

Reliable Data-Engineer-Associate Test Preparation & Data-Engineer-Associate Related Exams



P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by Fast2test:
https://drive.google.com/open?id=1jRyZBoVUpHAF8T9VP_m45AuBINiMM3gK

Facing the incoming Amazon Data-Engineer-Associate Exam, you may feel stained and anxious, suspicious whether you could pass the exam smoothly and successfully. Actually, you must not impoverish your ambition. Our suggestions are never boggle at difficulties. It is your right time to make your mark.

Amazon Data-Engineer-Associate Exam Syllabus Topics:

Section	Weight	Objectives
Data Security and Governance	18%	- Protect sensitive data - Encrypt data at rest and in transit - Implement access control and authentication • 1. IAM, Lake Formation, resource policies - Enforce compliance and data governance • 1. Data lineage, audit logging, regulatory requirements
Data Operations and Support	22%	- Backup, restore, and disaster recovery - Monitor and troubleshoot data pipelines • 1. CloudWatch, X-Ray, logging and metrics - Automate operational tasks - Ensure reliability and scalability
Data Store Management	26%	- Design and implement data storage solutions • 1. Data lakes, data warehouses, databases • 2. S3, Redshift, DynamoDB, RDS, Lake Formation - Optimize storage performance and cost - Manage data lifecycle and storage tiers

Data Ingestion and Transformation	34%	- Ingest data from various sources • 1. Use services like Kinesis, DMS, Glue, S3 • 2. Batch and streaming data ingestion - Transform and enrich data • 1. Use Spark, EMR, Step Functions • 2. Orchestrate data pipelines • 3. Apply data processing logic - Implement data quality and validation
-----------------------------------	-----	--

>> **Reliable Data-Engineer-Associate Test Preparation** <<

Data-Engineer-Associate Related Exams - Data-Engineer-Associate Certification Dumps

Many people want to be the competent people which can excel in the job in some area and be skillful in applying the knowledge to the practical working in some industry. But the thing is not so easy for them they need many efforts to achieve their goals. Passing the Data-Engineer-Associate test certification can make them become that kind of people and if you are one of them buying our Data-Engineer-Associate study materials will help you pass the Data-Engineer-Associate test smoothly with few efforts needed.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q35-Q40):

NEW QUESTION # 35

A company stores details about transactions in an Amazon S3 bucket. The company wants to log all writes to the S3 bucket into another S3 bucket that is in the same AWS Region.

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

- A. Configure an S3 Event Notifications rule for all activities on the transactions S3 bucket to invoke an AWS Lambda function. Program the Lambda function to write the events to the logs S3 bucket.
- **B. Create a trail of data events in AWS CloudTrail. Configure the trail to receive data from the transactions S3 bucket. Specify an empty prefix and write-only events. Specify the logs S3 bucket as the destination bucket.**
- C. Create a trail of management events in AWS CloudTrail. Configure the trail to receive data from the transactions S3 bucket. Specify an empty prefix and write-only events. Specify the logs S3 bucket as the destination bucket.
- D. Configure an S3 Event Notifications rule for all activities on the transactions S3 bucket to invoke an AWS Lambda function. Program the Lambda function to write the event to Amazon Kinesis Data Firehose. Configure Kinesis Data Firehose to write the event to the logs S3 bucket.

Answer: B

Explanation:

This solution meets the requirement of logging all writes to the S3 bucket into another S3 bucket with the least operational effort. AWS CloudTrail is a service that records the API calls made to AWS services, including Amazon S3. By creating a trail of data events, you can capture the details of the requests that are made to the transactions S3 bucket, such as the requester, the time, the IP address, and the response elements. By specifying an empty prefix and write-only events, you can filter the data events to only include the ones that write to the bucket. By specifying the logs S3 bucket as the destination bucket, you can store the CloudTrail logs in another S3 bucket that is in the same AWS Region. This solution does not require any additional coding or configuration, and it is more scalable and reliable than using S3 Event Notifications and Lambda functions. Reference:

Logging Amazon S3 API calls using AWS CloudTrail

Creating a trail for data events

Enabling Amazon S3 server access logging

NEW QUESTION # 36

A company wants to migrate data from an Amazon RDS for PostgreSQL DB instance in the eu-east-1 Region of an AWS account named Account_A.

The company will migrate the data to an Amazon Redshift cluster in the eu-west-1 Region of an AWS account named Account_B. Which solution will give AWS Database Migration Service (AWS DMS) the ability to replicate data between two data stores?

- A. Set up an AWS DMS replication instance in Account_A in eu-east-1.
- B. Set up an AWS DMS replication instance in a new AWS account in eu-west-1
- C. Set up an AWS DMS replication instance in Account_B in eu-east-1.
- **D. Set up an AWS DMS replication instance in Account_B in eu-west-1.**

Answer: D

Explanation:

To migrate data from an Amazon RDS for PostgreSQL DB instance in the eu-east-1 Region (Account_A) to an Amazon Redshift cluster in the eu-west-1 Region (Account_B), AWS DMS needs a replication instance located in the target region (in this case, eu-west-1) to facilitate the data transfer between regions.

Option A: Set up an AWS DMS replication instance in Account_B in eu-west-1.

Placing the DMS replication instance in the target account and region (Account_B in eu-west-1) is the most efficient solution. The replication instance can connect to the source RDS PostgreSQL in eu-east-1 and migrate the data to the Redshift cluster in eu-west-1. This setup ensures data is replicated across AWS accounts and regions.

Options B, C, and D place the replication instance in either the wrong account or region, which increases complexity without adding any benefit.

Reference:

AWS Database Migration Service (DMS) Documentation
Cross-Region and Cross-Account Replication

NEW QUESTION # 37

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 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.
- **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 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.
- 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 # 38

A company uses Amazon Redshift for its data warehouse. A data engineer must query a table named orders.

complete_orders_history, which contains 100 columns. The query must return all columns except columns named company_id and unique_system_id.

Which Amazon Redshift SQL statement will meet this requirement?

- **A. SELECT * EXCLUDE company_id, unique_system_id FROM orders.complete_orders_history;**
- B. SELECT * TRUNCATE company_id, unique_system_id FROM orders.complete_orders_history;
- C. SELECT * EXCEPT company_id, unique_system_id FROM orders.complete_orders_history;

- D. `SELECT * NOT IN company_id, unique_system_id FROM orders.complete_orders_history;`

Answer: A

Explanation:

Option A is correct because Amazon Redshift supports the EXCLUDE clause in the SELECT list for exactly this use case: returning all columns from a wide table except a small set of unwanted columns. AWS documentation states that "EXCLUDE column_list names the columns that are excluded from the query results" and notes that the option is helpful when only a subset of columns must be excluded from a wide table. That matches this question precisely because the table contains 100 columns and only company_id and unique_system_id need to be omitted.

Option B is invalid SQL syntax in Amazon Redshift because NOT IN is used in predicates, not in the select list to remove columns.

Option C is also incorrect because EXCEPT in Redshift is a set operator used between query result sets, not a select-list column exclusion feature. AWS documents EXCEPT alongside UNION and INTERSECT as result-set comparison operators, not as a way to omit columns from SELECT *.

Option D is not valid Redshift SQL syntax for column selection. Therefore, the correct and documented Redshift statement is `SELECT * EXCLUDE company_id, unique_system_id FROM orders.complete_orders_history;`

NEW QUESTION # 39

A data engineer needs to use AWS Step Functions to design an orchestration workflow. The workflow must parallel process a large collection of data files and apply a specific transformation to each file.

Which Step Functions state should the data engineer use to meet these requirements?

- A. Map state
- B. Parallel state
- C. Wait state
- D. Choice state

Answer: A

Explanation:

Option C is the correct answer because the Map state is designed to process a collection of data in parallel by applying the same transformation to each element. The Map state can invoke a nested workflow for each element, which can be another state machine or a Lambda function. The Map state will wait until all the parallel executions are completed before moving to the next state.

Option A is incorrect because the Parallel state is used to execute multiple branches of logic concurrently, not to process a collection of data. The Parallel state can have different branches with different logic and states, whereas the Map state has only one branch that is applied to each element of the collection.

Option B is incorrect because the Choice state is used to make decisions based on a comparison of a value to a set of rules. The Choice state does not process any data or invoke any nested workflows.

Option D is incorrect because the Wait state is used to delay the state machine from continuing for a specified time. The Wait state does not process any data or invoke any nested workflows.

References:

* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.3: AWS Step Functions, Pages 131-132

* Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.2: AWS Step Functions, Pages 9-10

* AWS Documentation Overview, AWS Step Functions Developer Guide, Step Functions Concepts, State Types, Map State, Pages 1-3

NEW QUESTION # 40

.....

The successful selection, development and Data-Engineer-Associate training of personnel are critical to our company's ability to provide a high standard of service to our customers and to respond their needs. That's the reason why we can produce the best Data-Engineer-Associate exam prep and can get so much praise in the international market. And we always believe first-class quality comes with the first-class service. You will find we are professional on the answering the questions on our Data-Engineer-Associate Study Materials.

Data-Engineer-Associate Related Exams: <https://www.fast2test.com/Data-Engineer-Associate-premium-file.html>

- BTW, DOWNLOAD part of Fast2test Data-Engineer-Associate dumps from Cloud Storage: https://drive.google.com/open?id=1jRyZBoVUpHAF8T9VP_m45AuBINiMM3gK

BTW, DOWNLOAD part of Fast2test Data-Engineer-Associate dumps from Cloud Storage: https://drive.google.com/open?id=1jRyZBoVUpHAF8T9VP_m45AuBINiMM3gK