

Data-Engineer-Associate Study Materials Review | Real Data-Engineer-Associate Dumps



2026 Latest Dumps Tests Data-Engineer-Associate PDF Dumps and Data-Engineer-Associate Exam Engine Free Share:
https://drive.google.com/open?id=1K2zSWvowz6tEmmSIAg_MgJDZ0nM5joQw

Each candidate will enjoy one-year free update after purchased our Data-Engineer-Associate dumps collection. We will send you the latest Data-Engineer-Associate dumps pdf to your email immediately once we have any updating about the certification exam. And there are free demo of Data-Engineer-Associate Exam Questions in our website for your reference. Our Amazon exam torrent is the best partner for your exam preparation.

As soon as you enter the learning interface of our system and start practicing our Amazon Data-Engineer-Associate learning materials on our Windows software, you will find small buttons on the interface. These buttons show answers, and you can choose to hide answers during your learning of our Amazon Data-Engineer-Associate Exam Quiz so as not to interfere with your learning process.

>> Data-Engineer-Associate Study Materials Review <<

Hot Data-Engineer-Associate Study Materials Review | Latest Real Data-Engineer-Associate Dumps: AWS Certified Data Engineer - Associate (DEA-C01) 100% Pass

In spite of the high-quality of our Amazon Data-Engineer-Associate study braindumps, our after-sales service can be the most attractive project in our Data-Engineer-Associate guide questions. We have free online service which means that if you have any trouble using our Amazon Data-Engineer-Associate Learning Materials or operate different versions on the platform mistakenly, we can provide help for you remotely in the shortest time.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q117-Q122):

NEW QUESTION # 117

A data engineer must use AWS services to ingest a dataset into an Amazon S3 data lake. The data engineer profiles the dataset and discovers that the dataset contains personally identifiable information (PII). The data engineer must implement a solution to profile the dataset and obfuscate the PII.

Which solution will meet this requirement with the LEAST operational effort?

- A. Ingest the dataset into Amazon DynamoDB. Create an AWS Lambda function to identify and obfuscate the PII in the DynamoDB table and to transform the data. Use the same Lambda function to ingest the data into the S3 data lake.
- B. Use an Amazon Kinesis Data Firehose delivery stream to process the dataset. Create an AWS Lambda transform function to identify the PII. Use an AWS SDK to obfuscate the PII. Set the S3 data lake as the target for the delivery stream.
- C. Use the Detect PII transform in AWS Glue Studio to identify the PII. Create a rule in AWS Glue Data Quality to obfuscate the PII. Use an AWS Step Functions state machine to orchestrate a data pipeline to ingest the data into the S3 data lake.

lake.

- D. Use the Detect PII transform in AWS Glue Studio to identify the PII. Obfuscate the PII. Use an AWS Step Functions state machine to orchestrate a data pipeline to ingest the data into the S3 data lake.

Answer: C

Explanation:

AWS Glue is a fully managed service that provides a serverless data integration platform for data preparation, data cataloging, and data loading. AWS Glue Studio is a graphical interface that allows you to easily author, run, and monitor AWS Glue ETL jobs.

AWS Glue Data Quality is a feature that enables you to validate, cleanse, and enrich your data using predefined or custom rules.

AWS Step Functions is a service that allows you to coordinate multiple AWS services into serverless workflows.

Using the Detect PII transform in AWS Glue Studio, you can automatically identify and label the PII in your dataset, such as names, addresses, phone numbers, email addresses, etc. You can then create a rule in AWS Glue Data Quality to obfuscate the PII, such as masking, hashing, or replacing the values with dummy data.

You can also use other rules to validate and cleanse your data, such as checking for null values, duplicates, outliers, etc. You can then use an AWS Step Functions state machine to orchestrate a data pipeline to ingest the data into the S3 data lake. You can use AWS Glue DataBrew to visually explore and transform the data, AWS Glue crawlers to discover and catalog the data, and AWS Glue jobs to load the data into the S3 data lake.

This solution will meet the requirement with the least operational effort, as it leverages the serverless and managed capabilities of AWS Glue, AWS Glue Studio, AWS Glue Data Quality, and AWS Step Functions.

You do not need to write any code to identify or obfuscate the PII, as you can use the built-in transforms and rules in AWS Glue Studio and AWS Glue Data Quality. You also do not need to provision or manage any servers or clusters, as AWS Glue and AWS Step Functions scale automatically based on the demand.

The other options are not as efficient as using the Detect PII transform in AWS Glue Studio, creating a rule in AWS Glue Data Quality, and using an AWS Step Functions state machine. Using an Amazon Kinesis Data Firehose delivery stream to process the dataset, creating an AWS Lambda transform function to identify the PII, using an AWS SDK to obfuscate the PII, and setting the S3 data lake as the target for the delivery stream will require more operational effort, as you will need to write and maintain code to identify and obfuscate the PII, as well as manage the Lambda function and its resources. Using the Detect PII transform in AWS Glue Studio to identify the PII, obfuscating the PII, and using an AWS Step Functions state machine to orchestrate a data pipeline to ingest the data into the S3 data lake will not be as effective as creating a rule in AWS Glue Data Quality to obfuscate the PII, as you will need to manually obfuscate the PII after identifying it, which can be error-prone and time-consuming. Ingesting the dataset into Amazon DynamoDB, creating an AWS Lambda function to identify and obfuscate the PII in the DynamoDB table and to transform the data, and using the same Lambda function to ingest the data into the S3 data lake will require more operational effort, as you will need to write and maintain code to identify and obfuscate the PII, as well as manage the Lambda function and its resources. You will also incur additional costs and complexity by using DynamoDB as an intermediate data store, which may not be necessary for your use case. References:

AWS Glue

AWS Glue Studio

AWS Glue Data Quality

[AWS Step Functions]

[AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide], Chapter 6: Data Integration and Transformation, Section 6.1: AWS Glue

NEW QUESTION # 118

A healthcare company uses Amazon Kinesis Data Streams to stream real-time health data from wearable devices, hospital equipment, and patient records.

A data engineer needs to find a solution to process the streaming data. The data engineer needs to store the data in an Amazon Redshift Serverless warehouse. The solution must support near real-time analytics of the streaming data and the previous day's data. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the Amazon Aurora zero-ETL integration with Amazon Redshift.
- B. Load the data into Amazon S3. Use the COPY command to load the data into Amazon Redshift.
- **C. Use the streaming ingestion feature of Amazon Redshift.**
- D. Load data into Amazon Kinesis Data Firehose. Load the data into Amazon Redshift.

Answer: C

Explanation:

The streaming ingestion feature of Amazon Redshift enables you to ingest data from streaming sources, such as Amazon Kinesis Data Streams, into Amazon Redshift tables in near real-time. You can use the streaming ingestion feature to process the streaming

data from the wearable devices, hospital equipment, and patient records. The streaming ingestion feature also supports incremental updates, which means you can append new data or update existing data in the Amazon Redshift tables. This way, you can store the data in an Amazon Redshift Serverless warehouse and support near real-time analytics of the streaming data and the previous day's data. This solution meets the requirements with the least operational overhead, as it does not require any additional services or components to ingest and process the streaming data. The other options are either not feasible or not optimal. Loading data into Amazon Kinesis Data Firehose and then into Amazon Redshift (option A) would introduce additional latency and cost, as well as require additional configuration and management. Loading data into Amazon S3 and then using the COPY command to load the data into Amazon Redshift (option C) would also introduce additional latency and cost, as well as require additional storage space and ETL logic. Using the Amazon Aurora zero-ETL integration with Amazon Redshift (option D) would not work, as it requires the data to be stored in Amazon Aurora first, which is not the case for the streaming data from the healthcare company. References: Using streaming ingestion with Amazon Redshift AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 3: Data Ingestion and Transformation, Section 3.5: Amazon Redshift Streaming Ingestion

NEW QUESTION # 119

A company uses Amazon S3 to store data and Amazon QuickSight to create visualizations. The company has an S3 bucket in an AWS account named Hub-Account. The S3 bucket is encrypted by an AWS Key Management Service (AWS KMS) key. The company's QuickSight instance is in a separate account named BI-Account. The company updates the S3 bucket policy to grant access to the QuickSight service role. The company wants to enable cross-account access to allow QuickSight to interact with the S3 bucket. Which combination of steps will meet this requirement? (Select TWO.)

- A. Add the KMS key as a resource that the QuickSight service role can access.
- B. Add the S3 bucket as a resource that the QuickSight service role can access.
- C. Add an IAM policy to the QuickSight service role to give QuickSight access to the KMS key that encrypts the S3 bucket.
- D. Use the existing AWS KMS key to encrypt connections from QuickSight to the S3 bucket.
- E. Use AWS Resource Access Manager (AWS RAM) to share the S3 bucket with the BI-Account account.

Answer: A,C

Explanation:

Problem Analysis:

The company needs cross-account access to allow QuickSight in BI-Account to interact with an S3 bucket in Hub-Account. The bucket is encrypted with an AWS KMS key.

Appropriate permissions must be set for both S3 access and KMS decryption.

Key Considerations:

QuickSight requires IAM permissions to access S3 data and decrypt files using the KMS key.

Both S3 and KMS permissions need to be properly configured across accounts.

Solution Analysis:

Option A: Use Existing KMS Key for Encryption

While the existing KMS key is used for encryption, it must also grant decryption permissions to QuickSight.

Option B: Add S3 Bucket to QuickSight Role

Granting S3 bucket access to the QuickSight service role is necessary for cross-account access.

Option C: AWS RAM for Bucket Sharing

AWS RAM is not required; bucket policies and IAM roles suffice for granting cross-account access.

Option D: IAM Policy for KMS Access

QuickSight's service role in BI-Account needs explicit permissions to use the KMS key for decryption.

Option E: Add KMS Key as Resource for Role

The KMS key must explicitly list the QuickSight role as an entity that can access it.

Implementation Steps:

S3 Bucket Policy in Hub-Account: Add a policy to the S3 bucket granting the QuickSight service role access:

json

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "AWS": "arn:aws:iam::<BI-Account-ID>:role/service-role/QuickSightRole" },
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::<Bucket-Name>/*"
```

```
}  
]  
}
```

KMS Key Policy in Hub-Account: Add permissions for the QuickSight role:

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Principal": { "AWS": "arn:aws:iam:<BI-Account-ID>:role/service-role/QuickSightRole" },  
      "Action": [  
        "kms:Decrypt",  
        "kms:DescribeKey"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

IAM Policy for QuickSight Role in BI-Account: Attach the following policy to the QuickSight service role:

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "s3:GetObject",  
        "kms:Decrypt"  
      ],  
      "Resource": [  
        "arn:aws:s3:::<Bucket-Name>/*",  
        "arn:aws:kms:<region>:<Hub-Account-ID>:key/<KMS-Key-ID>"  
      ]  
    }  
  ]  
}
```

Setting Up Cross-Account S3 Access

AWS KMS Key Policy Examples

Amazon QuickSight Cross-Account Access

NEW QUESTION # 120

A company has a data pipeline that processes transaction data in real time. The company needs a notification system that alerts different teams based on the type of processing error without any delay. For security-related errors, the system must immediately notify the security team. For data validation errors, the system must notify the data quality team. For system errors, the system must notify the operations team.

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

- A. Set up Amazon CloudWatch alarms with different metrics for each error type. Invoke a different Amazon Simple Notification Service (Amazon SNS) notification each time a metrics threshold is crossed.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic with an AWS Lambda function subscriber that evaluates the error type and forwards the error to the appropriate email addresses.
- C. Use Amazon Simple Queue Service (Amazon SQS) with message attributes to categorize errors. Allow each team to poll their respective SQS queue for relevant errors.
- **D. Configure Amazon EventBridge rules with distinct event patterns for each error type. Route each error type to a dedicated Amazon Simple Notification Service (Amazon SNS) topic for team-specific alerts.**

Answer: D

Explanation:

Option B is the best answer because Amazon EventBridge is designed to route events based on matching rules, and each rule can use a specific event pattern to match only the relevant error type and then send that event to the correct target. AWS states that

EventBridge event buses receive events and deliver them to targets, and that rules can match on event details and route matching events accordingly. This makes it ideal for real-time error classification with very little custom code. An SNS topic per team then provides immediate push-based notification delivery.

Option A adds unnecessary operational overhead because Lambda would be doing routing logic that EventBridge rules already handle natively. Option C is weaker because SQS is pull-based; teams must poll queues, which is not the best fit for "without any delay" notifications. AWS documentation explicitly describes SQS consumption through polling. Option D is also not the best fit because CloudWatch alarms are threshold-based metric monitors, not event routers for per-message error classification.

This aligns with the study guide's emphasis on using managed AWS services to automate and orchestrate pipeline actions with the least overhead.

NEW QUESTION # 121

A data engineer needs to build an extract, transform, and load (ETL) job. The ETL job will process daily incoming .csv files that users upload to an Amazon S3 bucket. The size of each S3 object is less than 100 MB.

Which solution will meet these requirements MOST cost-effectively?

- A. Write an AWS Glue Python shell job. Use pandas to transform the data.
- B. Write a PySpark ETL script. Host the script on an Amazon EMR cluster.
- C. Write a custom Python application. Host the application on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- D. Write an AWS Glue PySpark job. Use Apache Spark to transform the data.

Answer: A

Explanation:

AWS Glue is a fully managed serverless ETL service that can handle various data sources and formats, including .csv files in Amazon S3. AWS Glue provides two types of jobs: PySpark and Python shell. PySpark jobs use Apache Spark to process large-scale data in parallel, while Python shell jobs use Python scripts to process small-scale data in a single execution environment. For this requirement, a Python shell job is more suitable and cost-effective, as the size of each S3 object is less than 100 MB, which does not require distributed processing. A Python shell job can use pandas, a popular Python library for data analysis, to transform the .csv data as needed. The other solutions are not optimal or relevant for this requirement. Writing a custom Python application and hosting it on an Amazon EKS cluster would require more effort and resources to set up and manage the Kubernetes environment, as well as to handle the data ingestion and transformation logic. Writing a PySpark ETL script and hosting it on an Amazon EMR cluster would also incur more costs and complexity to provision and configure the EMR cluster, as well as to use Apache Spark for processing small data files. Writing an AWS Glue PySpark job would also be less efficient and economical than a Python shell job, as it would involve unnecessary overhead and charges for using Apache Spark for small data files. Reference:

AWS Glue

Working with Python Shell Jobs

pandas

[AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide]

NEW QUESTION # 122

.....

In order to serve you better, we have a complete system for Data-Engineer-Associate exam materials. We offer you free demo to have a try before buying, so that you can have a better understanding of what you are going to buy. If you want the Data-Engineer-Associate exam dumps after trying, just add to cart and pay for it. You will receive the downloading link and password within ten minutes and you can start your learning right now. If you don't receive, contact us, and we will check it for you. After you purchasing Data-Engineer-Associate Exam Materials, we also have after-sales, and if you have any questions, you can consult us.

Real Data-Engineer-Associate Dumps: <https://www.dumpstests.com/Data-Engineer-Associate-latest-test-dumps.html>

They work collectively and strive hard to ensure the top quality of DumpsTests Data-Engineer-Associate exam practice questions all the time, This feature will boost your confidence when taking the Amazon Data-Engineer-Associate certification exam, Amazon Data-Engineer-Associate Study Materials Review We are look forward to become your learning partner in the near future, Amazon Data-Engineer-Associate Study Materials Review First of all, we have various kinds of study guide for customers to buy.

Colligan is a popular writer and presenter on Data-Engineer-Associate Training Questions Internet technology topics and frequently speaks online, on the air, and before audiences about his passions, No senior executive can be expected Data-Engineer-Associate Certified to know the entire competitive terrain well enough to correctly call all the shots.

