100% Pass Quiz Amazon - Reliable MLA-C01 - AWS Certified Machine Learning Engineer - Associate Visual Cert Exam

Category	Associate
Exam duration	170 minutes
Exam format	85 questions
Cost	75 USD/10,000 JPY
Intended candidate	Individuals with at least 1 year of experience using Amazon SageMaker and other ML engineering AWS services
Candidate role examples	backend software developer, DevOps engineer, data engineer, MLOps engineer, and data scientist
Testing options	Pearson VUE testing center or online proctored exam
Languages offered	English, Japanese

DOWNLOAD the newest ActualCollection MLA-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1FV1n47ODY4KK9S_-p0XxM8vSGA47bArj

We have created a number of reports and learning functions for evaluating your proficiency for the AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam dumps. In preparation, you can optimize AWS Certified Machine Learning Engineer - Associate (MLA-C01) practice exam time and question type by utilizing our Amazon MLA-C01 Practice Test software. ActualCollection makes it easy to download AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam questions immediately after purchase.

Amazon MLA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	• ML Model Development: This section of the exam measures skills of Fraud Examiners and covers choosing and training machine learning models to solve business problems such as fraud detection. It includes selecting algorithms, using built-in or custom models, tuning parameters, and evaluating performance with standard metrics. The domain emphasizes refining models to avoid overfitting and maintaining version control to support ongoing investigations and audit trails.
Topic 2	• ML Solution Monitoring, Maintenance, and Security: This section of the exam measures skills of Fraud Examiners and assesses the ability to monitor machine learning models, manage infrastructure costs, and apply security best practices. It includes setting up model performance tracking, detecting drift, and using AWS tools for logging and alerts. Candidates are also tested on configuring access controls, auditing environments, and maintaining compliance in sensitive data environments like financial fraud detection.
Topic 3	• Deployment and Orchestration of ML Workflows: This section of the exam measures skills of Forensic Data Analysts and focuses on deploying machine learning models into production environments. It covers choosing the right infrastructure, managing containers, automating scaling, and orchestrating workflows through CI • CD pipelines. Candidates must be able to build and script environments that support consistent deployment and efficient retraining cycles in real-world fraud detection systems.
Topic 4	• Data Preparation for Machine Learning (ML): This section of the exam measures skills of Forensic Data Analysts and covers collecting, storing, and preparing data for machine learning. It focuses on understanding different data formats, ingestion methods, and AWS tools used to process and transform data. Candidates are expected to clean and engineer features, ensure data integrity, and address biases or compliance issues, which are crucial for preparing high-quality datasets in fraud analysis contexts.

>> MLA-C01 Visual Cert Exam <<

Why Do You Need to Trust on ActualCollection Amazon MLA-C01 Exam Questions?

Our technician will check the update of MLA-C01 exam questions every day, and we can guarantee that you can get a free update service from the date of purchase. Once you have any questions and doubts about the MLA-C01 exam questions we will provide you with our customer service before or after the sale, you can contact us if you have question or doubt about our MLA-C01 Exam Materials and the professional personnel can help you solve your issue about using MLA-C01 study materials.

Amazon AWS Certified Machine Learning Engineer - Associate Sample Questions (Q159-Q164):

NEW QUESTION # 159
An ML engineer needs to process thousands of existing CSV objects and new CSV objects that are uploaded.
The CSV objects are stored in a central Amazon S3 bucket and have the same number of columns. One of the columns is a transaction date. The ML engineer must query the data based on the transaction date.
Which solution will meet these requirements with the LEAST operational overhead?

- A. Use an Amazon Athena CREATE TABLE AS SELECT (CTAS) statement to create a table based on the transaction date from data in the central S3 bucket. Query the objects from the table.
- B. Create a new S3 bucket for processed data. Use AWS Glue for Apache Spark to create a job to query the CSV objects based on transaction date. Configure the job to store the results in the new S3 bucket.
  Query the objects from the new S3 bucket.
- C. Create a new S3 bucket for processed data. Set up S3 replication from the central S3 bucket to the new S3 bucket. Use S3 Object Lambda to query the objects based on transaction date.
- D. Create a new S3 bucket for processed data. Use Amazon Data Firehose to transfer the data from the central S3 bucket to the new S3 bucket. Configure Firehose to run an AWS Lambda function to query the data based on transaction date.

Answer: A

Explanation:
Scenario:The ML engineer needs a low-overhead solution to query thousands of existing and new CSV objects stored in Amazon S3 based on a transaction date.
Why Athena?
* Serverless:Amazon Athena is a serverless query service that allows direct querying of data stored in S3 using standard SQL, reducing operational overhead.
* Ease of Use:By using the CTAS statement, the engineer can create a table with optimized partitions based on the transaction date. Partitioning improves query performance and minimizes costs by scanning only relevant data.
* Low Operational Overhead:No need to manage or provision additional infrastructure. Athena integrates seamlessly with S3, and CTAS simplifies table creation and optimization.
Steps to Implement:
* Organize Data in S3:Store CSV files in a bucket in a consistent format and directory structure if possible.
* Configure Athena:Use the AWS Management Console or Athena CLI to set up Athena to point to the S3 bucket.
* Run CTAS Statement:

```
CREATE TABLE processed_data
WITH (
format = 'PARQUET',
external_location = 's3://processed-bucket/',
partitioned_by = ARRAY['transaction_date']
) AS
SELECT *
FROM input_data;
```

This creates a new table with data partitioned by transaction date.
* Query the Data:Use standard SQL queries to fetch data based on the transaction date.
References:
* Amazon Athena CTAS Documentation
* Partitioning Data in Athena

NEW QUESTION # 160

A company launches a feature that predicts home prices. An ML engineer trained a regression model using the SageMaker AI XGBoost algorithm. The model performs well on training data but underperforms on real- world validation data.
Which solution will improve the validation score with the LEAST implementation effort?

- A. Increase the num_round hyperparameter.
- B. Create a larger training dataset with more real-world data and retrain.
- C. Change the eval_metric from RMSE to Error.
- D. Increase the lambda hyperparameter.

Answer: D

Explanation:
This scenario indicates overfitting. AWS documentation for XGBoost recommends increasing the L2 regularization parameter (lambda) to reduce overfitting and improve generalization.
Increasing num_round worsens overfitting. Changing evaluation metrics does not change model behavior.
Collecting more data is effective but requires significant effort.
Regularization is a low-effort, high-impact fix.
Therefore, Option D is correct.

NEW QUESTION # 161
A company is building a web-based AI application by using Amazon SageMaker. The application will provide the following capabilities and features: ML experimentation, training, a central model registry, model deployment, and model monitoring.
The application must ensure secure and isolated use of training data during the ML lifecycle. The training data is stored in Amazon S3.
The company needs to run an on-demand workflow to monitor bias drift for models that are deployed to real- time endpoints from the application.
Which action will meet this requirement?

- A. Use SageMaker notebooks to compare the bias.
- B. Use AWS Glue Data Quality to monitor bias.
- C. Invoke an AWS Lambda function to pull the sagemaker-model-monitor-analyzer built-in SageMaker image.
- D. Configure the application to invoke an AWS Lambda function that runs a SageMaker Clarify job.

Answer: D

Explanation:
Monitoring bias drift in deployed machine learning models is crucial to ensure fairness and accuracy over time. Amazon SageMaker Clarify provides tools to detect bias in ML models, both during training and after deployment. To monitor bias drift for models deployed to real-time endpoints, an effective approach involves orchestrating SageMaker Clarify jobs using AWS Lambda functions.
Implementation Steps:
* Set Up Data Capture:
* Enable data capture on the SageMaker endpoint to record input data and model predictions. This captured data serves as the basis for bias analysis.
* Develop a Lambda Function:
* Create an AWS Lambda function configured to initiate a SageMaker Clarify job. This function will process the captured data to assess bias metrics.
* Schedule or Trigger the Lambda Function:
* Configure the Lambda function to run on-demand or at scheduled intervals using Amazon CloudWatch Events or EventBridge. This setup allows for regular bias monitoring as per the application's requirements.
* Analyze and Respond to Results:
* After each Clarify job completes, review the generated bias reports. If bias drift is detected, take appropriate actions, such as retraining the model or adjusting data preprocessing steps.
Advantages of This Approach:
* Automation: Utilizing AWS Lambda for orchestrating Clarify jobs enables automated and scalable bias monitoring without manual intervention.
* Cost-Effectiveness: AWS Lambda's serverless nature ensures that you only pay for the compute time consumed during the execution of the function, optimizing resource usage.
* Flexibility: The solution can be tailored to specific monitoring needs, allowing for adjustments in monitoring frequency and analysis parameters.

By implementing this solution, the company can effectively monitor bias drift in real-time, ensuring that the AI application maintains fairness and accuracy throughout its lifecycle.
References:
Bias drift for models in production - Amazon SageMaker
Schedule Bias Drift Monitoring Jobs - Amazon SageMaker

NEW QUESTION # 162
Hotspot Question
An ML engineer is building a generative AI application on Amazon Bedrock by using large language models (LLMs).
Select the correct generative AI term from the following list for each description. Each term should be selected one time or not at all. (Select three.)
- Embedding
- Retrieval Augmented Generation (RAG)
- Temperature
- Token

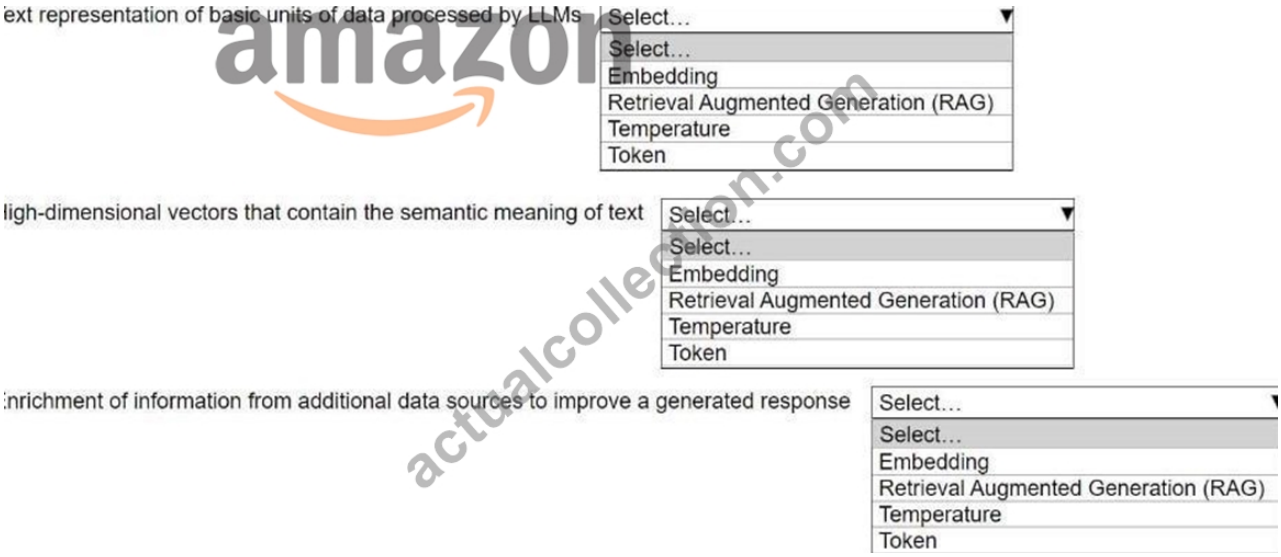

Answer:

Explanation:

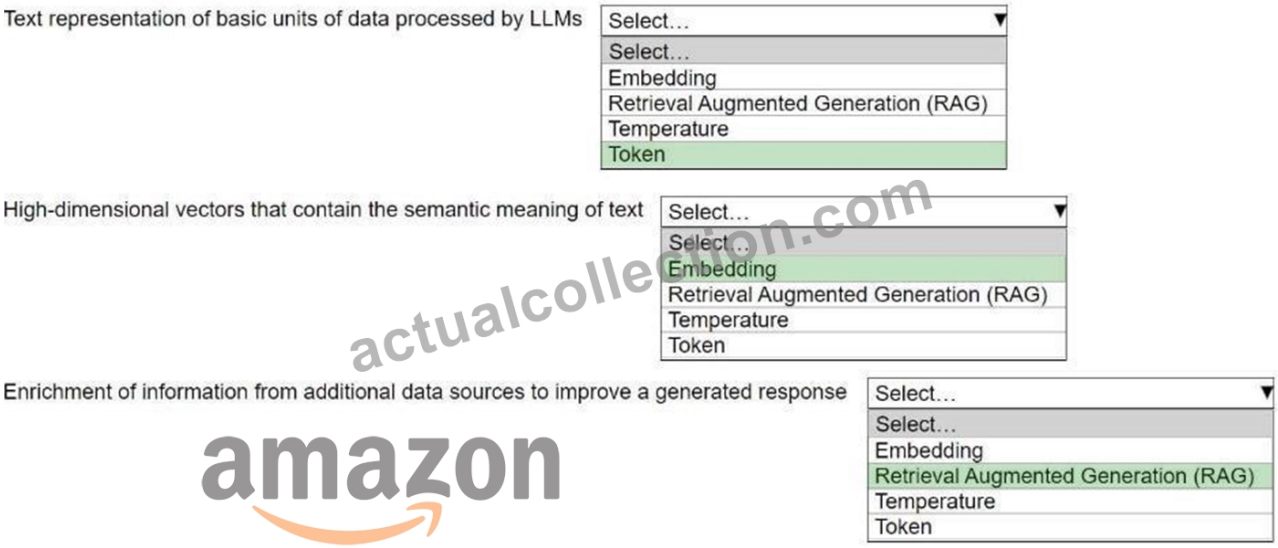

NEW QUESTION # 163
An ML engineer needs to ensure that a dataset complies with regulations for personally identifiable information (PII). The ML engineer will use the data to train an ML model on Amazon SageMaker instances. SageMaker must not use any of the PII.
Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use the Amazon Comprehend DetectPiiEntities API call to redact the PII from the data. Store the data in an Amazon S3 bucket. Access the S3 bucket from the SageMaker instances for model training.
- B. Use AWS Glue DataBrew to cleanse the dataset of PII. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system to the SageMaker instances for model training.
- C. Use the Amazon Comprehend DetectPiiEntities API call to redact the PII from the data. Store the data in an Amazon Elastic File System (Amazon EFS) file system. Mount the EFS file system to the SageMaker instances for model training.
- D. Use Amazon Macie for automatic discovery of PII in the data. Remove the PII. Store the data in an Amazon S3 bucket. Mount the S3 bucket to the SageMaker instances for model training.

Answer: A

NEW QUESTION # 164

.....