

# Newest DEA-C02 Exam Questions: SnowPro Advanced: Data Engineer (DEA-C02) supply you high-quality Preparation Dump - ValidDumps

## Snowflake SnowPro Advanced Data Engineer 42 Real Exam Question and Answers



Snowflake  
DEA-C01



ShapingPixel

P.S. Free & New DEA-C02 dumps are available on Google Drive shared by ValidDumps: <https://drive.google.com/open?id=1OrKhbWJB-043UTEyM3xfYVNHJpxtCjAq>

You must hold an optimistic belief for your life. There always have solutions to the problems. We really hope that our DEA-C02 study materials will greatly boost your confidence. In fact, many people are confused about their future and have no specific aims. Then our DEA-C02 practice quiz can help you find your real interests. Just think about that you will get more opportunities to bigger enterprise and better position in your career with the DEA-C02 certification. It is quite encouraging!

While making revisions and modifications to the Snowflake DEA-C02 practice exam, our team takes reports from over 90,000 professionals worldwide to make the Snowflake DEA-C02 Exam Questions foolproof. To make you capable of preparing for the DEA-C02 exam smoothly, we provide actual Snowflake DEA-C02 exam dumps.

>> DEA-C02 Latest Version <<

## Certification DEA-C02 Test Answers - Reliable DEA-C02 Exam Cram

Our DEA-C02 training materials are regarded as the most excellent practice materials by authority. Our company is dedicated to researching, manufacturing, selling and service of the DEA-C02 study guide. Also, we have our own research center and experts team. So our products can quickly meet the new demands of customers. That is why our DEA-C02 Exam Questions are popular among candidates. we have strong strenght to support our DEA-C02 practice engine.

## Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q287-Q292):

### NEW QUESTION # 287

You have data residing in AWS S3 in Parquet format, which is updated daily with new columns being added occasionally. The data is rarely accessed, but when it is, it needs to be queried using SQL within Snowflake. You want to minimize storage costs within Snowflake while ensuring the data can be queried without requiring manual table schema updates every time a new column is added to the S3 data'. Which approach is MOST suitable?

- ☐ Load the data into a Snowflake internal table using Snowpipe and regularly update the table schema using ALTER TABLE ADD COLUMN.
- ☐ Create an external table in Snowflake pointing to the S3 bucket, using the AUTO\_REFRESH=TRUE and FILE\_FORMAT = (TYPE = PARQUET) settings and rely on schema evolution support.
- ☐ Create a view on top of the S3 data using a Spark cluster and query the view from Snowflake via JDBC.
- ☐ Use Snowflake's COPY INTO command to periodically load the data into a Snowflake table and rely on schema evolution by specifying 'ON\_ERROR = CONTINUE'.
- ☐ Create an external table in Snowflake pointing to the S3 bucket, using the 'AUTO\_REFRESH=FALSE' and 'FILE\_FORMAT = (TYPE = PARQUET)' settings. Refresh the metadata manually using 'ALTER EXTERNAL TABLE ... REFRESH' after each file update.

- A. Option C
- B. Option D
- C. Option A
- D. Option E
- **E. Option B**

**Answer: E**

Explanation:

Option B is the most suitable because external tables with AUTO REFRESH enabled and properly configured file format (Parquet in this case) support schema evolution automatically. When new columns are added to the S3 data, Snowflake will detect these changes and update the external table's metadata accordingly. This eliminates the need for manual schema updates and minimizes storage costs because the data remains in S3. AUTO\_REFRESH needs a configured notification integration. Option A will require manual ALTER TABLE ADD COLUMN commands which are not scalable, option C introduces unnecessary complexity with Spark. Option D using COPY INTO does not support external storage locations. Option E requires manual refresh.

### NEW QUESTION # 288

You are developing a data pipeline that extracts data from an on-premise PostgreSQL database, transforms it, and loads it into Snowflake. You want to use the Snowflake Python connector in conjunction with a secure method for accessing the PostgreSQL database. Which of the following approaches provides the MOST secure and manageable way to handle the PostgreSQL connection credentials in your Python script when deploying to a production environment?

- A. Store the PostgreSQL username and password in a configuration file (e.g., JSON or YAML) and load the file in the Python script.
- B. Prompt the user for the PostgreSQL username and password each time the script is executed.
- C. Hardcode the PostgreSQL username and password directly into the Python script.
- D. Store the PostgreSQL username and password in environment variables and retrieve them in the Python script using 'os.environ'
- **E. Store the PostgreSQL username and password in a dedicated secrets management service (e.g., AWS Secrets Manager, HashiCorp Vault, Azure Key Vault) and retrieve them in the Python script using the appropriate API.**

**Answer: E**

Explanation:

Option D, using a dedicated secrets management service, provides the most secure and manageable approach. Secrets management services are designed to securely store and manage sensitive information like database credentials. They offer features like encryption, access control, auditing, and versioning, making them the best choice for production environments. Option A is highly insecure. Options B and C are better than A but still less secure than using a secrets management service, as environment variables and configuration files can be accidentally exposed or committed to version control. Option E is impractical and insecure for automated pipelines.

### NEW QUESTION # 289

You're tasked with building a data pipeline using Snowpark Python to incrementally load data into a target table 'SALES SUMMARY' from a source table 'RAW SALES'. The pipeline needs to ensure that only new or updated records from 'RAW SALES' are merged into 'SALES SUMMARY' based on a 'TRANSACTION ID'. You want to use Snowpark's 'MERGE' operation for this, but you also need to handle potential conflicts and log any rejected records to an error table 'SALES SUMMARY ERRORS'. Which of the following approaches offers the MOST robust and efficient solution for handling errors and ensuring data integrity within the MERGE statement?

- A. Use the 'WHEN MATCHED THEN UPDATE' clause to update existing records and the 'WHEN NOT MATCHED THEN INSERT' clause to insert new records. Implement a separate process to periodically compare 'SALES\_SUMMARY' with 'RAW\_SALES' to identify and log any inconsistencies.
- B. Use a single 'MERGE' statement with 'WHEN MATCHED THEN UPDATE' and 'WHEN NOT MATCHED THEN INSERT' clauses. Capture rejected records by leveraging the 'SYSTEM\$PIPE STATUS' function after the 'MERGE' operation to identify rows that failed during the merge.
- **C. Employ the 'MERGE' statement with 'WHEN MATCHED THEN UPDATE' and 'WHEN NOT MATCHED THEN INSERT' clauses, and use a stored procedure that executes the 'MERGE' statement and then conditionally inserts rejected records into the 'SALES SUMMARY ERRORS' table based on criteria defined within the stored procedure. This will use the table function on the output.**

- D. Incorporate an 'ELSE' clause in the 'MERGE' statement to capture records that do not satisfy the update or insert conditions due to data quality issues. Use this 'ELSE' clause to insert rejected records into 'SALES SUMMARY ERRORS'
- E. Utilize the 'WHEN MATCHED THEN UPDATE' and 'WHEN NOT MATCHED THEN INSERT' clauses with a 'WHERE' condition in each clause to filter out potentially problematic records. Log these filtered records to using a separate 'INSERT' statement after the 'MERGE' operation.

**Answer: C**

Explanation:

Option E provides the most robust solution. Using a stored procedure to execute the MERGE allows for more complex error handling logic. Critically, the result\_scan function of the MERGE query can then be used to identify and analyze the success or failure of each individual record processed within the MERGE. This avoids separate processes or post-merge comparisons and is therefore more robust. Option A requires a separate process for inconsistency checking, which is less efficient and may miss real-time errors. Options B, C, and D do not offer a reliable and atomic way to capture and log all rejected records. The SYSTEM\$PIPE\_STATUS function is relevant for Snowpipe, not direct MERGE operations.

### NEW QUESTION # 290

You have a Snowflake table 'CUSTOMER DATA' with a column 'EMAIL' containing customer email addresses. You need to classify this column as 'PII' using a tag named 'SENSITIVITY' and value 'CONFIDENTIAL'. Also, you want all queries accessing this 'EMAIL' column to be logged, with specific details about who accessed it and when. You already have appropriate roles and privileges to perform the required operations. Which of the following SQL statements, when executed in sequence, will achieve this goal, assuming appropriate logging mechanisms are already configured to read from the Snowflake ACCESS HISTORY view?

- A.

```
CREATE OR REPLACE TAG SENSITIVITY COMMENT = 'Data Sensitivity';
CREATE OR REPLACE TAG VALUE CONFIDENTIAL COMMENT = 'Confidential Data';
ALTER TABLE CUSTOMER_DATA MODIFY COLUMN EMAIL SET TAG SENSITIVITY = 'CONFIDENTIAL';
GRANT USAGE ON TAG SENSITIVITY TO ROLE data_steward;
```

- B.

```
CREATE OR REPLACE TAG SENSITIVITY COMMENT = 'Data Sensitivity';
CREATE OR REPLACE TAG SENSITIVITY.CONFIDENTIAL COMMENT = 'Confidential Data';
ALTER TABLE CUSTOMER_DATA MODIFY COLUMN EMAIL SET TAG SENSITIVITY = 'CONFIDENTIAL';
```

- C.

```
CREATE OR REPLACE TAG SENSITIVITY COMMENT = 'Data Sensitivity';
CREATE OR REPLACE TAG SENSITIVITY.CONFIDENTIAL COMMENT = 'Confidential Data';
ALTER TABLE CUSTOMER_DATA MODIFY COLUMN EMAIL SET TAG SENSITIVITY = 'CONFIDENTIAL';
GRANT USAGE ON TAG SENSITIVITY TO ROLE data_steward;
```

- D.

```
CREATE OR REPLACE TAG SENSITIVITY COMMENT = 'Data Sensitivity';
ALTER TABLE CUSTOMER_DATA MODIFY COLUMN EMAIL SET TAG SENSITIVITY = 'CONFIDENTIAL';
GRANT USAGE ON TAG SENSITIVITY TO ROLE data_steward;
```

- E.

```
ALTER TABLE CUSTOMER_DATA MODIFY COLUMN EMAIL SET TAG SENSITIVITY = 'CONFIDENTIAL';
```

**Answer: A**

Explanation:

The correct sequence of SQL statements is in Option D. First, the tag 'SENSITIVITY' needs to be created. Second, although not explicitly a 'tag value' object, you still need to create a SENSITIVITY tag object. This makes the tag value 'CONFIDENTIAL' available. Then, the tag is applied to the column using the 'ALTER TABLE MODIFY COLUMN ... SET TAG' command with the desired value ('CONFIDENTIAL'). Granting USAGE permissions to the 'data\_steward' role (though not directly related to the tagging itself) allows that role to manage and view tag assignments, which is important for governance. Options A, B, C, and E are incorrect because they do not correctly create and assign the tag and its value or have some of the GRANT statement.

### NEW QUESTION # 291

You are tasked with loading a large dataset (50TB) of JSON files into Snowflake. The JSON files are complex, deeply nested, and irregularly structured. You want to maximize loading performance while minimizing storage costs and ensuring data integrity. You have a dedicated Snowflake virtual warehouse (X-Large).

Which combination of approaches would be MOST effective?

- A. Load the JSON data using the COPY INTO command with no pre-processing. Create a VIEW on top of the raw VARIANT column to flatten the data for querying.
- **B. Use Snowpipe with auto-ingest, create a raw VARIANT column alongside projected relational columns for frequently accessed fields, and use search optimization on those projected columns.**
- C. Use Snowpipe with auto-ingest, create a single VARIANT column in your target table, and rely solely on Snowflake's automatic schema detection.
- D. Pre-process the JSON data using a Python script with Pandas to flatten the structure and convert it into a relational format like CSV. Then, load the CSV files using the COPY INTO command with gzip compression.
- E. Load the JSON data using the COPY INTO command with gzip compression. Create a raw VARIANT column alongside projected relational columns for frequently accessed fields, and use materialized views to improve query performance.

**Answer: B**

Explanation:

Option C is the most effective. Snowpipe provides continuous loading. A raw VARIANT column captures all data, and projecting commonly accessed fields into relational columns optimizes query performance. Search optimization on the projected columns allows for faster filtering and lookups. Options A, B, D, and E have trade-offs. A lacks optimized querying and can lead to expensive computations on the variant column. B requires pre-processing and may lose data fidelity. D impacts query performance due to runtime flattening. E introduces complexities with materialized view maintenance.

## NEW QUESTION # 292

.....

If you're still studying hard to pass the Snowflake DEA-C02 exam, ValidDumps help you to achieve your dream. We provide you with the best Snowflake DEA-C02 exam materials. It passed the test of practice, and with the best quality. It is better than Snowflake DEA-C02 tutorials and any other related materials. It can help you to pass the Snowflake DEA-C02 exam, and help you to become a strong IT expert.

**Certification DEA-C02 Test Answers:** <https://www.validdumps.top/DEA-C02-exam-torrent.html>

Snowflake DEA-C02 Latest Version Passing industry-relevant certification exams is always worthwhile to make your place or consolidate the existing one in the IT industry. With the help of our Certification DEA-C02 Test Answers - SnowPro Advanced: Data Engineer (DEA-C02) renewal version during the year, I assure that you will stand out in the crowd, Snowflake DEA-C02 Latest Version But you know good thing always need time and energy.

Steven Leon covers a wide spectrum of social, Valid DEA-C02 Test Duration economic and environmental issues, as well as new areas such as closed-loop supply chains, The SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam questions are checked and verified by experienced and qualified SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam trainers.

## Realistic DEA-C02 Latest Version - Certification SnowPro Advanced: Data Engineer (DEA-C02) Test Answers Free PDF Quiz

Passing industry-relevant certification exams Reliable DEA-C02 Exam Cram is always worthwhile to make your place or consolidate the existing one in the IT industry. With the help of our SnowPro Advanced: Data Engineer (DEA-C02) Valid DEA-C02 Test Duration renewal version during the year, I assure that you will stand out in the crowd.

But you know good thing always need time and energy, The first step DEA-C02 is choosing right SnowPro Advanced: Data Engineer (DEA-C02) free dumps, which will save your time and money in the preparation of SnowPro Advanced: Data Engineer (DEA-C02) passleader review.

Our passing rate for DEA-C02 certification is high up to 99.26%.

- Precious SnowPro Advanced: Data Engineer (DEA-C02) Guide Dumps Will be Your Best Choice - [www.testsdumps.com](http://www.testsdumps.com)  
☐ Search on ☐ [www.testsdumps.com](http://www.testsdumps.com) ☐ for ➤ DEA-C02 ☐ to obtain exam materials for free download ☐ DEA-C02 Related Content
- Fast Download DEA-C02 Latest Version - Pass DEA-C02 in One Time - Useful Certification DEA-C02 Test Answers ☐ Easily obtain ➡ DEA-C02 ☐ for free download through **【 [www.pdfvce.com](http://www.pdfvce.com) 】** ☐ Valid DEA-C02 Test Simulator
- DEA-C02 Valid Exam Topics ☐ New DEA-C02 Braindumps ☐ New DEA-C02 Braindumps ☐ Open website “

www.passcollection.com” and search for ➤ DEA-C02 ☐ for free download ☐DEA-C02 Valid Exam Experience

- Free PDF Valid Snowflake - DEA-C02 - SnowPro Advanced: Data Engineer (DEA-C02) Latest Version ☐ Search for ☐ DEA-C02 ☐ on ☐ www.pdfvce.com ☐ immediately to obtain a free download ☐New DEA-C02 Braindumps
- SnowPro Advanced: Data Engineer (DEA-C02) exam certification - DEA-C02 exam reviews ☐ Search for ⇒ DEA-C02 ⇐ and easily obtain a free download on [ www.passcollection.com ] ☐Latest DEA-C02 Exam Duration
- Precious SnowPro Advanced: Data Engineer (DEA-C02) Guide Dumps Will be Your Best Choice - Pdfvce ⇨ Go to website 《 www.pdfvce.com 》 open and search for ☐ DEA-C02 ☐ to download for free ☐DEA-C02 Latest Exam Papers
- DEA-C02 Related Content ☐ DEA-C02 Related Content ☐ Reliable DEA-C02 Exam Pattern ☐ Download { DEA-C02 } for free by simply searching on ✓ www.exams4collection.com ☐✓☐ ☐DEA-C02 Valid Exam Experience
- Latest DEA-C02 Exam Duration ☐ DEA-C02 Exam Question ☐ Well DEA-C02 Prep ☐ ☐ www.pdfvce.com ☐ is best website to obtain ☐ DEA-C02 ☐ for free download ☐Exam DEA-C02 Tutorials
- Pass Guaranteed Quiz Reliable Snowflake - DEA-C02 Latest Version ☐ Open website “www.testsdumps.com” and search for ☼ DEA-C02 ☐☼☐ for free download ☐New DEA-C02 Test Testking
- Pdfvce Snowflake DEA-C02 Dumps - Improve Your Exam Preparation Quickly ☐ Open ✓ www.pdfvce.com ☐✓☐ enter ☐ DEA-C02 ☐ and obtain a free download ☐DEA-C02 Exam Reference
- New DEA-C02 Braindumps ⇐ DEA-C02 Valid Exam Topics ☐ DEA-C02 Related Content ☐ The page for free download of 【 DEA-C02 】 on 「 www.exams4collection.com 」 will open immediately ☐Test DEA-C02 Questions Pdf
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, paulhun512.free-blogz.com, sahuwebgrowmaster.online, massageben.com, www.stes.tyc.edu.tw, worksmarterpinoy.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, school.technovators.co.za, Disposable vapes

BTW, DOWNLOAD part of ValidDumps DEA-C02 dumps from Cloud Storage: <https://drive.google.com/open?id=10rKhbWJB-043UTEyM3xfYVNHJpxtCjAq>