

Reliable ARA-C01 Braindumps Free | New ARA-C01 Test Notes



P.S. Free & New ARA-C01 dumps are available on Google Drive shared by Easy4Engine: <https://drive.google.com/open?id=13vOoLeoAChWk1xz9ER5BvPo-JVKI3P1m>

What is more difficult is not only passing the Snowflake ARA-C01 Certification Exam, but the acute anxiety and the excessive burden also make the candidate nervous to qualify for the SnowPro Advanced Architect Certification certification. If you are going through the same tough challenge, do not worry because Snowflake is here to assist you.

It was never so easy to make your way to the world's most rewarding professional qualification as it has become now! Easy4Engine' ARA-C01 practice test questions answers are the best option to secure your success in just one go. You can easily answer all exam questions by doing our ARA-C01 exam dumps repeatedly. For further sharpening your skills, practice mock tests using our ARA-C01 Brain Dumps Testing Engine software and overcome your fear of failing exam. Our SnowPro Advanced Architect Certification dumps are the most trustworthy, reliable and the best helpful study content that will prove the best alternative to your time and money.

>> Reliable ARA-C01 Braindumps Free <<

Free PDF Quiz ARA-C01 - Authoritative Reliable SnowPro Advanced Architect Certification Braindumps Free

Many customers want to check the content and quality of our ARA-C01 exam braindumps. So we developed trial versions for you. After you have used a trial version, you will have an overview of the content of the ARA-C01 simulating exam. This is enough to convince you that this is a product with high quality. If you are sure that you want this product, but we are not sure which version to buy, we can let you try multiple versions of ARA-C01 learning guide. And there are three varied versions on our website.

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.

Snowflake SnowPro Advanced Architect Certification Sample Questions

(Q179-Q184):

NEW QUESTION # 179

A company has an external vendor who puts data into Google Cloud Storage. The company's Snowflake account is set up in Azure. What would be the MOST efficient way to load data from the vendor into Snowflake?

- A. Create an external stage on Google Cloud Storage and use the external table to load the data into Snowflake.
- B. Create a Snowflake Account in the Google Cloud Platform (GCP), ingest data into this account and use data replication to move the data from GCP to Azure.
- C. Copy the data from Google Cloud Storage to Azure Blob storage using external tools and load data from Blob storage to Snowflake.
- D. Ask the vendor to create a Snowflake account, load the data into Snowflake and create a data share.

Answer: A

Explanation:

The most efficient way to load data from the vendor into Snowflake is to create an external stage on Google Cloud Storage and use the external table to load the data into Snowflake (Option B). This way, you can avoid copying or moving the data across different cloud platforms, which can incur additional costs and latency.

You can also leverage the external table feature to query the data directly from Google Cloud Storage without loading it into Snowflake tables, which can save storage space and improve performance. Option A is not efficient because it requires the vendor to create a Snowflake account and a data share, which can be complicated and costly. Option C is not efficient because it involves copying the data from Google Cloud Storage to Azure Blob storage using external tools, which can be slow and expensive. Option D is not efficient because it requires creating a Snowflake account in the Google Cloud Platform (GCP), ingesting data into this account, and using data replication to move the data from GCP to Azure, which can be complex and time-consuming. References: The answer can be verified from Snowflake's official documentation on external stages and external tables available on their website. Here are some relevant links:

- * [Using External Stages | Snowflake Documentation](#)
- * [Using External Tables | Snowflake Documentation](#)
- * [Loading Data from a Stage | Snowflake Documentation](#)

NEW QUESTION # 180

A company's Architect needs to find an efficient way to get data from an external partner, who is also a Snowflake user. The current solution is based on daily JSON extracts that are placed on an FTP server and uploaded to Snowflake manually. The files are changed several times each month, and the ingestion process needs to be adapted to accommodate these changes.

What would be the MOST efficient solution?

- A. Ask the partner to create a share and add the company's account.
- B. Keep the current structure but request that the partner stop changing files, instead only appending new files.
- C. Ask the partner to use the data lake export feature and place the data into cloud storage where Snowflake can natively ingest it (schema-on-read).
- D. Ask the partner to set up a Snowflake reader account and use that account to get the data for ingestion.

Answer: A

Explanation:

The most efficient solution is to ask the partner to create a share and add the company's account (Option A). This way, the company can access the live data from the partner without any data movement or manual intervention. Snowflake's secure data sharing feature allows data providers to share selected objects in a database with other Snowflake accounts. The shared data is read-only and does not incur any storage or compute costs for the data consumers. The data consumers can query the shared data directly or create local copies of the shared objects in their own databases. Option B is not efficient because it involves using the data lake export feature, which is intended for exporting data from Snowflake to an external data lake, not for importing data from another Snowflake account. The data lake export feature also requires the data provider to create an external stage on cloud storage and use the COPY INTO <location> command to export the data into parquet files. The data consumer then needs to create an external table or a file format to load the data from the cloud storage into Snowflake. This process can be complex and costly, especially if the data changes frequently. Option C is not efficient because it does not solve the problem of manual data ingestion and adaptation. Keeping the current structure of daily JSON extracts on an FTP server and requesting the partner to stop changing files, instead only appending new files, does not improve the efficiency or reliability of the data ingestion process. The company still needs to upload the data to Snowflake manually and deal with any schema changes or data quality issues. Option D is not efficient because it requires the partner to set up a Snowflake reader account and use that account to get the data for ingestion. A reader account is a special

type of account that can only consume data from the provider account that created it. It is intended for data consumers who are not Snowflake customers and do not have a licensing agreement with Snowflake. A reader account is not suitable for data ingestion from another Snowflake account, as it does not allow uploading, modifying, or unloading data. The company would need to use external tools or interfaces to access the data from the reader account and load it into their own account, which can be slow and expensive. Reference: The answer can be verified from Snowflake's official documentation on secure data sharing, data lake export, and reader accounts available on their website. Here are some relevant links:

[Introduction to Secure Data Sharing | Snowflake Documentation](#)

[Data Lake Export Public Preview Is Now Available on Snowflake | Snowflake Blog](#) [Managing Reader Accounts | Snowflake Documentation](#)

NEW QUESTION # 181

Which steps are recommended best practices for prioritizing cluster keys in Snowflake? (Choose two.)

- A. Choose lower cardinality columns to support clustering keys and cost effectiveness.
- B. Choose cluster columns that are most actively used in selective filters.
- C. Choose columns that are frequently used in join predicates.
- D. Choose cluster columns that are actively used in the GROUP BY clauses.
- E. Choose TIMESTAMP columns with nanoseconds for the highest number of unique rows.

Answer: B,C

Explanation:

According to the Snowflake documentation, the best practices for choosing clustering keys are:

Choose columns that are frequently used in join predicates. This can improve the join performance by reducing the number of micro-partitions that need to be scanned and joined.

Choose columns that are most actively used in selective filters. This can improve the scan efficiency by skipping micro-partitions that do not match the filter predicates.

Avoid using low cardinality columns, such as gender or country, as clustering keys. This can result in poor clustering and high maintenance costs.

Avoid using TIMESTAMP columns with nanoseconds, as they tend to have very high cardinality and low correlation with other columns. This can also result in poor clustering and high maintenance costs.

Avoid using columns with duplicate values or NULLs, as they can cause skew in the clustering and reduce the benefits of pruning. Cluster on multiple columns if the queries use multiple filters or join predicates. This can increase the chances of pruning more micro-partitions and improve the compression ratio.

Clustering is not always useful, especially for small or medium-sized tables, or tables that are not frequently queried or updated.

Clustering can incur additional costs for initially clustering the data and maintaining the clustering over time.

Reference:

[Clustering Keys & Clustered Tables | Snowflake Documentation](#)

[\[Considerations for Choosing Clustering for a Table | Snowflake Documentation\]](#)

NEW QUESTION # 182

Which Snowflake data modeling approach is designed for BI queries?

- A. Star schema
- B. 3 NF
- C. Snowflake schema
- D. Data Vault

Answer: A

Explanation:

A star schema is a Snowflake data modeling approach that is designed for BI queries. A star schema is a type of dimensional modeling that organizes data into fact tables and dimension tables. A fact table contains the measures or metrics of the business process, such as sales amount, order quantity, or profit margin. A dimension table contains the attributes or descriptors of the business process, such as product name, customer name, or order date. A star schema is called so because it resembles a star, with one fact table in the center and multiple dimension tables radiating from it. A star schema can improve the performance and simplicity of BI queries by reducing the number of joins, providing fast access to aggregated data, and enabling intuitive query syntax. A star schema can also support various types of analysis, such as trend analysis, slice and dice, drill down, and roll up.

References:

* Snowflake Documentation: Dimensional Modeling

* Snowflake Documentation: Star Schema

NEW QUESTION # 183

An Architect uses COPY INTO with the ON_ERROR=SKIP_FILE option to bulk load CSV files into a table called TABLEA, using its table stage. One file named file5.csv fails to load. The Architect fixes the file and re-loads it to the stage with the exact same file name it had previously.

Which commands should the Architect use to load only file5.csv file from the stage? (Choose two.)

- A. COPY INTO tablea FROM @%tablea NEW_FILES_ONLY = TRUE;
- B. COPY INTO tablea FROM @%tablea MERGE = TRUE;
- C. COPY INTO tablea FROM @%tablea RETURN_FAILED_ONLY = TRUE;
- D. COPY INTO tablea FROM @%tablea FORCE = TRUE;
- E. COPY INTO tablea FROM @%tablea;
- F. COPY INTO tablea FROM @%tablea FILES = ('file5.csv');

Answer: E,F

Explanation:

* Option A (RETURN_FAILED_ONLY) will only load files that previously failed to load. Since file5.csv already exists in the stage with the same name, it will not be considered a new file and will not be loaded.

* Option D (FORCE) will overwrite any existing data in the table. This is not desired as we only want to load the data from file5.csv.

* Option E (NEW_FILES_ONLY) will only load files that have been added to the stage since the last COPY command. This will not work because file5.csv was already in the stage before it was fixed.

* Option F (MERGE) is used to merge data from a stage into an existing table, creating new rows for any data not already present. This is not needed in this case as we simply want to load the data from file5.csv.

Therefore, the architect can use either COPY INTO tablea FROM @%tablea or COPY INTO tablea FROM @%tablea FILES = ('file5.csv') to load only file5.csv from the stage. Both options will load the data from the specified file without overwriting any existing data or requiring additional configuration

NEW QUESTION # 184

.....

About the oncoming ARA-C01 exam, every exam candidates are wishing to utilize all intellectual and technical skills to solve the obstacles ahead of them to go as well as it possibly could. So the pending exam causes a panic among the exam candidates. The ARA-C01 exam prepare of our website is completed by experts who has a good understanding of real exams and have many years of experience writing ARA-C01 Study Materials. They know very well what candidates really need most when they prepare for the exam. They also understand the real exam situation very well. So they compiled ARA-C01 exam prepare that they hope to do their utmost to help candidates pass the exam and get what job they want.

New ARA-C01 Test Notes: <https://www.easy4engine.com/ARA-C01-test-engine.html>

- ARA-C01 Examcollection □ ARA-C01 Exam Flashcards □ Printable ARA-C01 PDF ■ Search for ➤ ARA-C01 □ and download exam materials for free through □ www.vce4dumps.com □ □ ARA-C01 Latest Study Notes
- Test ARA-C01 King □ ARA-C01 Exam Flashcards □ ARA-C01 Latest Exam Notes □ Enter (www.pdfvce.com) and search for ✓ ARA-C01 □✓□ to download for free □Printable ARA-C01 PDF
- Explore www.dumpsmaterials.com's Top Three Formats for Snowflake ARA-C01 Exam □ Search for [ARA-C01] and easily obtain a free download on ▶ www.dumpsmaterials.com ▲ □Latest ARA-C01 Test Cram
- ARA-C01 Latest Exam Notes □ ARA-C01 Latest Study Notes □ ARA-C01 Latest Exam Notes □ Search for (ARA-C01) and easily obtain a free download on ▷ www.pdfvce.com □ □Pass ARA-C01 Guaranteed
- ARA-C01 Interactive Questions □ ARA-C01 Related Exams □ Exam ARA-C01 Objectives Pdf □ The page for free download of ➡ ARA-C01 □ on ➡ www.pdfdumps.com □ will open immediately □ ARA-C01 Latest Exam Notes
- 2026 Snowflake Unparalleled Reliable ARA-C01 Braindumps Free Pass Guaranteed Quiz □ Search on ✓ www.pdfvce.com □✓□ for [ARA-C01] to obtain exam materials for free download □Exam ARA-C01 Simulator
- Do You Want To Pass Snowflake ARA-C01 Exam Successfully And Effectively □ Open { www.practicevce.com } enter (ARA-C01) and obtain a free download □ARA-C01 Latest Test Dumps
- Explore Pdfvce's Top Three Formats for Snowflake ARA-C01 Exam □ 《 www.pdfvce.com 》 is best website to obtain { ARA-C01 } for free download □ARA-C01 Latest Test Dumps

DOWNLOAD the newest Easy4Engine ARA-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=13vOoLeoAChWk1xz9ER5BvPo-JVKI3P1m>