders table, capturing both inserts and updates. The first task would apply these changes to the `f_orders` fact table, ensuring all dimensions are accurately represented. The second task would use the same stream to insert relevant changes into the `order_repairs` table, which is critical for the Accounting department's monthly review. This method ensures efficient processing by minimizing the overhead of managing multiple streams and synchronizing between them, while also allowing specific tasks to optimize for their target operations. References: Snowflake's documentation on streams and tasks for handling data changes efficiently.

NEW QUESTION # 167

Which feature provides the capability to define an alternate cluster key for a table with an existing cluster key?

- A. Materialized view
- B. Search optimization
- C. Result cache
- D. External table

Answer: A

Explanation:

A materialized view is a feature that provides the capability to define an alternate cluster key for a table with an existing cluster key. A materialized view is a pre-computed result set that is stored in Snowflake and can be queried like a regular table. A materialized view can have a different cluster key than the base table, which can improve the performance and efficiency of queries on the materialized view. A materialized view can also support aggregations, joins, and filters on the base table data. A materialized view is automatically refreshed when the underlying data in the base table changes, as long as the `AUTO_REFRESH` parameter is set to `true`.

Reference:

Materialized Views | Snowflake Documentation

NEW QUESTION # 168

A company's daily Snowflake workload consists of a huge number of concurrent queries triggered between 9pm and 11pm. At the individual level, these queries are smaller statements that get completed within a short time period.

What configuration can the company's Architect implement to enhance the performance of this workload?

(Choose two.)

- A. Reduce the amount of data that is being processed through this workload.
- B. Increase the size of the virtual warehouse to size X-Large.
- C. Enable a multi-clustered virtual warehouse in maximized mode during the workload duration.
- D. Set the connection timeout to a higher value than its default.
- E. Set the `MAX_CONCURRENCY_LEVEL` to a higher value than its default value of 8 at the virtual warehouse level.

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

DOWNLOAD the newest ITexamReview ARA-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1MhMlj49eR1PBpCy--xPsGgc_a9Y4o49C