

DEA-C01熱門考古題 - DEA-C01證照考試



P.S. PDFExamDumps在Google Drive上分享了免費的、最新的DEA-C01考試題庫：<https://drive.google.com/open?id=1r5QF9gRtUem8-QYI3m2PtEQ6GvhVeimv>

大家都知道，PDFExamDumps Snowflake的DEA-C01考試培訓資料的知名度非常高，在全球範圍類也是赫赫有名的，為什麼會產生這麼大的連鎖反映呢，因為PDFExamDumps Snowflake的DEA-C01考試培訓資料確實很適用，而且真的可以幫助我們取得優異的成績。

Snowflake DEA-C01 考試大綱：

主題	簡介
主題 1	<ul style="list-style-type: none">• Data Transformation: The SnowPro Advanced: Data Engineer exam evaluates skills in using User-Defined Functions (UDFs), external functions, and stored procedures. It assesses the ability to handle semi-structured data and utilize Snowpark for transformations. This section ensures Snowflake engineers can effectively transform data within Snowflake environments, critical for data manipulation tasks.
主題 2	<ul style="list-style-type: none">• Performance Optimization: This topic assesses the ability to optimize and troubleshoot underperforming queries in Snowflake. Candidates must demonstrate knowledge in configuring optimal solutions, utilizing caching, and monitoring data pipelines. It focuses on ensuring engineers can enhance performance based on specific scenarios, crucial for Snowflake Data Engineers and Software Engineers.
主題 3	<ul style="list-style-type: none">• Data Movement: Snowflake Data Engineers and Software Engineers are assessed on their proficiency to load, ingest, and troubleshoot data in Snowflake. It evaluates skills in building continuous data pipelines, configuring connectors, and designing data sharing solutions.
主題 4	<ul style="list-style-type: none">• Security: The Security topic of the DEA-C01 test covers the principles of Snowflake security, including the management of system roles and data governance. It measures the ability to secure data and ensure compliance with policies, crucial for maintaining secure data environments for Snowflake Data Engineers and Software Engineers.
主題 5	<ul style="list-style-type: none">• Storage and Data Protection: The topic tests the implementation of data recovery features and the understanding of Snowflake's Time Travel and micro-partitions. Engineers are evaluated on their ability to create new environments through cloning and ensure data protection, highlighting essential skills for maintaining Snowflake data integrity and accessibility.

>> DEA-C01熱門考古題 <<

DEA-C01證照考試，DEA-C01下載

作為一位 Snowflake DEA-C01 考生而言，作好充分的準備可以幫助您通過考試。首先您必須去當地考試中心諮詢

相關考試信息，然後挑選最新的 DEA-C01 考試題庫，因為擁有了最新的 DEA-C01 考試題庫可以有利的提高通過考試的機率。使用 PDFExamDumps 的題庫可以節省您寶貴的時間，保證你順利通過 DEA-C01 考試。既能幫您節省時間，又可以順利幫助您通過考試，這將是您的最佳選擇。

最新的 SnowPro Advanced DEA-C01 免費考試真題 (Q137-Q142):

問題 #137

Data Engineer is looking out to delete staged files automatically/periodically when the data is successfully loaded into tables by the Snowpipe. For achieving the same, which options/command is best suited: [Select 2]

- A. To remove staged files that no longer needed, periodically REMOVE command can be executed to delete the files.
- B. To remove staged files that no longer needed, periodically DELETE command can be executed to delete the files.
- C. REMOVE_STAGE_FILES option can be set as True in the COPY INTO Command embedded in PIPE objects definition.
- D. PURGE option can be set as True in the COPY INTO Command embedded in PIPE objects definition.

答案: A,D

解題說明:

Explanation

Deleting Staged Files After Snowpipe Loads the Data

Pipe objects do not support the PURGE copy option. Snowpipe cannot delete staged files automatically when the data is successfully loaded into tables.

To remove staged files that you no longer need, It is recommended to periodically executing the REMOVE command to delete the files.

Alternatively, configure any lifecycle management features provided by cloud storage service provider.

問題 #138

Streams record the differences between two offsets. If a row is added and then updated in the current offset, what will be the value of METADATA\$ISUPDATE Columns in this scenario?

- A. TRUE
- B. FALSE
- C. UPDATE
- D. INSERT

答案: B

解題說明:

Explanation

Stream Columns

A stream stores an offset for the source object and not any actual table columns or data. When queried, a stream accesses and returns the historic data in the same shape as the source object (i.e. the same column names and ordering) with the following additional columns:

METADATA\$ACTION

Indicates the DML operation (INSERT, DELETE) recorded.

METADATA\$ISUPDATE

Indicates whether the operation was part of an UPDATE statement. Updates to rows in the source object are represented as a pair of DELETE and INSERT records in the stream with a metadata column METADATA\$ISUPDATE values set to TRUE.

METADATA\$ROW_ID

Specifies the unique and immutable ID for the row, which can be used to track changes to specific rows over time.

Note that streams record the differences between two offsets. If a row is added and then updated in the current offset, the delta change is a new row. The METADATA\$ISUPDATE row records a FALSE value.

問題 #139

A company uses Amazon Redshift as its data warehouse. Data encoding is applied to the existing tables of the data warehouse. A data engineer discovers that the compression encoding applied to some of the tables is not the best fit for the data.

The data engineer needs to improve the data encoding for the tables that have sub-optimal encoding.

Which solution will meet this requirement?

- A. Run the VACUUM REINDEX command against the identified tables.
- B. Run the VACUUM RECLUSTER command against the identified tables.
- **C. Run the ANALYZE COMPRESSION command against the identified tables. Manually update the compression encoding of columns based on the output of the command.**
- D. Run the ANALYZE command against the identified tables. Manually update the compression encoding of columns based on the output of the command.

答案： C

解題說明：

The ANALYZE COMPRESSION command in Amazon Redshift evaluates the existing data in a table and suggests the most optimal compression encoding for each column. After running this command, the data engineer can manually update the table to apply the recommended compression encodings. This approach ensures the best fit for data compression, improving storage efficiency and query performance.

The ANALYZE command collects statistics for query optimization, but it does not provide compression encoding recommendations. It is focused on query performance, not data compression.

The VACUUM REINDEX command does not exist in Amazon Redshift. VACUUM commands are generally used for reclaiming disk space and sorting tables, not for compression optimization.

VACUUM RECLUSTER does not exist in Amazon Redshift. VACUUM operations focus on reorganizing and sorting data but do not address compression encoding.

問題 #140

A company processes 500 GB of audience and advertising data daily, storing CSV files in Amazon S3 with schemas registered in AWS Glue Data Catalog. They need to convert these files to Apache Parquet format and store them in an S3 bucket.

The solution requires a long-running workflow with 15 GiB memory capacity to process the data concurrently, followed by a correlation process that begins only after the first two processes complete.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the workflow by using AWS Glue. Configure AWS Glue to begin the third process after the first two processes have finished.
- **B. Use AWS Glue workflows to run the first two processes in parallel. Ensure that the third process starts after the first two processes have finished.**
- C. Use AWS Step Functions to orchestrate a workflow that uses multiple AWS Lambda functions. Ensure that the third process starts after the first two processes have finished.
- D. Use Amazon EMR to run each process in the workflow. Create an Amazon Simple Queue Service (Amazon SQS) queue to handle messages that indicate the completion of the first two processes. Configure an AWS Lambda function to process the SQS queue by running the third process.

答案： B

解題說明：

AWS Glue workflows natively orchestrate multiple Glue jobs in parallel and support dependency triggers so the third job starts only after the first two finish. Glue jobs provide the required ~15 GiB memory capacity and easily convert CSV to Parquet using the Glue Data Catalog, delivering a fully managed, low-operations solution.

問題 #141

Two data engineering teams use separate AWS accounts. Both teams request access to the same datashare in an Amazon Redshift cluster that is in a third AWS account. The datashare is named salesshare.

A data engineer must use the Amazon Redshift SQL interface to grant both data engineering teams' access to the datashare.

Which command or commands will meet this requirement?

- A. GRANT USAGE ON DATASHARE salesshare TO ACCOUNT '<account ID 1>';
GRANT USAGE ON DATASHARE salesshare TO ACCOUNT '<account ID 2>';
- **B. GRANT USAGE ON DATASHARE salesshare TO NAMESPACES '<account ID 1>' AND '<account ID 2>';**
- C. GRANT USAGE ON DATASHARE salesshare TO ACCOUNTS '<account ID 1>' AND '<account ID 2>';
- D. GRANT USAGE ON DATASHARE salesshare TO NAMESPACE '<account ID 1>';

