

Braindumps Snowflake DEA-C01 Downloads - DEA-C01 Test Questions Fee

Pass Snowflake DEA-C01 Exam with Real Questions

Snowflake DEA-C01 Exam

SnowPro Advanced Data Engineer Certification

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



35% OFF on All, Including DEA-C01 Questions and Answers

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

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

P.S. Free & New DEA-C01 dumps are available on Google Drive shared by Exam4Free: <https://drive.google.com/open?id=1utz0iykj8KqpPvb1NURnQNul5vi9EegX>

Get the test DEA-C01 certification requires the user to have extremely high concentration will all test sites in mind, and this is definitely a very difficult. Our DEA-C01 learning questions can successfully solve this question for you for the content are exactly close to the changes of the DEA-C01 Real Exam. When you grasp the key points, nothing will be difficult for you anymore. Our professional experts are good at compiling the DEA-C01 training guide with the most important information. Believe in us, and your success is 100% guaranteed!

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 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.
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 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.

>> **Braindumps Snowflake DEA-C01 Downloads** <<

Braindumps DEA-C01 Downloads offer you accurate Test Questions Fee to pass Snowflake SnowPro Advanced: Data Engineer Certification Exam exam

The format name of Channel Partner Program DEA-C01 practice test questions is Snowflake PDF Questions file, desktop practice test software, and web-based practice test software. Choose the nay type of Channel Partner Program SnowPro Advanced: Data Engineer Certification Exam DEA-C01 Practice Exam Questions that fit your Snowflake DEA-C01 exam preparation requirement and budget and start preparation without wasting further time.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q295-Q300):

NEW QUESTION # 295

A company has a gaming application that stores data in Amazon DynamoDB tables. A data engineer needs to ingest the game data into an Amazon OpenSearch Service cluster. Data updates must occur in near real time.

Which solution will meet these requirements?

- A. Use AWS Step Functions to periodically export data from the Amazon DynamoDB tables to an Amazon S3 bucket. Use an AWS Lambda function to load the data into Amazon OpenSearch Service.
- **B. Use Amazon DynamoDB Streams to capture table changes. Use an AWS Lambda function to process and update the data in Amazon OpenSearch Service.**
- C. Configure an AWS Glue job to have a source of Amazon DynamoDB and a destination of Amazon OpenSearch Service to transfer data in near real time.
- D. Use a custom OpenSearch plugin to sync data from the Amazon DynamoDB tables.

Answer: B

Explanation:

DynamoDB Streams can capture changes to items in DynamoDB tables (such as inserts, updates, and deletes) in near real-time. An AWS Lambda function can be triggered by these streams, and the function can process the changes and update the data in Amazon OpenSearch Service. This solution meets the requirement for near real-time updates with minimal latency.

The "Use AWS Step Functions to periodically export data from the Amazon DynamoDB tables to an Amazon S3 bucket. Use an AWS Lambda function to load the data into Amazon OpenSearch Service." solution involves exporting data periodically, which does not meet the requirement for near real-time updates. Periodic exports could introduce delays.

AWS Glue is generally used for batch processing, and it is not designed for near real-time data ingestion. It would not be an ideal solution for real-time or low-latency requirements.

The "Use a custom OpenSearch plugin to sync data from the Amazon DynamoDB tables." option suggests building a custom

solution, which would involve significant development effort.

Moreover, OpenSearch does not natively support such plugins, making this solution less practical compared to using DynamoDB Streams and Lambda, which are designed to work together.

NEW QUESTION # 296

A Data Engineer wants to create a new development database (DEV) as a clone of the permanent production database (PROD)

There is a requirement to disable Fail-safe for all tables.

Which command will meet these requirements?

- A. CREATE DATABASE DEV
CLOSE PROD
DATA_RETENTION_TIME_IN_DAYS =0L
- B. CREATE DATABASE DEV
CLONE PROD
FAIL_SAFE=FALSE;
- C. CREATE DATABASE DEV
CLONE PROD;
- **D. CREATE TRANSIENT DATABASE DEV
CLONE RPOD**

Answer: D

Explanation:

Explanation

This option will meet the requirements of creating a new development database (DEV) as a clone of the permanent production database (PROD) and disabling Fail-safe for all tables. By using the CREATE TRANSIENT DATABASE command, the Data Engineer can create a transient database that does not have Fail-safe enabled by default. Fail-safe is a feature in Snowflake that provides additional protection against data loss by retaining historical data for seven days beyond the time travel retention period. Transient databases do not have Fail-safe enabled, which means that they do not incur additional storage costs for historical data beyond their time travel retention period. By using the CLONE option, the Data Engineer can create an exact copy of the PROD database, including its schemas, tables, views, and other objects.

NEW QUESTION # 297

A company uses Amazon DataZone as a data governance and business catalog solution. The company stores data in an Amazon S3 data lake. The company uses AWS Glue with an AWS Glue Data Catalog.

A data engineer needs to publish AWS Glue Data Quality scores to the Amazon DataZone portal.

Which solution will meet this requirement?

- A. Configure AWS Glue ETL jobs to use an Evaluate Data Quality transform. Define a data quality ruleset inside the jobs. Configure the Amazon DataZone project to have an AWS Glue data source. Enable the data quality configuration for the data source.
- B. Create a data quality ruleset with Data Quality Definition language (DQDL) rules that apply to a specific AWS Glue table. Schedule the ruleset to run daily. Configure the Amazon DataZone project to have an Amazon Redshift data source. Enable the data quality configuration for the data source.
- C. Configure AWS Glue ETL jobs to use an Evaluate Data Quality transform. Define a data quality ruleset inside the jobs. Configure the Amazon DataZone project to have an Amazon Redshift data source. Enable the data quality configuration for the data source.
- **D. Create a data quality ruleset with Data Quality Definition language (DQDL) rules that apply to a specific AWS Glue table. Schedule the ruleset to run daily. Configure the Amazon DataZone project to have an AWS Glue data source. Enable the data quality configuration for the data source.**

Answer: D

Explanation:

Data Quality Ruleset: Creating a ruleset with Data Quality Definition Language (DQDL) rules allows for defining and evaluating data quality on specific AWS Glue tables, enabling automated checks on data quality.

Scheduled Execution: Running the ruleset daily ensures that data quality scores are regularly updated.

AWS Glue Data Source in Amazon DataZone: Configuring Amazon DataZone with an AWS Glue data source enables seamless integration, allowing data quality scores from AWS Glue Data Quality to be published to the Amazon DataZone portal.

NEW QUESTION # 298

A company stores its processed data in an S3 bucket. The company has a strict data access policy. The company uses IAM roles to grant teams within the company different levels of access to the S3 bucket.

The company wants to receive notifications when a user violates the data access policy. Each notification must include the username of the user who violated the policy.

Which solution will meet these requirements?

- A. Use Amazon S3 server access logs to monitor access to the bucket. Forward the access logs to an Amazon CloudWatch log group. Use metric filters on the log group to set up CloudWatch alarms.
- B. Use Amazon CloudWatch metrics to gather object-level metrics. Set up CloudWatch alarms.
- C. Use AWS CloudTrail to track object-level events for the S3 bucket. Forward events to Amazon CloudWatch to set up CloudWatch alarms.
- D. Use AWS Config rules to detect violations of the data access policy. Set up compliance alarms.

Answer: C

Explanation:

AWS CloudTrail provides detailed logging of AWS API calls, including object-level events for Amazon S3. You can configure CloudTrail to track data events for the S3 bucket, which will log each access to the objects in the bucket, including the username of the IAM entity (user or role) that accessed the data. By forwarding these CloudTrail events to Amazon CloudWatch, you can set up alarms to trigger notifications when policy violations are detected, and include the violating user's identity.

AWS Config is primarily designed for monitoring resource configurations, not for tracking real-time access events like data access violations. It does not directly provide information about which user accessed the S3 objects.

CloudWatch metrics can monitor S3 storage or request metrics, but it does not provide detailed logging of object-level access, including the username of the violator. You need CloudTrail to get detailed event information.

While S3 server access logs can track access to objects, they lack real-time processing and do not provide the same level of detail as CloudTrail.

Additionally, processing access logs with CloudWatch metric filters adds more complexity compared to using CloudTrail, which is designed for this use case.

NEW QUESTION # 299

A company needs a solution to manage costs for an existing Amazon DynamoDB table. The company also needs to control the size of the table. The solution must not disrupt any ongoing read or write operations. The company wants to use a solution that automatically deletes data from the table after 1 month.

Which solution will meet these requirements with the LEAST ongoing maintenance?

- A. Configure a stream on the DynamoDB table to invoke an AWS Lambda function. Configure the Lambda function to delete data in the table that is older than 1 month.
- B. Configure a scheduled Amazon EventBridge rule to invoke an AWS Lambda function to check for data that is older than 1 month. Configure the Lambda function to delete old data.
- C. Use an AWS Lambda function to periodically scan the DynamoDB table for data that is older than 1 month. Configure the Lambda function to delete old data.
- D. Use the DynamoDB TTL feature to automatically expire data based on timestamps.

Answer: D

Explanation:

DynamoDB TTL (Time to Live) is a built-in feature that allows you to automatically delete expired items from your table based on a timestamp attribute.

Once TTL is enabled, DynamoDB automatically removes the items after they reach the expiration time, without disrupting any ongoing read or write operations. This solution provides the least ongoing maintenance because it's fully managed by DynamoDB and doesn't require manual interventions or additional Lambda functions.

NEW QUESTION # 300

.....

The desktop-based practice exam is customizable, tracks your progress, and creates a real SnowPro Advanced: Data Engineer

