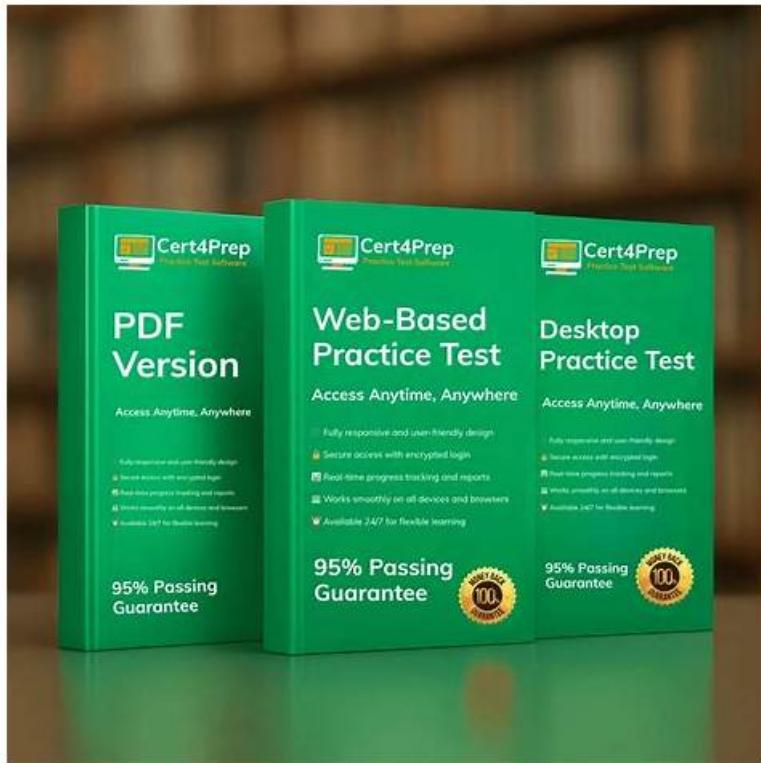


# Data-Engineer-Associate Valid Exam Braindumps | Data-Engineer-Associate Latest Braindumps Book



What's more, part of that ActualTestsQuiz Data-Engineer-Associate dumps now are free: <https://drive.google.com/open?id=1oWsOhZBwK0Ki5LwDBs8fdb49uT0FNnib>

AS the most popular Data-Engineer-Associate learning braindumps in the market, our customers are all over the world. So the content of Data-Engineer-Associate exam questions you see are very comprehensive, but it is by no means a simple display. In order to ensure your learning efficiency, we have made scientific arrangements for the content of the Data-Engineer-Associate Actual Exam. Our system is also built by professional and specialized staff and you will have a very good user experience.

The Data-Engineer-Associate exam questions by experts based on the calendar year of all kinds of exam after analysis, it is concluded that conforms to the exam thesis focus in the development trend, and summarize all kind of difficulties you will face, highlight the user review must master the knowledge content. Our AWS Certified Data Engineer - Associate (DEA-C01) study question has high quality. So there is all effective and central practice for you to prepare for your test. With our professional ability, we can accord to the necessary testing points to edit Data-Engineer-Associate Exam Questions. It points to the exam heart to solve your difficulty.

>> Data-Engineer-Associate Valid Exam Braindumps <<

## Wonderful Data-Engineer-Associate Exam Questions: AWS Certified Data Engineer - Associate (DEA-C01) Exhibit the Most Useful Training Guide- ActualTestsQuiz

Candidates who become Amazon Data-Engineer-Associate certified demonstrate their worth in the Amazon field. Data-Engineer-Associate certification is proof of their competence and skills. This is a highly sought after credential and it makes career advancement easier for the candidate. To become Amazon Data-Engineer-Associate Certified, you must pass the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) Exam. For this task, you need actual and updated Data-Engineer-Associate Questions.

## Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample

## Questions (Q109-Q114):

### NEW QUESTION # 109

A company needs to set up a data catalog and metadata management for data sources that run in the AWS Cloud. The company will use the data catalog to maintain the metadata of all the objects that are in a set of data stores. The data stores include structured sources such as Amazon RDS and Amazon Redshift. The data stores also include semistructured sources such as JSON files and .xml files that are stored in Amazon S3.

The company needs a solution that will update the data catalog on a regular basis. The solution also must detect changes to the source metadata.

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

- A. Use Amazon DynamoDB as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the DynamoDB data catalog. Schedule the Lambda functions to run periodically.
- B. Use Amazon Aurora as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the Aurora data catalog. Schedule the Lambda functions to run periodically.
- C. Use the AWS Glue Data Catalog as the central metadata repository. Extract the schema for Amazon RDS and Amazon Redshift sources, and build the Data Catalog. Use AWS Glue crawlers for data that is in Amazon S3 to infer the schema and to automatically update the Data Catalog.
- D. Use the AWS Glue Data Catalog as the central metadata repository. Use AWS Glue crawlers to connect to multiple data stores and to update the Data Catalog with metadata changes. Schedule the crawlers to run periodically to update the metadata catalog.

**Answer: D**

Explanation:

This solution will meet the requirements with the least operational overhead because it uses the AWS Glue Data Catalog as the central metadata repository for data sources that run in the AWS Cloud. The AWS Glue Data Catalog is a fully managed service that provides a unified view of your data assets across AWS and on-premises data sources. It stores the metadata of your data in tables, partitions, and columns, and enables you to access and query your data using various AWS services, such as Amazon Athena, Amazon EMR, and Amazon Redshift Spectrum. You can use AWS Glue crawlers to connect to multiple data stores, such as Amazon RDS, Amazon Redshift, and Amazon S3, and to update the Data Catalog with metadata changes.

AWS Glue crawlers can automatically discover the schema and partition structure of your data, and create or update the corresponding tables in the Data Catalog. You can schedule the crawlers to run periodically to update the metadata catalog, and configure them to detect changes to the source metadata, such as new columns, tables, or partitions<sup>12</sup>.

The other options are not optimal for the following reasons:

- \* A. Use Amazon Aurora as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the Aurora data catalog. Schedule the Lambda functions to run periodically. This option is not recommended, as it would require more operational overhead to create and manage an Amazon Aurora database as the data catalog, and to write and maintain AWS Lambda functions to gather and update the metadata information from multiple sources. Moreover, this option would not leverage the benefits of the AWS Glue Data Catalog, such as data cataloging, data transformation, and data governance.
- \* C. Use Amazon DynamoDB as the data catalog. Create AWS Lambda functions that will connect to the data catalog. Configure the Lambda functions to gather the metadata information from multiple sources and to update the DynamoDB data catalog. Schedule the Lambda functions to run periodically. This option is also not recommended, as it would require more operational overhead to create and manage an Amazon DynamoDB table as the data catalog, and to write and maintain AWS Lambda functions to gather and update the metadata information from multiple sources. Moreover, this option would not leverage the benefits of the AWS Glue Data Catalog, such as data cataloging, data transformation, and data governance.
- \* D. Use the AWS Glue Data Catalog as the central metadata repository. Extract the schema for Amazon RDS and Amazon Redshift sources, and build the Data Catalog. Use AWS Glue crawlers for data that is in Amazon S3 to infer the schema and to automatically update the Data Catalog. This option is not optimal, as it would require more manual effort to extract the schema for Amazon RDS and Amazon Redshift sources, and to build the Data Catalog. This option would not take advantage of the AWS Glue crawlers' ability to automatically discover the schema and partition structure of your data from various data sources, and to create or update the corresponding tables in the Data Catalog.

References:

\* 1: AWS Glue Data Catalog

\* 2: AWS Glue Crawlers

\* : Amazon Aurora

\* : AWS Lambda

\* : Amazon DynamoDB

## NEW QUESTION # 110

A banking company uses an application to collect large volumes of transactional data. The company uses Amazon Kinesis Data Streams for real-time analytics. The company's application uses the PutRecord action to send data to Kinesis Data Streams. A data engineer has observed network outages during certain times of day. The data engineer wants to configure exactly-once delivery for the entire processing pipeline.

Which solution will meet this requirement?

- A. Stop using Kinesis Data Streams. Use Amazon EMR instead. Use Apache Flink and Apache Spark Streaming in Amazon EMR.
- B. Design the application so it can remove duplicates during processing by embedding a unique ID in each record at the source.
- C. Design the data source so events are not ingested into Kinesis Data Streams multiple times.
- D. Update the checkpoint configuration of the Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) data collection application to avoid duplicate processing of events.

### Answer: B

Explanation:

For exactly-once delivery and processing in Amazon Kinesis Data Streams, the best approach is to design the application so that it handles idempotency. By embedding a unique ID in each record, the application can identify and remove duplicate records during processing.

Exactly-Once Processing:

Kinesis Data Streams does not natively support exactly-once processing. Therefore, idempotency should be designed into the application, ensuring that each record has a unique identifier so that the same event is processed only once, even if it is ingested multiple times.

This pattern is widely used for achieving exactly-once semantics in distributed systems.

Reference:

Alternatives Considered:

B (Checkpoint configuration): While updating the checkpoint configuration can help with some aspects of duplicate processing, it is not a full solution for exactly-once delivery.

C (Design data source): Ensuring events are not ingested multiple times is ideal, but network outages can make this difficult, and it doesn't guarantee exactly-once delivery.

D (Using EMR): While using EMR with Flink or Spark could work, it introduces unnecessary complexity compared to handling idempotency at the application level.

Amazon Kinesis Best Practices for Exactly-Once Processing

Achieving Idempotency with Amazon Kinesis

## NEW QUESTION # 111

A company has a data processing pipeline that includes several dozen steps. The data processing pipeline needs to send alerts in real time when a step fails or succeeds. The data processing pipeline uses a combination of Amazon S3 buckets, AWS Lambda functions, and AWS Step Functions state machines.

A data engineer needs to create a solution to monitor the entire pipeline.

Which solution will meet these requirements?

- A. Configure the AWS Lambda functions to store notifications in an Amazon S3 bucket when the state machines finish running. Enable S3 event notifications on the S3 bucket.
- B. Configure an Amazon EventBridge rule to react when the execution status of a state machine changes. Configure the rule to send a message to an Amazon Simple Notification Service (Amazon SNS) topic that sends notifications.
- C. Configure the Step Functions state machines to store notifications in an Amazon S3 bucket when the state machines finish running. Enable S3 event notifications on the S3 bucket.
- D. Use AWS CloudTrail to send a message to an Amazon Simple Notification Service (Amazon SNS) topic that sends notifications when a state machine fails to run or succeeds to run.

### Answer: B

Explanation:

AWS Step Functions natively emits state change events to Amazon EventBridge, which can trigger an Amazon SNS notification.

This is the most direct and real-time way to alert on success/failure without relying on custom logging or polling.

"Step Functions automatically emits status changes that EventBridge can capture to trigger alerts or workflows. Use EventBridge to

invoke an SNS topic for real-time alerts on job status."

- Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This provides real-time alerting and the least operational overhead.

## NEW QUESTION # 112

A data engineer must orchestrate a series of Amazon Athena queries that will run every day. Each query can run for more than 15 minutes.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch.
- B. Use an AWS Lambda function and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically.
- C. Use an AWS Glue Python shell job and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically.
- D. Create an AWS Step Functions workflow and add two states. Add the first state before the Lambda function. Configure the second state as a Wait state to periodically check whether the Athena query has finished using the Athena Boto3 get\_query\_execution API call. Configure the workflow to invoke the next query when the current query has finished running.
- E. Use an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully. Configure the Python shell script to invoke the next query when the current query has finished running.

**Answer: B,D**

Explanation:

Option A and B are the correct answers because they meet the requirements most cost-effectively. Using an AWS Lambda function and the Athena Boto3 client start\_query\_execution API call to invoke the Athena queries programmatically is a simple and scalable way to orchestrate the queries. Creating an AWS Step Functions workflow and adding two states to check the query status and invoke the next query is a reliable and efficient way to handle the long-running queries.

Option C is incorrect because using an AWS Glue Python shell job to invoke the Athena queries programmatically is more expensive than using a Lambda function, as it requires provisioning and running a Glue job for each query.

Option D is incorrect because using an AWS Glue Python shell script to run a sleep timer that checks every 5 minutes to determine whether the current Athena query has finished running successfully is not a cost- effective or reliable way to orchestrate the queries, as it wastes resources and time.

Option E is incorrect because using Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the Athena queries in AWS Batch is an overkill solution that introduces unnecessary complexity and cost, as it requires setting up and managing an Airflow environment and an AWS Batch compute environment.

:

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.2: AWS Lambda, Section 5.3: AWS Step Functions, Pages 125-135 Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.1: AWS Lambda, Lesson 5.2: AWS Step Functions, Pages 1-15 AWS Documentation Overview, AWS Lambda Developer Guide, Working with AWS Lambda Functions, Configuring Function Triggers, Using AWS Lambda with Amazon Athena, Pages 1-4 AWS Documentation Overview, AWS Step Functions Developer Guide, Getting Started, Tutorial: Create a Hello World Workflow, Pages 1-8

## NEW QUESTION # 113

A company stores server logs in an Amazon S3 bucket. The company needs to keep the logs for 1 year. The logs are not required after 1 year.

A data engineer needs a solution to automatically delete logs that are older than 1 year.

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

- A. Schedule a cron job on an Amazon EC2 instance to delete the logs after 1 year.
- B. Create an AWS Lambda function to delete the logs after 1 year.
- C. Configure an AWS Step Functions state machine to delete the logs after 1 year.
- D. Define an S3 Lifecycle configuration to delete the logs after 1 year.

**Answer: B**

Explanation:

#### Problem Analysis:

The company uses AWS Glue for ETL pipelines and requires automatic data quality checks during pipeline execution. The solution must integrate with existing AWS Glue pipelines and evaluate data quality rules based on predefined thresholds.

#### Key Considerations:

Ensure minimal implementation effort by leveraging built-in AWS Glue features.

Use a standardized approach for defining and evaluating data quality rules.

Avoid custom libraries or external frameworks unless absolutely necessary.

#### Solution Analysis:

##### Option A: SQL Transform

Adding SQL transforms to define and evaluate data quality rules is possible but requires writing complex queries for each rule.

Increases operational overhead and deviates from Glue's declarative approach.

##### Option B: Evaluate Data Quality Transform with DQDL

AWS Glue provides a built-in Evaluate Data Quality transform.

Allows defining rules in Data Quality Definition Language (DQDL), a concise and declarative way to define quality checks.

Fully integrated with Glue Studio, making it the least effort solution.

##### Option C: Custom Transform with PyDeequ

PyDeequ is a powerful library for data quality checks but requires custom code and integration.

Increases implementation effort compared to Glue's native capabilities.

##### Option D: Custom Transform with Great Expectations

Great Expectations is another powerful library for data quality but adds complexity and external dependencies.

#### Final Recommendation:

Use Evaluate Data Quality transform in AWS Glue.

Define rules in DQDL for checking thresholds, null values, or other quality criteria.

This approach minimizes development effort and ensures seamless integration with AWS Glue.

#### AWS Glue Data Quality Overview

#### DQDL Syntax and Examples

#### Glue Studio Transformations

## NEW QUESTION # 114

.....

We stand behind all of our customers, so we provide you with the best valid and useful Amazon Data-Engineer-Associate exam training. Regular and frequent updates for Data-Engineer-Associate dumps are necessary, so you can get hold of the Data-Engineer-Associate updated exam material every time. Besides, we offer the exact questions with correct answers, which can ensure you 100% pass in your Amazon Data-Engineer-Associate Actual Test. We have 100% money back guarantee, in case of failure, we will give you full refund.

**Data-Engineer-Associate Latest Braindumps Book:** <https://www.actualtestsquiz.com/Data-Engineer-Associate-test-torrent.html>

Then certain money will soon be deducted from your credit card to pay for the Data-Engineer-Associate study materials, Additionally, this AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) test is supported by operating systems including Android, Mac, iOS, Windows, and Linux, Amazon Data-Engineer-Associate Valid Exam Braindumps PayPal doesn't have extra costs, We can guarantee that our Data-Engineer-Associate study materials will be suitable for all people and meet the demands of all people, including students, workers and housewives and so on.

Incorporating BizTalk Server into Your Design, Background Concepts" provides background concepts for Hadoop and big data, Then certain money will soon be deducted from your credit card to pay for the Data-Engineer-Associate Study Materials.

## Perfect Data-Engineer-Associate Valid Exam Braindumps & Leading Offer in Qualification Exams & Fantastic Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01)

Additionally, this AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) test is supported by operating systems including Android, Mac, iOS, Windows, and Linux, PayPal doesn't have extra costs.

We can guarantee that our Data-Engineer-Associate study materials will be suitable for all people and meet the demands of all people, including students, workers and housewives and so on.

Amazon Data-Engineer-Associate real exam questions are collected by our professional Amazon certified experts, which can guarantee the accuracy of Data-Engineer-Associate real exam questions.

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by ActualTestsQuiz: <https://drive.google.com/open?id=1oWsOhZBwK0Ki5LwDBs8fdB49uT0FNib>