Pass Guaranteed Quiz Snowflake - ARA-C01 - Perfect SnowPro Advanced Architect Certification Trustworthy Exam Content

BTW, DOWNLOAD part of CramPDF ARA-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1uioX0Bg2G7g7wKt_YUR0lARRqf2hZkNS

First and foremost, we have high class operation system so we can assure you that you can start to prepare for the ARA-C01 exam with our study materials only 5 to 10 minutes after payment. Fortunately, you need not to worry about this sort of question any more, since you can find the best solution in this website--our ARA-C01 Training Materials. With our continued investment in technology, people and facilities, the future of our company has never looked so bright. There are so many advantages of our ARA-C01 practice test and I would like to give you a brief introduction now.

Snowflake ARA-C01 (SnowPro Advanced Architect Certification) Certification Exam is a highly specialized certification program designed for professionals who want to demonstrate their advanced technical knowledge and skills in designing, deploying, and managing Snowflake solutions. SnowPro Advanced Architect Certification certification exam is designed to assess the knowledge and skills required to build, design, and manage a Snowflake architecture at an expert level.

>> ARA-C01 Trustworthy Exam Content <<

Snowflake - ARA-C01 –High-quality Trustworthy Exam Content

If you are ready for the ARA-C01 exam for a long time, but lack of a set of suitable ARA-C01 learning materials, I will tell you that you are so lucky to enter this page. We are such ARA-C01 exam questions that you can use our products to prepare the exam and obtain your dreamed ARA-C01certificates. We all know that if you desire a better job post, you have to be equipped with appropriate professional quality and an attitude of keeping forging ahead. And we can give what you need!

The SnowPro Advanced Architect Certification is a valuable asset for individuals who want to advance their careers in the data warehousing and data analytics fields. SnowPro Advanced Architect Certification certification demonstrates that the holder has a deep understanding of Snowflake architecture and can design, build, and manage high-performance Snowflake data warehouses and data analytics solutions. SnowPro Advanced Architect Certification certification also shows that the holder is capable of optimizing performance, ensuring security, and managing the administration of Snowflake environments.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q45-Q50):

NEW QUESTION # 45

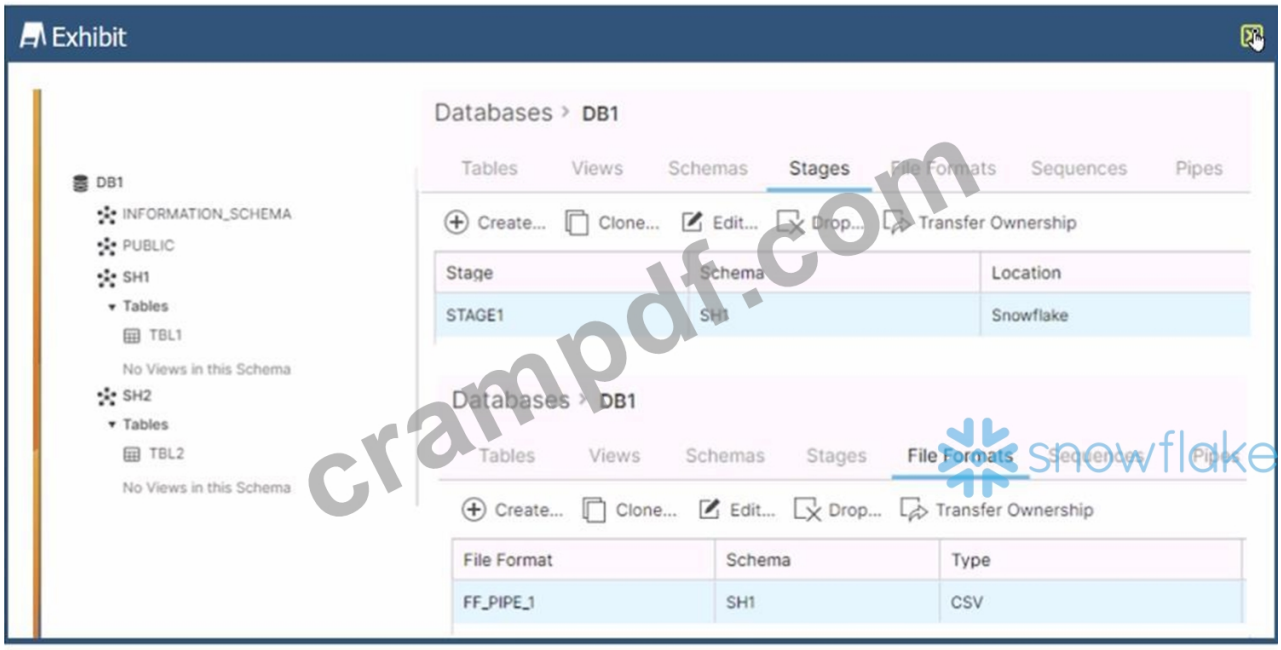

Based on the architecture in the image, how can the data from DB1 be copied into TBL2? (Select TWO).

- A. A white background with black text Description automatically generated

```
use database DB1;
use schema SH2;
copy into TBL2
  from @STAGE1
 file_format = (format_name = FF_PIPE_1);
```

- B. A white background with black text Description automatically generated
- C. A close-up of a computer code Description automatically generated

```
use database DB1;
use schema SH2;
copy into DB1.SH2.TBL2
  from @STAGE1
 file_format = (format_name = DB1.SH1.FF_PIPE_1);
```

- D. A computer code with black text Description automatically generated

```
use database DB1;
use schema SH2;
copy into DB1.SH2.TBL2
  from @DB1.SH1.STAGE1
 file_format = (format_name = FF_PIPE_1);
```

- E. A close-up of a computer code Description automatically generated

Answer: A,E

Explanation:
* The architecture in the image shows a Snowflake data platform with two databases, DB1 and DB2, and two schemas, SH1 and SH2. DB1 contains a table TBL1 and a stage STAGE1. DB2 contains a table TBL2. The image also shows a snippet of code written in SQL language that copies data from STAGE1 to TBL2 using a file format FF PIPE 1.
* To copy data from DB1 to TBL2, there are two possible options among the choices given:
* Option B: Use a named external stage that references STAGE1. This option requires creating an external stage object in DB2.SH2 that points to the same location as STAGE1 in DB1.SH1. The external stage can be created using the CREATE STAGE command with the URL parameter specifying the location of STAGE11. For example:

SQLAI-generated code. Review and use carefully. More info on FAQ.
use database DB2;
use schema SH2;
create stage EXT_STAGE1
url = @DB1.SH1.STAGE1;
* Then, the data can be copied from the external stage to TBL2 using the COPY INTO command with the FROM parameter specifying the external stage name and the FILE FORMAT parameter specifying the file format name2. For example:
SQLAI-generated code. Review and use carefully. More info on FAQ.
copy into TBL2
from @EXT_STAGE1
file format = (format name = DB1.SH1.FF PIPE 1);
* Option E: Use a cross-database query to select data from TBL1 and insert into TBL2. This option requires using the INSERT INTO command with the SELECT clause to query data from TBL1 in DB1.SH1 and insert it into TBL2 in DB2.SH2. The query must use the fully-qualified names of the tables, including the database and schema names3. For example:
SQLAI-generated code. Review and use carefully. More info on FAQ.
use database DB2;
use schema SH2;
insert into TBL2
select * from DB1.SH1.TBL1;
* The other options are not valid because:
* Option A: It uses an invalid syntax for the COPY INTO command. The FROM parameter cannot specify a table name, only a stage name or a file location2.
* Option C: It uses an invalid syntax for the COPY INTO command. The FILE FORMAT parameter cannot specify a stage name, only a file format name or options2.
* Option D: It uses an invalid syntax for the CREATE STAGE command. The URL parameter cannot specify a table name, only a file location1.
References:
* 1: CREATE STAGE | Snowflake Documentation
* 2: COPY INTO table | Snowflake Documentation
* 3: Cross-database Queries | Snowflake Documentation

NEW QUESTION # 46
What is a key consideration when setting up search optimization service for a table?

- A. Search optimization service works best with a column that has a minimum of 100 K distinct values.
- B. The table must be clustered with a key having multiple columns for effective search optimization.
- C. Search optimization service can help to optimize storage usage by compressing the data into a GZIP format.
- D. Search optimization service can significantly improve query performance on partitioned external tables.

Answer: A

Explanation:
Search optimization service is a feature of Snowflake that can significantly improve the performance of certain types of lookup and analytical queries on tables. Search optimization service creates and maintains a persistent data structure called a search access path, which keeps track of which values of the table's columns might be found in each of its micro-partitions, allowing some micro-partitions to be skipped when scanning the table1.
Search optimization service can significantly improve query performance on partitioned external tables, which are tables that store data in external locations such as Amazon S3 or Google Cloud Storage. Partitioned external tables can leverage the search access path to prune the partitions that do not contain the relevant data, reducing the amount of data that needs to be scanned and transferred from the external location2.
The other options are not correct because:
* A. Search optimization service works best with a column that has a high cardinality, which means that the column has many distinct values. However, there is no specific minimum number of distinct values required for search optimization service to work effectively. The actual performance improvement depends on the selectivity of the queries and the distribution of the data1.
* C. Search optimization service does not help to optimize storage usage by compressing the data into a GZIP format. Search optimization service does not affect the storage format or compression of the data, which is determined by the file format options of the table. Search optimization service only creates an additional data structure that is stored separately from the table data1.
* D. The table does not need to be clustered with a key having multiple columns for effective search optimization. Clustering is a feature of Snowflake that allows ordering the data in a table or a partitioned external table based on one or more clustering keys. Clustering can improve the performance of queries that filter on the clustering keys, as it reduces the number of micro-partitions that

need to be scanned. However, clustering is not required for search optimization service to work, as search optimization service can skip micro-partitions based on any column that has a search access path, regardless of the clustering key3.
References:
* 1: Search Optimization Service | Snowflake Documentation
* 2: Partitioned External Tables | Snowflake Documentation
* 3: Clustering Keys | Snowflake Documentation

NEW QUESTION # 47
A user named USER_01 needs access to create a materialized view on a schema EDW. STG_SCHEMA. How can this access be provided?

- A. GRANT CREATE MATERIALIZED VIEW ON DATABASE EDW TO USER USERJD1;
- B. GRANT CREATE MATERIALIZED VIEW ON SCHEMA EDW.STG_SCHEMA TO USER USER_01;
- C. GRANT ROLE NEW_ROLE TO USER USER_01;
  GRANT CREATE MATERIALIZED VIEW ON SCHEMA ECW.STG_SCHEKA TO NEW_ROLE;
- D. GRANT ROLE NEW_ROLE TO USER_01;
  GRANT CREATE MATERIALIZED VIEW ON EDW.STG_SCHEMA TO NEW_ROLE;

Answer: B

Explanation:
* The correct answer is A because it grants the specific privilege to create a materialized view on the schema EDW.STG_SCHEMA to the user USER_01 directly.
* Option B is incorrect because it grants the privilege to create a materialized view on the entire database EDW, which is too broad and unnecessary. Also, there is a typo in the user name (USERJD1 instead of USER_01).
* Option C is incorrect because it grants the privilege to create a materialized view on a different schema (ECW.STG_SCHEKA instead of EDW.STG_SCHEMA). Also, there is no need to create a new role for this purpose.
* Option D is incorrect because it grants the privilege to create a materialized view on an invalid object (EDW.STG_SCHEMA is not a valid schema name, it should be EDW.STG_SCHEMA). Also, there is no need to create a new role for this purpose.
References:
* Snowflake Documentation: CREATE MATERIALIZED VIEW
* Snowflake Documentation: Working with Materialized Views
* [Snowflake Documentation: GRANT Privileges on a Schema]

NEW QUESTION # 48
A healthcare company is deploying a Snowflake account that may include Personal Health Information (PHI).
The company must ensure compliance with all relevant privacy standards.
Which best practice recommendations will meet data protection and compliance requirements? (Choose three.)

- A. Create Dynamic Data Masking policies and apply them to columns that contain PHI.
- B. Use the External Tokenization feature to obfuscate sensitive data.
- C. Rewrite SQL queries to eliminate projections of PHI data based on current_role().
- D. Avoid sharing data with partner organizations.
- E. Use the Internal Tokenization feature to obfuscate sensitive data.
- F. Use, at minimum, the Business Critical edition of Snowflake.

Answer: A,B,F

Explanation:
* A healthcare company that handles PHI data must ensure compliance with relevant privacy standards, such as HIPAA, HITRUST, and GDPR. Snowflake provides several features and best practices to help customers meet their data protection and compliance requirements1.
* One best practice recommendation is to use, at minimum, the Business Critical edition of Snowflake. This edition provides the highest level of data protection and security, including end-to-end encryption with customer-managed keys, enhanced object-level security, and HIPAA and HITRUST compliance2. Therefore, option A is correct.
* Another best practice recommendation is to create Dynamic Data Masking policies and apply them to columns that contain PHI. Dynamic Data Masking is a feature that allows masking or redacting sensitive data based on the current user's role. This way, only authorized users can view the unmasked data, while others will see masked values, such as NULL, asterisks, or random

characters3. Therefore, option B is correct.
* A third best practice recommendation is to use the External Tokenization feature to obfuscate sensitive data. External Tokenization is a feature that allows replacing sensitive data with tokens that are
* generated and stored by an external service, such as Protegrity. This way, the original data is never stored or processed by Snowflake, and only authorized users can access the tokenized data through the external service4. Therefore, option D is correct.
* Option C is incorrect, because the Internal Tokenization feature is not available in Snowflake. Snowflake does not provide any native tokenization functionality, but only supports integration with external tokenization services4.
* Option E is incorrect, because rewriting SQL queries to eliminate projections of PHI data based on current_role() is not a best practice. This approach is error-prone, inefficient, and hard to maintain. A better alternative is to use Dynamic Data Masking policies, which can automatically mask data based on the user's role without modifying the queries3.
* Option F is incorrect, because avoiding sharing data with partner organizations is not a best practice.
Snowflake enables secure and governed data sharing with internal and external consumers, such as business units, customers, or partners. Data sharing does not involve copying or moving data, but only granting access privileges to the shared objects. Data sharing can also leverage Dynamic Data Masking and External Tokenization features to protect sensitive data5.
References: : Snowflake's Security & Compliance Reports : Snowflake Editions : Dynamic Data Masking : External Tokenization : Secure Data Sharing

NEW QUESTION # 49
An Architect has been asked to clone schema STAGING as it looked one week ago, Tuesday June 1st at 8:00 AM, to recover some objects.
The STAGING schema has 50 days of retention.
The Architect runs the following statement:
CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-06-01 08:00:00'); The Architect receives the following error: Time travel data is not available for schema STAGING. The requested time is either beyond the allowed time travel period or before the object creation time.
The Architect then checks the schema history and sees the following:
CREATED_ON|NAME|DROPPED_ON
2021-06-02 23:00:00 | STAGING | NULL
2021-05-01 10:00:00 | STAGING | 2021-06-02 23:00:00
How can cloning the STAGING schema be achieved?

- A. Rename the STAGING schema and perform an UNDROP to retrieve the previous STAGING schema version, then run the CLONE statement.
- B. Modify the statement: CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-05-01 10:00:00');
- C. Cloning cannot be accomplished because the STAGING schema version was not active during the proposed Time Travel time period.
- D. Undrop the STAGING schema and then rerun the CLONE statement.

Answer: A

NEW QUESTION # 50
.....

Syllabus

- ARA-C01 Valid Test Papers □ ARA-C01 New Exam Materials □ Reliable ARA-C01 Test Cram □ Search for { ARA-C01 } and download it for free on ➥ www.pdfvce.com □ website □ARA-C01 New Soft Simulations
- ARA-C01 Exam Torrent - ARA-C01 Exam Bootcamp - ARA-C01 Exam Cram □ Download ▸ ARA-C01 ◂ for free by simply searching on □ www.testkingpass.com □ □Guaranteed ARA-C01 Success
- Trustworthy ARA-C01 Trustworthy Exam Content - Latest Updated Exam ARA-C01 Bootcamp - High Pass-Rate Snowflake SnowPro Advanced Architect Certification □ Search for 「 ARA-C01 」 on □ www.pdfvce.com □ immediately to obtain a free download □ARA-C01 Exam Collection
- Exam ARA-C01 Quiz □ Reliable ARA-C01 Test Cram □ ARA-C01 Official Study Guide □ Immediately open 《 www.vce4dumps.com 》 and search for ▸ ARA-C01 ◂ to obtain a free download □ARA-C01 Exam Brain Dumps
- ARA-C01 Exam Torrent - ARA-C01 Exam Bootcamp - ARA-C01 Exam Cram □ Search for 《 ARA-C01 》 and obtain a free download on □ www.pdfvce.com □ □ARA-C01 Valid Test Papers
- ARA-C01 New Exam Materials □ Guaranteed ARA-C01 Success □ ARA-C01 Prepaway Dumps □ Easily obtain （ ARA-C01 ） for free download through 【 www.pdfdumps.com 】 □Reliable ARA-C01 Exam Syllabus
- ycs.instructure.com, www.notebook.ai, wibki.com, www.stes.tyc.edu.tw, cssoxfordgrammar.site, guangai.nx567.cn, backloggd.com, ycs.instructure.com, justpaste.me, disqus.com, Disposable vapes

P.S. Free & New ARA-C01 dumps are available on Google Drive shared by CramPDF: https://drive.google.com/open?id=1uioX0Bg2G7g7wKt_YUR0lARRqf2hZkNS