

Reliable Data-Engineer-Associate Practice Questions, Data-Engineer-Associate Authorized Exam Dumps

AWS Data Engineer Associate (DEA-C01) Exam Dumps 2024
AWS Data Engineer Associate (DEA-C01) Practice Tests 2024. Contains 390+ exam questions to pass the exam in first attempt.
[SkillCertPro](#) offers real exam questions for practice for all major IT certifications.

- For a full set of 390+ questions. Go to <https://skillcertpro.com/product/aws-data-engineer-associate-dea-c01-exam-questions/>
- SkillCertPro offers detailed explanations to each question which helps to understand the concepts better.
- It is recommended to score above 85% in SkillCertPro exams before attempting a real exam.
- SkillCertPro updates exam questions every 2 weeks.
- You will get life time access and life time free updates
- SkillCertPro assures 100% pass guarantee in first attempt.

Below are the free 10 sample questions.

Question 1:

A data engineer working for an analytics company is working on a consumer to a Kinesis Data Streams application. They have written the consumer using Kinesis Client Library (KCL), however, currently they are receiving an ExpiredIteratorException when reading records from Kinesis Data Streams. What would you recommend to the engineer to solve their issue?

- A. Change the capacity mode of the Kinesis Data Stream to on-demand.
- B. Increase WCU in DynamoDB checkpointing table.
- C. Increase the amount of shards in Kinesis Data Streams.
- D. Increase RCU in DynamoDB checkpointing table.

Answer: B

2026 Latest ExamsReviews Data-Engineer-Associate PDF Dumps and Data-Engineer-Associate Exam Engine Free Share:
https://drive.google.com/open?id=1_hyR9sTEUfwZ8vqYCCAj9u81aojDn0v

Before you buy Data-Engineer-Associate exam torrent, you can log in to our website to download a free trial question bank, and fully experience the convenience of PDF, APP, and PC three models of Data-Engineer-Associate quiz guide. During the trial period, you can fully understand Data-Engineer-Associate practice test ' learning mode, completely eliminate any questions you have about Data-Engineer-Associate exam torrent, and make your purchase without any worries. If you are a student, Data-Engineer-Associate Quiz guide will also make your study time more flexible. With Data-Engineer-Associate exam torrent, you don't need to think about studying at the time of playing. You can study at any time you want to study and get the best learning results with the best learning status.

The value of professional qualification has been shown to rise with time. For the advancement of your profession, exams like the Amazon exam given by Amazon are crucial. Candidates aim to pass the AWS Certified Data Engineer - Associate (DEA-C01) exam on their first attempt. With Amazon Data-Engineer-Associate Exam Questions, applicants may study for and pass their desired certification exam on the first attempt. You may use ExamsReviews's top Data-Engineer-Associate study resources to prepare for the AWS Certified Data Engineer - Associate (DEA-C01) exam. The Amazon Data-Engineer-Associate exam questions offered by ExamsReviews are dependable and trustworthy sources of preparation.

>> **Reliable Data-Engineer-Associate Practice Questions** <<

Data-Engineer-Associate Authorized Exam Dumps - Most Data-Engineer-Associate Reliable Questions

ExamsReviews has put emphasis on providing our Data-Engineer-Associate exam questions with high quality products with high passing rate. Many exam candidates are uninformed about the fact that our Data-Engineer-Associate preparation materials can help them with higher chance of getting success than others. It is all about efficiency and accuracy. And what is more charming than our Data-Engineer-Associate Study Guide with a passing rate as 98% to 100%? The answer is no. Our Data-Engineer-Associate practice quiz is unique in the market.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q111-Q116):

NEW QUESTION # 111

A company is planning to upgrade its Amazon Elastic Block Store (Amazon EBS) General Purpose SSD storage from gp2 to gp3. The company wants to prevent any interruptions in its Amazon EC2 instances that will cause data loss during the migration to the upgraded storage.

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

- A. Change the volume type of the existing gp2 volumes to gp3. Enter new values for volume size, IOPS, and throughput.
- B. Use AWS DataSync to create new gp3 volumes. Transfer the data from the original gp2 volumes to the new gp3 volumes.
- C. Create new gp3 volumes. Gradually transfer the data to the new gp3 volumes. When the transfer is complete, mount the new gp3 volumes to the EC2 instances to replace the gp2 volumes.
- D. Create snapshots of the gp2 volumes. Create new gp3 volumes from the snapshots. Attach the new gp3 volumes to the EC2 instances.

Answer: A

Explanation:

Changing the volume type of the existing gp2 volumes to gp3 is the easiest and fastest way to migrate to the new storage type without any downtime or data loss. You can use the AWS Management Console, the AWS CLI, or the Amazon EC2 API to modify the volume type, size, IOPS, and throughput of your gp2 volumes. The modification takes effect immediately, and you can monitor the progress of the modification using CloudWatch. The other options are either more complex or require additional steps, such as creating snapshots, transferring data, or attaching new volumes, which can increase the operational overhead and the risk of errors. Reference:

Migrating Amazon EBS volumes from gp2 to gp3 and save up to 20% on costs (Section: How to migrate from gp2 to gp3)
Switching from gp2 Volumes to gp3 Volumes to Lower AWS EBS Costs (Section: How to Switch from GP2 Volumes to GP3 Volumes) Modifying the volume type, IOPS, or size of an EBS volume - Amazon Elastic Compute Cloud (Section: Modifying the volume type)

NEW QUESTION # 112

A company stores daily records of the financial performance of investment portfolios in .csv format in an Amazon S3 bucket. A data engineer uses AWS Glue crawlers to crawl the S3 data.

The data engineer must make the S3 data accessible daily in the AWS Glue Data Catalog.

Which solution will meet these requirements?

- A. Create an IAM role that includes the AWSGlueServiceRole policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Create a daily schedule to run the crawler. Specify a database name for the output.
- B. Create an IAM role that includes the AmazonS3FullAccess policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Allocate data processing units (DPUs) to run the crawler every day. Specify a database name for the output.
- C. Create an IAM role that includes the AmazonS3FullAccess policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Create a daily schedule to run the crawler. Configure the output destination to a new path in the existing S3 bucket.
- D. Create an IAM role that includes the AWSGlueServiceRole policy. Associate the role with the crawler. Specify the S3 bucket path of the source data as the crawler's data store. Allocate data processing units (DPUs) to run the crawler every day. Configure the output destination to a new path in the existing S3 bucket.

Answer: A

Explanation:

To make the S3 data accessible daily in the AWS Glue Data Catalog, the data engineer needs to create a crawler that can crawl the S3 data and write the metadata to the Data Catalog. The crawler also needs to run on a daily schedule to keep the Data Catalog

updated with the latest data. Therefore, the solution must include the following steps:

- * Create an IAM role that has the necessary permissions to access the S3 data and the Data Catalog. The AWSServiceRoleForGlue policy is a managed policy that grants these permissions¹.
 - * Associate the role with the crawler.
 - * Specify the S3 bucket path of the source data as the crawler's data store. The crawler will scan the data and infer the schema and format².
 - * Create a daily schedule to run the crawler. The crawler will run at the specified time every day and update the Data Catalog with any changes in the data³.
 - * Specify a database name for the output. The crawler will create or update a table in the Data Catalog under the specified database. The table will contain the metadata about the data in the S3 bucket, such as the location, schema, and classification.
- Option B is the only solution that includes all these steps. Therefore, option B is the correct answer.
- Option A is incorrect because it configures the output destination to a new path in the existing S3 bucket. This is unnecessary and may cause confusion, as the crawler does not write any data to the S3 bucket, only metadata to the Data Catalog.
- Option C is incorrect because it allocates data processing units (DPUs) to run the crawler every day. This is also unnecessary, as DPUs are only used for AWS Glue ETL jobs, not crawlers.
- Option D is incorrect because it combines the errors of option A and C. It configures the output destination to a new path in the existing S3 bucket and allocates DPUs to run the crawler every day, both of which are irrelevant for the crawler.

:

- 1: AWS managed (predefined) policies for AWS Glue - AWS Glue
- 2: Data Catalog and crawlers in AWS Glue - AWS Glue
- 3: Scheduling an AWS Glue crawler - AWS Glue
- [4]: Parameters set on Data Catalog tables by crawler - AWS Glue
- [5]: AWS Glue pricing - Amazon Web Services (AWS)

NEW QUESTION # 113

A company is setting up a data pipeline in AWS. The pipeline extracts client data from Amazon S3 buckets, performs quality checks, and transforms the data. The pipeline stores the processed data in a relational database. The company will use the processed data for future queries.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS Glue Studio to extract the data from the S3 buckets. Use AWS Glue DataBrew to perform the transformations and quality checks. Load the processed data into an Amazon RDS for MySQL instance. Load the quality check results into a new S3 bucket.
- B. Use AWS Glue Studio to extract the data from the S3 buckets. Use AWS Glue DataBrew to perform the transformations and quality checks. Load the processed data and quality check results into an Amazon RDS for MySQL instance.
- **C. Use AWS Glue ETL to extract the data from the S3 buckets and perform the transformations. Use AWS Glue Data Quality to enforce suggested quality rules. Load the data and the quality check results into an Amazon RDS for MySQL instance.**
- D. Use AWS Glue ETL to extract the data from the S3 buckets and perform the transformations. Use AWS Glue DataBrew to perform quality checks. Load the processed data and the quality check results into a new S3 bucket.

Answer: C

Explanation:

AWS Glue ETL is designed for scalable and serverless data processing, and it supports integrated quality enforcement using AWS Glue Data Quality, which makes it the most cost-effective and integrated option when combined with Amazon RDS for MySQL as the relational database.

"AWS Glue can perform data validation as part of the ETL process, ensuring data quality before storing the data in the target data store."

-Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf Using AWS Glue Data Quality directly in the ETL workflow is simpler and more cost-effective than separating transformation (Glue) and validation (DataBrew) into different services.

NEW QUESTION # 114

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 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.**
- B. 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.
- C. 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.
- D. 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.

Answer: A

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 partitions¹².

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 # 115

A company uses AWS Glue Data Catalog to index data that is uploaded to an Amazon S3 bucket every day. The company uses a daily batch processes in an extract, transform, and load (ETL) pipeline to upload data from external sources into the S3 bucket.

The company runs a daily report on the S3 dat

a. Some days, the company runs the report before all the daily data has been uploaded to the S3 bucket. A data engineer must be

able to send a message that identifies any incomplete data to an existing Amazon Simple Notification Service (Amazon SNS) topic. Which solution will meet this requirement with the LEAST operational overhead?

- A. Create AWS Lambda functions that run data quality queries on the columns data type and the presence of null values. Orchestrate the ETL pipeline by using an AWS Step Functions workflow that runs the Lambda functions. Configure the Step Functions workflow to send an email notification that informs the data engineer about the incomplete datasets to the SNS topic.
- B. Create data quality checks on the source datasets that the daily reports use. Create a new Amazon EMR cluster. Use Apache Spark SQL to create Apache Spark jobs in the EMR cluster that run data quality queries on the columns data type and the presence of null values. Orchestrate the ETL pipeline by using an AWS Step Functions workflow. Configure the workflow to send an email notification that informs the data engineer about the incomplete datasets to the SNS topic.
- C. Create data quality checks for the source datasets that the daily reports use. Create a new AWS managed Apache Airflow cluster. Run the data quality checks by using Airflow tasks that run data quality queries on the columns data type and the presence of null values. Configure Airflow Directed Acyclic Graphs (DAGs) to send an email notification that informs the data engineer about the incomplete datasets to the SNS topic.
- **D. Create data quality checks on the source datasets that the daily reports use. Create data quality actions by using AWS Glue workflows to confirm the completeness and consistency of the datasets. Configure the data quality actions to create an event in Amazon EventBridge if a dataset is incomplete. Configure EventBridge to send the event that informs the data engineer about the incomplete datasets to the Amazon SNS topic.**

Answer: D

Explanation:

AWS Glue workflows are designed to orchestrate the ETL pipeline, and you can create data quality checks to ensure the uploaded datasets are complete before running reports. If there is an issue with the data, AWS Glue workflows can trigger an Amazon EventBridge event that sends a message to an SNS topic.

AWS Glue Workflows:

AWS Glue workflows allow users to automate and monitor complex ETL processes. You can include data quality actions to check for null values, data types, and other consistency checks.

In the event of incomplete data, an EventBridge event can be generated to notify via SNS.

Reference:

Alternatives Considered:

A (Airflow cluster): Managed Airflow introduces more operational overhead and complexity compared to Glue workflows.

B (EMR cluster): Setting up an EMR cluster is also more complex compared to the Glue-centric solution.

D (Lambda functions): While Lambda functions can work, using Glue workflows offers a more integrated and lower operational overhead solution.

AWS Glue Workflow Documentation

NEW QUESTION # 116

.....

As we all know, it is a must for all of the candidates to pass the exam if they want to get the related Data-Engineer-Associate certification which serves as the best evidence for them to show their knowledge and skills. If you want to simplify the preparation process, here comes a piece of good news for you. Our Data-Engineer-Associate Exam Question has been widely praised by all of our customers in many countries and our company has become the leader in this field. Our Data-Engineer-Associate exam questions are very accurate for you to pass the Data-Engineer-Associate exam. Once you buy our Data-Engineer-Associate practice guide, you will have high pass rate.

Data-Engineer-Associate Authorized Exam Dumps: <https://www.examsreviews.com/Data-Engineer-Associate-pass4sure-exam-review.html>

Amazon Reliable Data-Engineer-Associate Practice Questions If you fail the exam we will unconditionally refund the full dumps cost to you, Amazon Reliable Data-Engineer-Associate Practice Questions You will receive the e-mails only from us, Amazon Reliable Data-Engineer-Associate Practice Questions Our study materials are different from common study materials, which can motivate you to concentrate on study, But you don't need to be anxious about this issue once you study with our Data-Engineer-Associate latest dumps: AWS Certified Data Engineer - Associate (DEA-C01).

This is the S" part of the user logon process, An introduction to the Data-Engineer-Associate table, If you fail the exam we will unconditionally refund the full dumps cost to you, You will receive the e-mails only from us.

