

Snowflake DEA-C01 Exam Questions - 1 year of Free Updates

Pass Snowflake DEA-C01 Exam with Real Questions

Snowflake DEA-C01 Exam

SnowPro Advanced Data Engineer Certification

<https://www.passquestion.com/DEA-C01.html>



Pass Snowflake DEA-C01 Exam with PassQuestion DEA-C01 questions and answers in the first attempt.

<https://www.passquestion.com/>

1 / 7

P.S. Free & New DEA-C01 dumps are available on Google Drive shared by PrepAwayETE: https://drive.google.com/open?id=1AWHoYmw2xaJXqJpL1gG7hcP_0tEHIhd9

PrepAwayETE's product is prepared for people who participate in the Snowflake certification DEA-C01 exam. PrepAwayETE's training materials include not only Snowflake certification DEA-C01 exam training materials which can consolidate your expertise, but also high degree of accuracy of practice questions and answers about Snowflake Certification DEA-C01 Exam. PrepAwayETE can guarantee you pass the Snowflake certification DEA-C01 exam with high score the even if you are the first time to participate in this exam.

Snowflake PDF Questions can be used anywhere or at any time. You can download DEA-C01 dumps pdf files on your laptop, tablet, smartphone, or any other device. Practicing with Web-based and desktop DEA-C01 practice test software, you will get a strong grip on every Snowflake DEA-C01 exam topic. You can take multiple Snowflake DEA-C01 Practice Exam attempts and identify and overcome your mistakes. Furthermore, through Snowflake DEA-C01 practice test software you will improve your time-management skills. You will easily manage your time while attempting the actual DEA-C01 test.

>> **DEA-C01 Exam Torrent <<**

DEA-C01 Exam Cram Pdf | DEA-C01 Practice Exams Free

Elaborately designed and developed DEA-C01 test guide as well as good learning support services are the key to assisting our customers to realize their dreams. Our DEA-C01 study braindumps have a variety of self-learning and self-assessment functions to detect learners' study outcomes, and the statistical reporting function of our DEA-C01 Test Guide is designed for students to figure out their weaknesses and tackle the causes, thus seeking out specific methods dealing with them. Our DEA-C01 exam guide have also set a series of explanation about the complicated parts certificated.

Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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.
Topic 2	<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.
Topic 3	<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.
Topic 4	<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.
Topic 5	<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.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q83-Q88):

NEW QUESTION # 83

Mark the correct statements about Cache?

- A. The size of the warehouse cache is determined by the compute resources in the ware-house.
- B. Materialized views are faster than tables because of their "cache" (i.e. the query results for the view); in addition, if data has changed, they can use their "cache" for data that hasn't changed and use the base table for any data that has changed.
- C. Warehouse cache is dropped when the warehouse is suspended, which may result in slower initial performance for some queries after the warehouse is resumed.
- D. Materialized views are more flexible than, but typically slower than, cached results.
- E. For persisted query results of all sizes, the cache expires after 24 hours.

Answer: A,B,C,D,E

Explanation:

Explanation

How Does Warehouse Caching Impact Queries?

Each warehouse, when running, maintains a cache of table data accessed as queries are processed by the warehouse. This enables improved performance for subsequent queries if they are able to read from the cache instead of from the table(s) in the query. The size of the cache is determined by the compute resources in the warehouse (i.e. the larger the warehouse and, therefore, more compute re-sources in the warehouse), the larger the cache.

This cache is dropped when the warehouse is suspended, which may result in slower initial performance for some queries after the warehouse is resumed. As the resumed warehouse runs and pro-cesses more queries, the cache is rebuilt, and queries that are able

to take advantage of the cache will experience improved performance.

Keep this in mind when deciding whether to suspend a warehouse or leave it running. In other words, consider the trade-off between saving credits by suspending a warehouse versus maintaining the cache of data from previous queries to help with performance.

Using Persisted Query Results

When a query is executed, the result is persisted (i.e. cached) for a period of time. At the end of the time period, the result is purged from the system.

Snowflake uses persisted query results to avoid re-generating results when nothing has changed (i.e. "retrieval optimization"). In addition, you can use persisted query results to post-process the results (e.g. layering a new query on top of the results already calculated).

For persisted query results of all sizes, the cache expires after 24 hours.

Both materialized views and cached query results provide query performance benefits:

Materialized views are more flexible than, but typically slower than, cached results.

Materialized views are faster than tables because of their "cache" (i.e. the query results for the view); in addition, if data has changed, they can use their "cache" for data that hasn't changed and use the base table for any data that has changed.

Regular views do not cache data, and therefore cannot improve performance by caching.

NEW QUESTION # 84

A company plans to provision a log delivery stream within a VPC. The company configured the VPC flow logs to publish to Amazon CloudWatch Logs. The company needs to send the flow logs to Splunk in near real time for further analysis.

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

- A. Create an Amazon Kinesis Data Firehose delivery stream to use Splunk as the destination. Create a CloudWatch Logs subscription filter to send log events to the delivery stream.
- B. Configure an Amazon Kinesis Data Streams data stream to use Splunk as the destination. Create an AWS Lambda function to send the flow logs from CloudWatch Logs to the data stream.
- C. Create an Amazon Kinesis Data Firehose delivery stream to use Splunk as the destination. Create an AWS Lambda function to send the flow logs from CloudWatch Logs to the delivery stream.
- D. Configure an Amazon Kinesis Data Streams data stream to use Splunk as the destination. Create a CloudWatch Logs subscription filter to send log events to the data stream.

Answer: A

Explanation:

Kinesis Data Firehose has built-in support for Splunk as a destination, making the integration straightforward. Using a CloudWatch Logs subscription filter directly to Firehose simplifies the data flow, eliminating the need for additional Lambda functions or custom integrations.

NEW QUESTION # 85

The company stores a large volume of customer records in Amazon S3. To comply with regulations, the company must be able to access new customer records immediately for the first

30 days after the records are created. The company accesses records that are older than 30 days infrequently.

The company needs to cost-optimize its Amazon S3 storage.

Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Standard-Infrequent Access (S3 Standard-IA) storage for all customer records.
- B. Use S3 Intelligent-Tiering storage.
- C. Transition records to S3 Glacier Deep Archive storage after 30 days.
- D. Apply a lifecycle policy to transition records to S3 Standard-Infrequent-Access (S3 Standard-IA) storage after 30 days.

Answer: D

Explanation:

S3 Standard-IA (Infrequent Access) is ideal for data that is accessed less frequently but requires immediate access when needed. Since the company needs immediate access to customer records for the first 30 days and infrequent access afterward, applying an S3 lifecycle policy to automatically transition the data to S3 Standard-IA after 30 days provides a cost-effective solution. This way, records remain easily accessible during the initial 30 days and are moved to a cheaper storage tier when they are accessed less often.

S3 Intelligent-Tiering automatically moves data between frequent and infrequent access tiers based on access patterns, but it incurs

additional monitoring and automation charges. This option is typically more expensive compared to S3 Standard-IA if the access pattern is well understood (like in this case, where data is accessed frequently for the first 30 days and infrequently thereafter). S3 Glacier Deep Archive is a low-cost storage option for archival data that is rarely accessed and requires hours for retrieval. It does not meet the requirement for immediate access to records after 30 days.

Storing all data in S3 Standard-IA is more expensive during the first 30 days when the data is accessed frequently. S3 Standard is more cost-effective for frequently accessed data during that period.

NEW QUESTION # 86

To troubleshoot data load failure in one of your Copy Statement, Data Engineer have Executed a COPY statement with the VALIDATION_MODE copy option set to RETURN_ALL_ERRORS with reference to the set of files he had attempted to load. Which below function can facilitate analysis of the problematic records on top of the Results produced? [Select 2]

- A. LOAD_ERROR
- B. RESULT_SCAN
- C. Rejected_record
- D. LAST_QUERY_ID

Answer: B,D

Explanation:

Explanation

LAST_QUERY_ID() Function

Returns the ID of a specified query in the current session. If no query is specified, the most recently executed query is returned.

RESULT_SCAN() Function

Returns the result set of a previous command (within 24 hours of when you executed the query) as if the result was a table.

The following example validates a set of files (SFfile.csv.gz) that contain errors. To facilitate analysis of the errors, a COPY INTO <location> statement then unloads the problematic records into a text file so they could be analyzed and fixed in the original data files. The statement queries the RESULT_SCAN table.

```
1.#copy into Snowtable
2.from @SFstage/SFfile.csv.gz
3.validation_mode=return_all_errors;
4.#set qid=last_query_id();
5.#copy into @SFstage/errors/load_errors.txt from(select rejected_record from table(result_scan($qid))); Note: Other options are not valid functions.
```

NEW QUESTION # 87

A company uses Amazon EMR as an extract, transform, and load (ETL) pipeline to transform data that comes from multiple sources. A data engineer must orchestrate the pipeline to maximize performance.

Which AWS service will meet this requirement MOST cost effectively?

- A. Amazon EventBridge
- B. AWS Glue Workflows
- C. Amazon Managed Workflows for Apache Airflow (Amazon MWAA)
- D. AWS Step Functions

Answer: D

Explanation:

Glue Workflows is for Glue job orchestration. C is for orchestration with different AWS services.

NEW QUESTION # 88

.....

Having a good command of professional knowledge in this line, they devised our high quality and high effective DEA-C01 study materials by unremitting effort and studious research. They are meritorious and unsuspecting experts with professional background. By concluding quintessential points into DEA-C01 Preparation engine, you can pass the exam with the least time while huge progress. And our pass rate of the DEA-C01 exam questions is high as 98% to 100%.

DEA-C01 Exam Cram Pdf: <https://www.prepawayete.com/Snowflake/DEA-C01-practice-exam-dumps.html>

P.S. Free & New DEA-C01 dumps are available on Google Drive shared by PrepAwayETE: https://drive.google.com/open?id=1AWHoYmw2xaJXqJpL1gG7hcP_0tEHIhd9