Amazon-Web-Services

Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate

NEW QUESTION 1
A company's ML engineer has deployed an ML model for sentiment analysis to an Amazon SageMaker endpoint. The ML engineer needs to explain to company stakeholders how the model makes predictions.
Which solution will provide an explanation for the model's predictions?

A. Use SageMaker Model Monitor on the deployed model.
B. Use SageMaker Clarify on the deployed model.
C. Show the distribution of inferences from A/ testing in Amazon CloudWatch.
D. Add a shadow endpoin
E. Analyze prediction differences on samples.

Answer: B

NEW QUESTION 2
A company is running ML models on premises by using custom Python scripts and proprietary datasets. The company is using PyTorch. The model building requires unique domain knowledge. The company needs to move the models to AWS.
Which solution will meet these requirements with the LEAST effort?

A. Use SageMaker built-in algorithms to train the proprietary datasets.
B. Use SageMaker script mode and premade images for ML frameworks.
C. Build a container on AWS that includes custom packages and a choice of ML frameworks.
D. Purchase similar production models through AWS Marketplace.

Answer: B

NEW QUESTION 3
A company has a team of data scientists who use Amazon SageMaker notebook instances to test ML models. When the data scientists need new permissions, the company attaches the permissions to each individual role that was created during the creation of the SageMaker notebook instance.
The company needs to centralize management of the team's permissions. Which solution will meet this requirement?

A. Create a single IAM role that has the necessary permission
B. Attach the role to each notebook instance that the team uses.
C. Create a single IAM grou
D. Add the data scientists to the grou
E. Associate the group with each notebook instance that the team uses.
F. Create a single IAM use
G. Attach the AdministratorAccess AWS managed IAM policy to the use
H. Configure each notebook instance to use the IAM user.
I. Create a single IAM grou
J. Add the data scientists to the grou
K. Create an IAM rol
L. Attach the AdministratorAccess AWS managed IAM policy to the rol
M. Associate the role with the grou
N. Associate the group with each notebook instance that the team uses.

Answer: A

NEW QUESTION 4
A company is gathering audio, video, and text data in various languages. The company needs to use a large language model (LLM) to summarize the gathered data that is in Spanish.
Which solution will meet these requirements in the LEAST amount of time?

A. Train and deploy a model in Amazon SageMaker to convert the data into English tex
B. Train and deploy an LLM in SageMaker to summarize the text.
C. Use Amazon Transcribe and Amazon Translate to convert the data into English tex
D. Use Amazon Bedrock with the Jurassic model to summarize the text.
E. Use Amazon Rekognition and Amazon Translate to convert the data into English tex
F. Use Amazon Bedrock with the Anthropic Claude model to summarize the text.
G. Use Amazon Comprehend and Amazon Translate to convert the data into English tex
H. Use Amazon Bedrock with the Stable Diffusion model to summarize the text.

Answer: B

NEW QUESTION 5
A company has trained and deployed an ML model by using Amazon SageMaker. The company needs to implement a solution to record and monitor all the API call events for the SageMaker endpoint. The solution also must provide a notification when the number of API call events breaches a threshold.
Use SageMaker Debugger to track the inferences and to report metrics. Create a custom rule to provide a notification when the threshold is breached.
Which solution will meet these requirements?

A. Use SageMaker Debugger to track the inferences and to report metric
B. Create a custom rule to provide a notification when the threshold is breached.
C. Use SageMaker Debugger to track the inferences and to report metric
D. Use the tensor_variance built-in rule to provide a notification when the threshold is breached.
E. Log all the endpoint invocation API events by using AWS CloudTrai
F. Use an Amazon CloudWatch dashboard for monitorin
G. Set up a CloudWatch alarm to provide notification when the threshold is breached.
H. Add the Invocations metric to an Amazon CloudWatch dashboard for monitorin

I. Set up a CloudWatch alarm to provide notification when the threshold is breached.

Answer: D

NEW QUESTION 6
HOTSPOT
A company wants to host an ML model on Amazon SageMaker. An ML engineer is configuring a continuous integration and continuous delivery (CI/CD) pipeline in AWS CodePipeline to deploy the model. The pipeline must run automatically when new training data for the model is uploaded to an Amazon S3 bucket.
Select and order the pipeline's correct steps from the following list. Each step should be selected one time or not at all. (Select and order three.)
• An S3 event notification invokes the pipeline when new data is uploaded.
• S3 Lifecycle rule invokes the pipeline when new data is uploaded.
• SageMaker retrains the model by using the data in the S3 bucket.
• The pipeline deploys the model to a SageMaker endpoint.
• The pipeline deploys the model to SageMaker Model Registry.

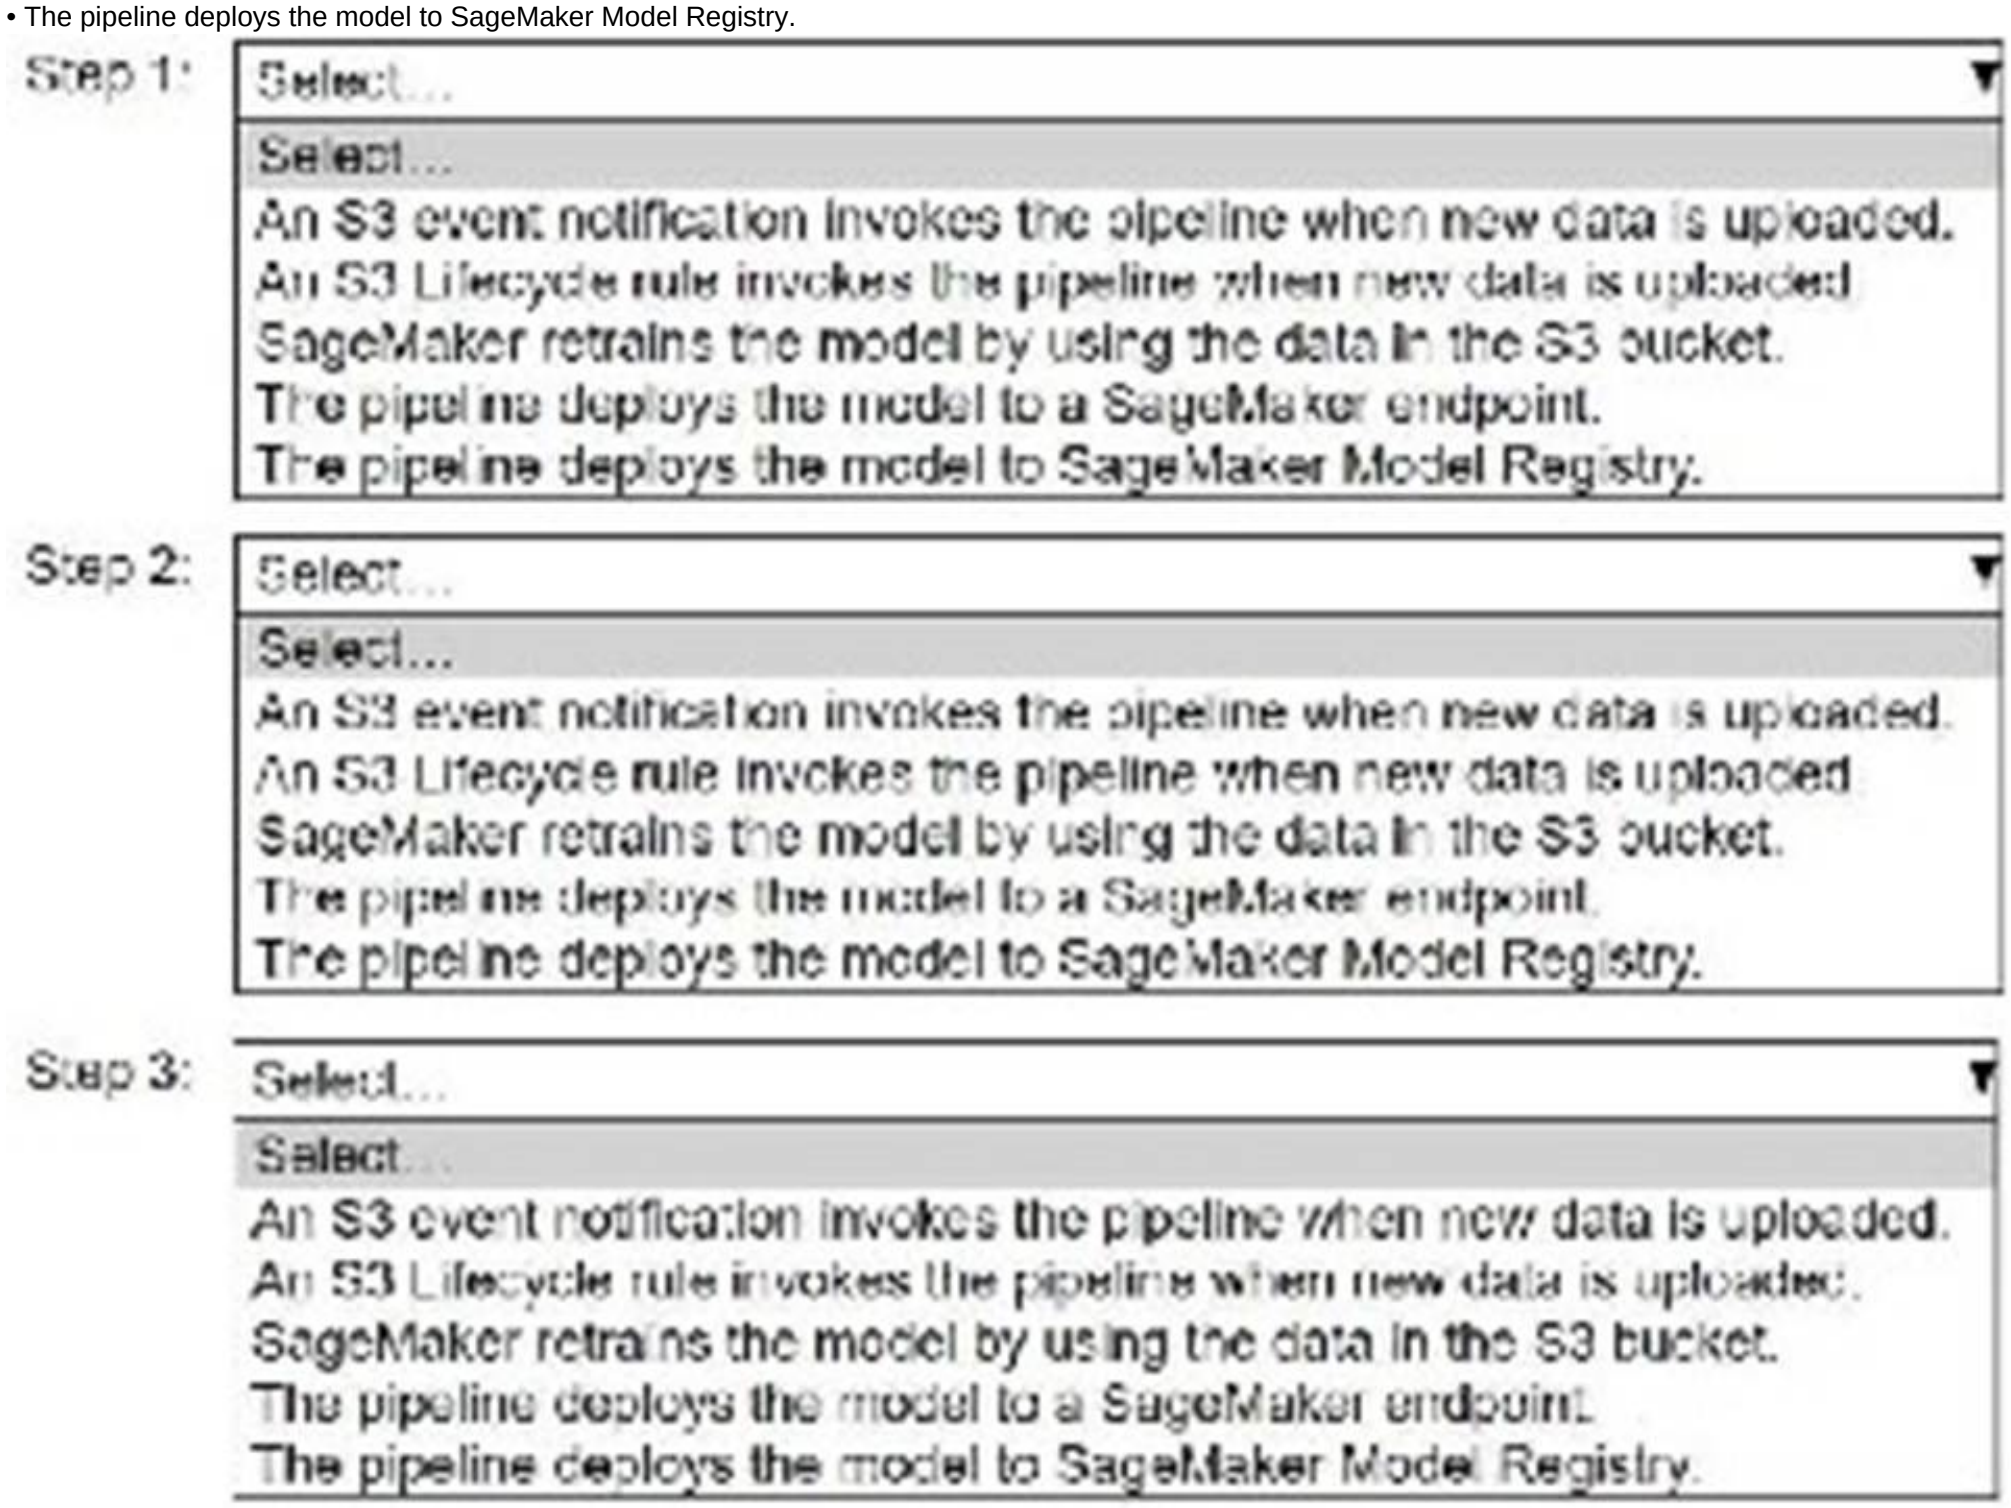

A. Mastered
B. Not Mastered

Answer: A

Explanation:

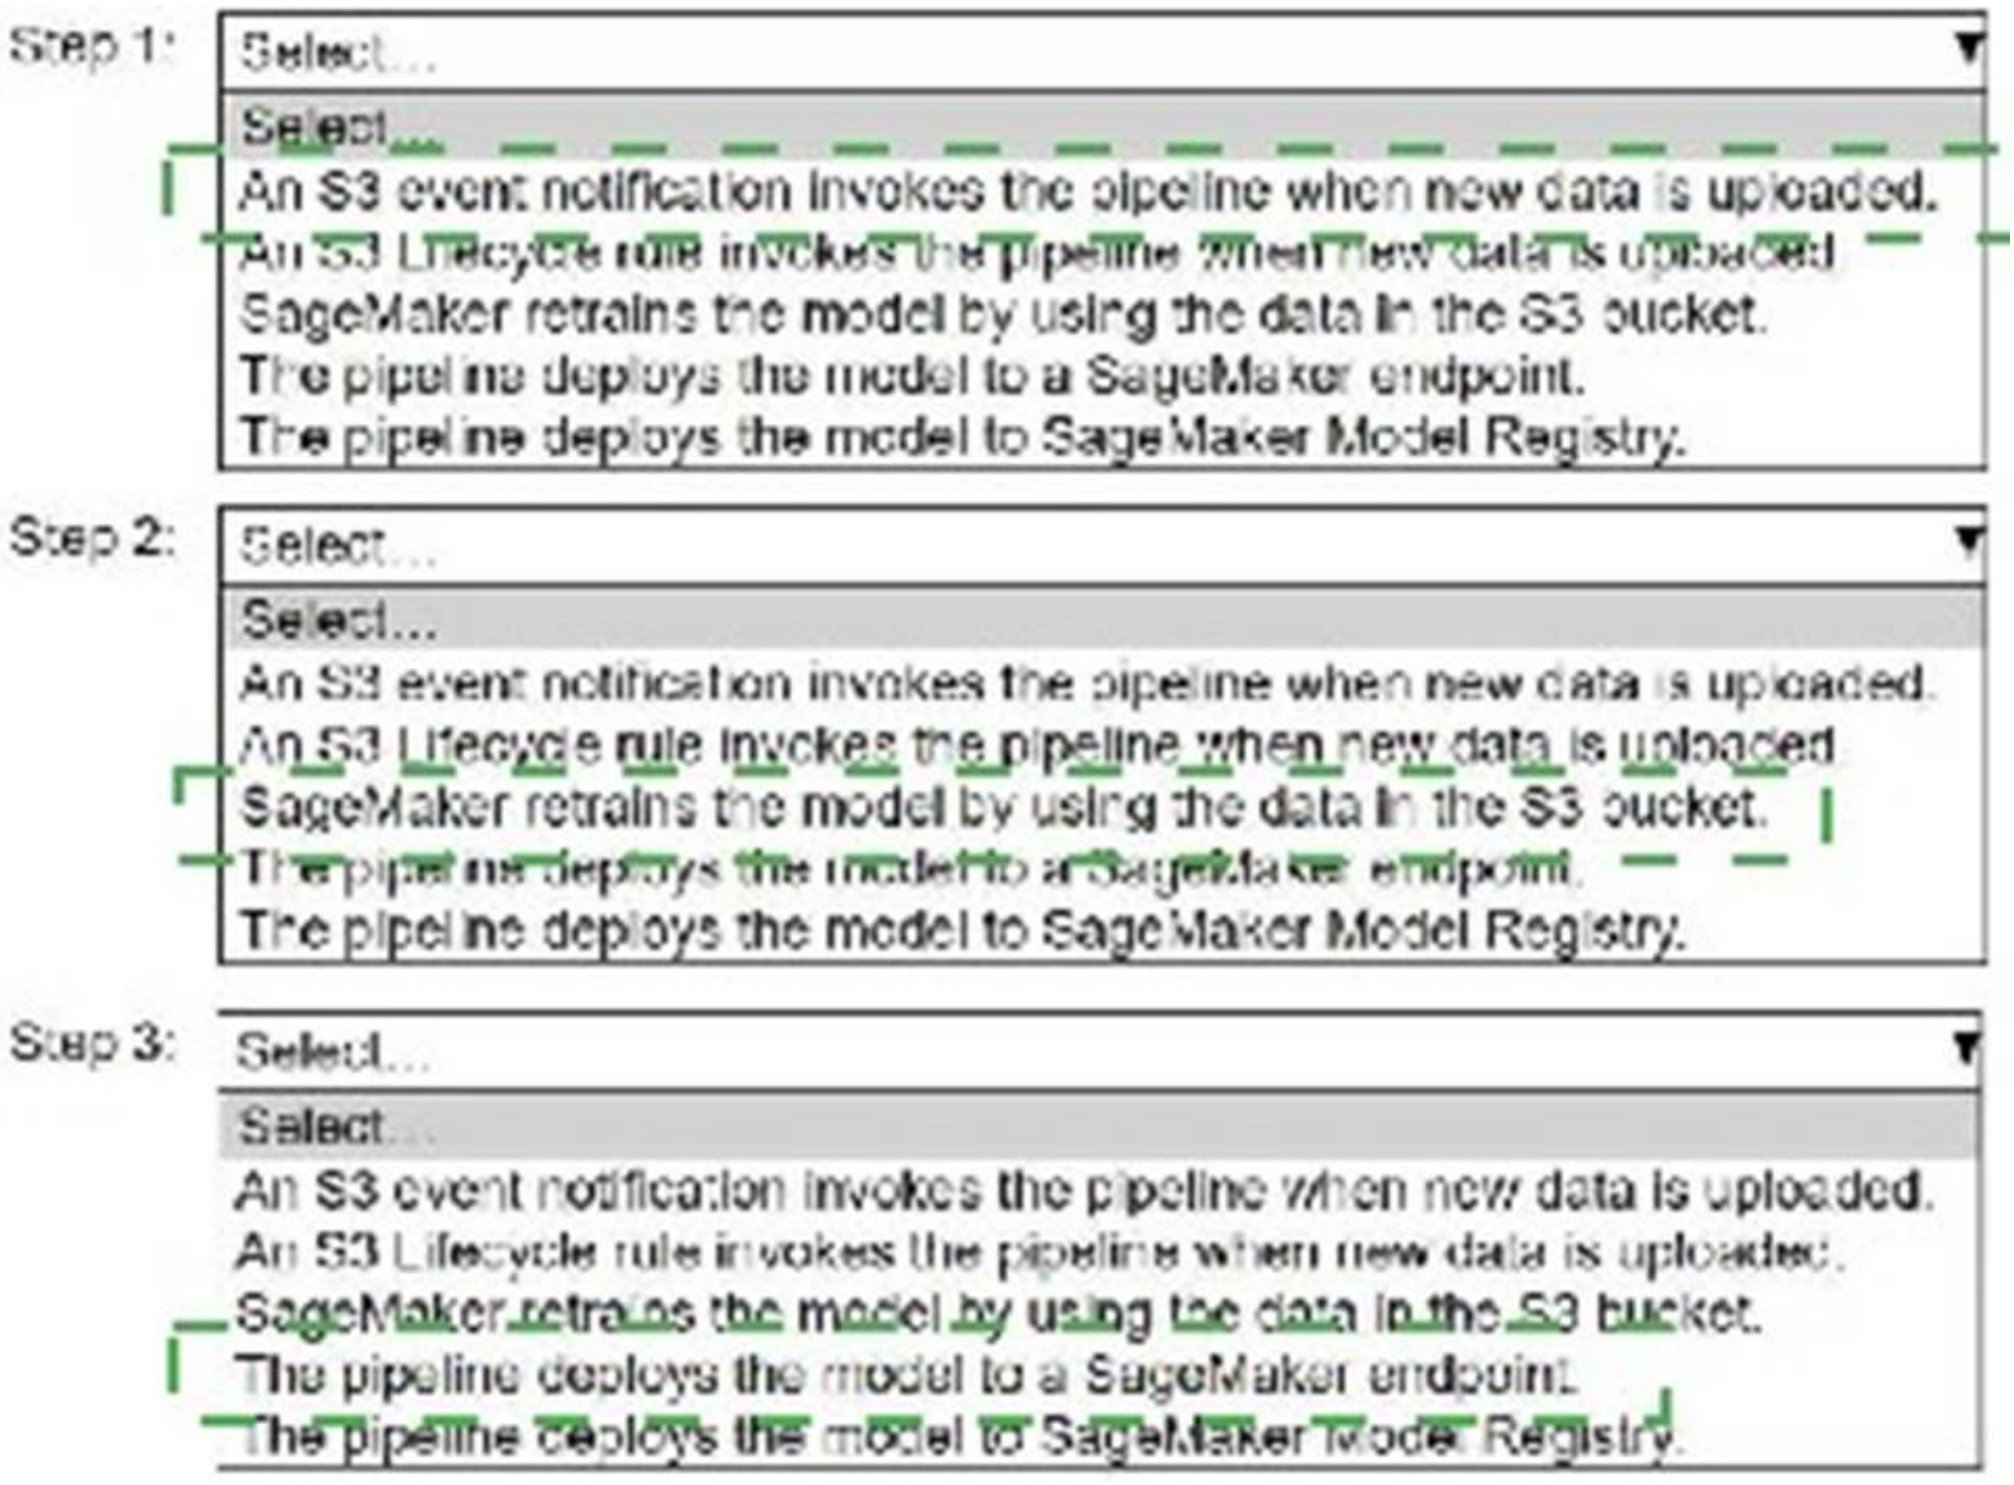

NEW QUESTION 7
An ML engineer is training a simple neural network model. The ML engineer tracks the performance of the model over time on a validation dataset. The model's performance improves substantially at first and then degrades after a specific number of epochs.
Which solutions will mitigate this problem? (Choose two.)

A. Enable early stopping on the model.
B. Increase dropout in the layers.
C. Increase the number of layers.
D. Increase the number of neurons.
E. Investigate and reduce the sources of model bias.

Answer: AB

NEW QUESTION 8
A company has used Amazon SageMaker to deploy a predictive ML model in production. The company is using SageMaker Model Monitor on the model. After a model update, an ML engineer notices data quality issues in the Model Monitor checks.
What should the ML engineer do to mitigate the data quality issues that Model Monitor has identified?

A. Adjust the model's parameters and hyperparameters.
B. Initiate a manual Model Monitor job that uses the most recent production data.
C. Create a new baseline from the latest datase
D. Update Model Monitor to use the new baseline for evaluations.
E. Include additional data in the existing training set for the mode
F. Retrain and redeploy the model.

Answer: C

NEW QUESTION 9
A company has a Retrieval Augmented Generation (RAG) application that uses a vector database to store embeddings of documents. The company must migrate the application to AWS and must implement a solution that provides semantic search of text files. The company has already migrated the text repository to an Amazon S3 bucket.
Which solution will meet these requirements?

A. Use an AWS Batch job to process the files and generate embedding
B. Use AWS Glue to store the embedding

C. Use SQL queries to perform the semantic searches.
D. Use a custom Amazon SageMaker notebook to run a custom script to generate embedding
E. Use SageMaker Feature Store to store the embedding
F. Use SQL queries to perform the semantic searches.
G. Use the Amazon Kendra S3 connector to ingest the documents from the S3 bucket into Amazon Kendr
H. Query Amazon Kendra to perform the semantic searches.
I. Use an Amazon Textract asynchronous job to ingest the documents from the S3 bucke
J. Query Amazon Textract to perform the semantic searches.

Answer: C

NEW QUESTION 10
An ML engineer needs to create data ingestion pipelines and ML model deployment pipelines on AWS. All the raw data is stored in Amazon S3 buckets.
Which solution will meet these requirements?

A. Use Amazon Data Firehose to create the data ingestion pipeline
B. Use Amazon SageMaker Studio Classic to create the model deployment pipelines.
C. Use AWS Glue to create the data ingestion pipeline
D. Use Amazon SageMaker Studio Classic to create the model deployment pipelines.
E. Use Amazon Redshift ML to create the data ingestion pipeline
F. Use Amazon SageMaker Studio Classic to create the model deployment pipelines.
G. Use Amazon Athena to create the data ingestion pipeline
H. Use an Amazon SageMaker notebook to create the model deployment pipelines.

Answer: B

NEW QUESTION 10
A company uses Amazon SageMaker for its ML workloads. The company's ML engineer receives a 50 MB Apache Parquet data file to build a fraud detection model. The file includes several correlated columns that are not required.
What should the ML engineer do to drop the unnecessary columns in the file with the LEAST effort?

A. Download the file to a local workstatio
B. Perform one-hot encoding by using a custom Python script.
C. Create an Apache Spark job that uses a custom processing script on Amazon EMR.
D. Create a SageMaker processing job by calling the SageMaker Python SDK.
E. Create a data flow in SageMaker Data Wrangle
F. Configure a transform step.

Answer: D

NEW QUESTION 11
An ML engineer is evaluating several ML models and must choose one model to use in production. The cost of false negative predictions by the models is much higher than the cost of false positive predictions.
Which metric finding should the ML engineer prioritize the MOST when choosing the model?

A. Low precision
B. High precision
C. Low recall
D. High recall

Answer: D

NEW QUESTION 13
HOTSPOT
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

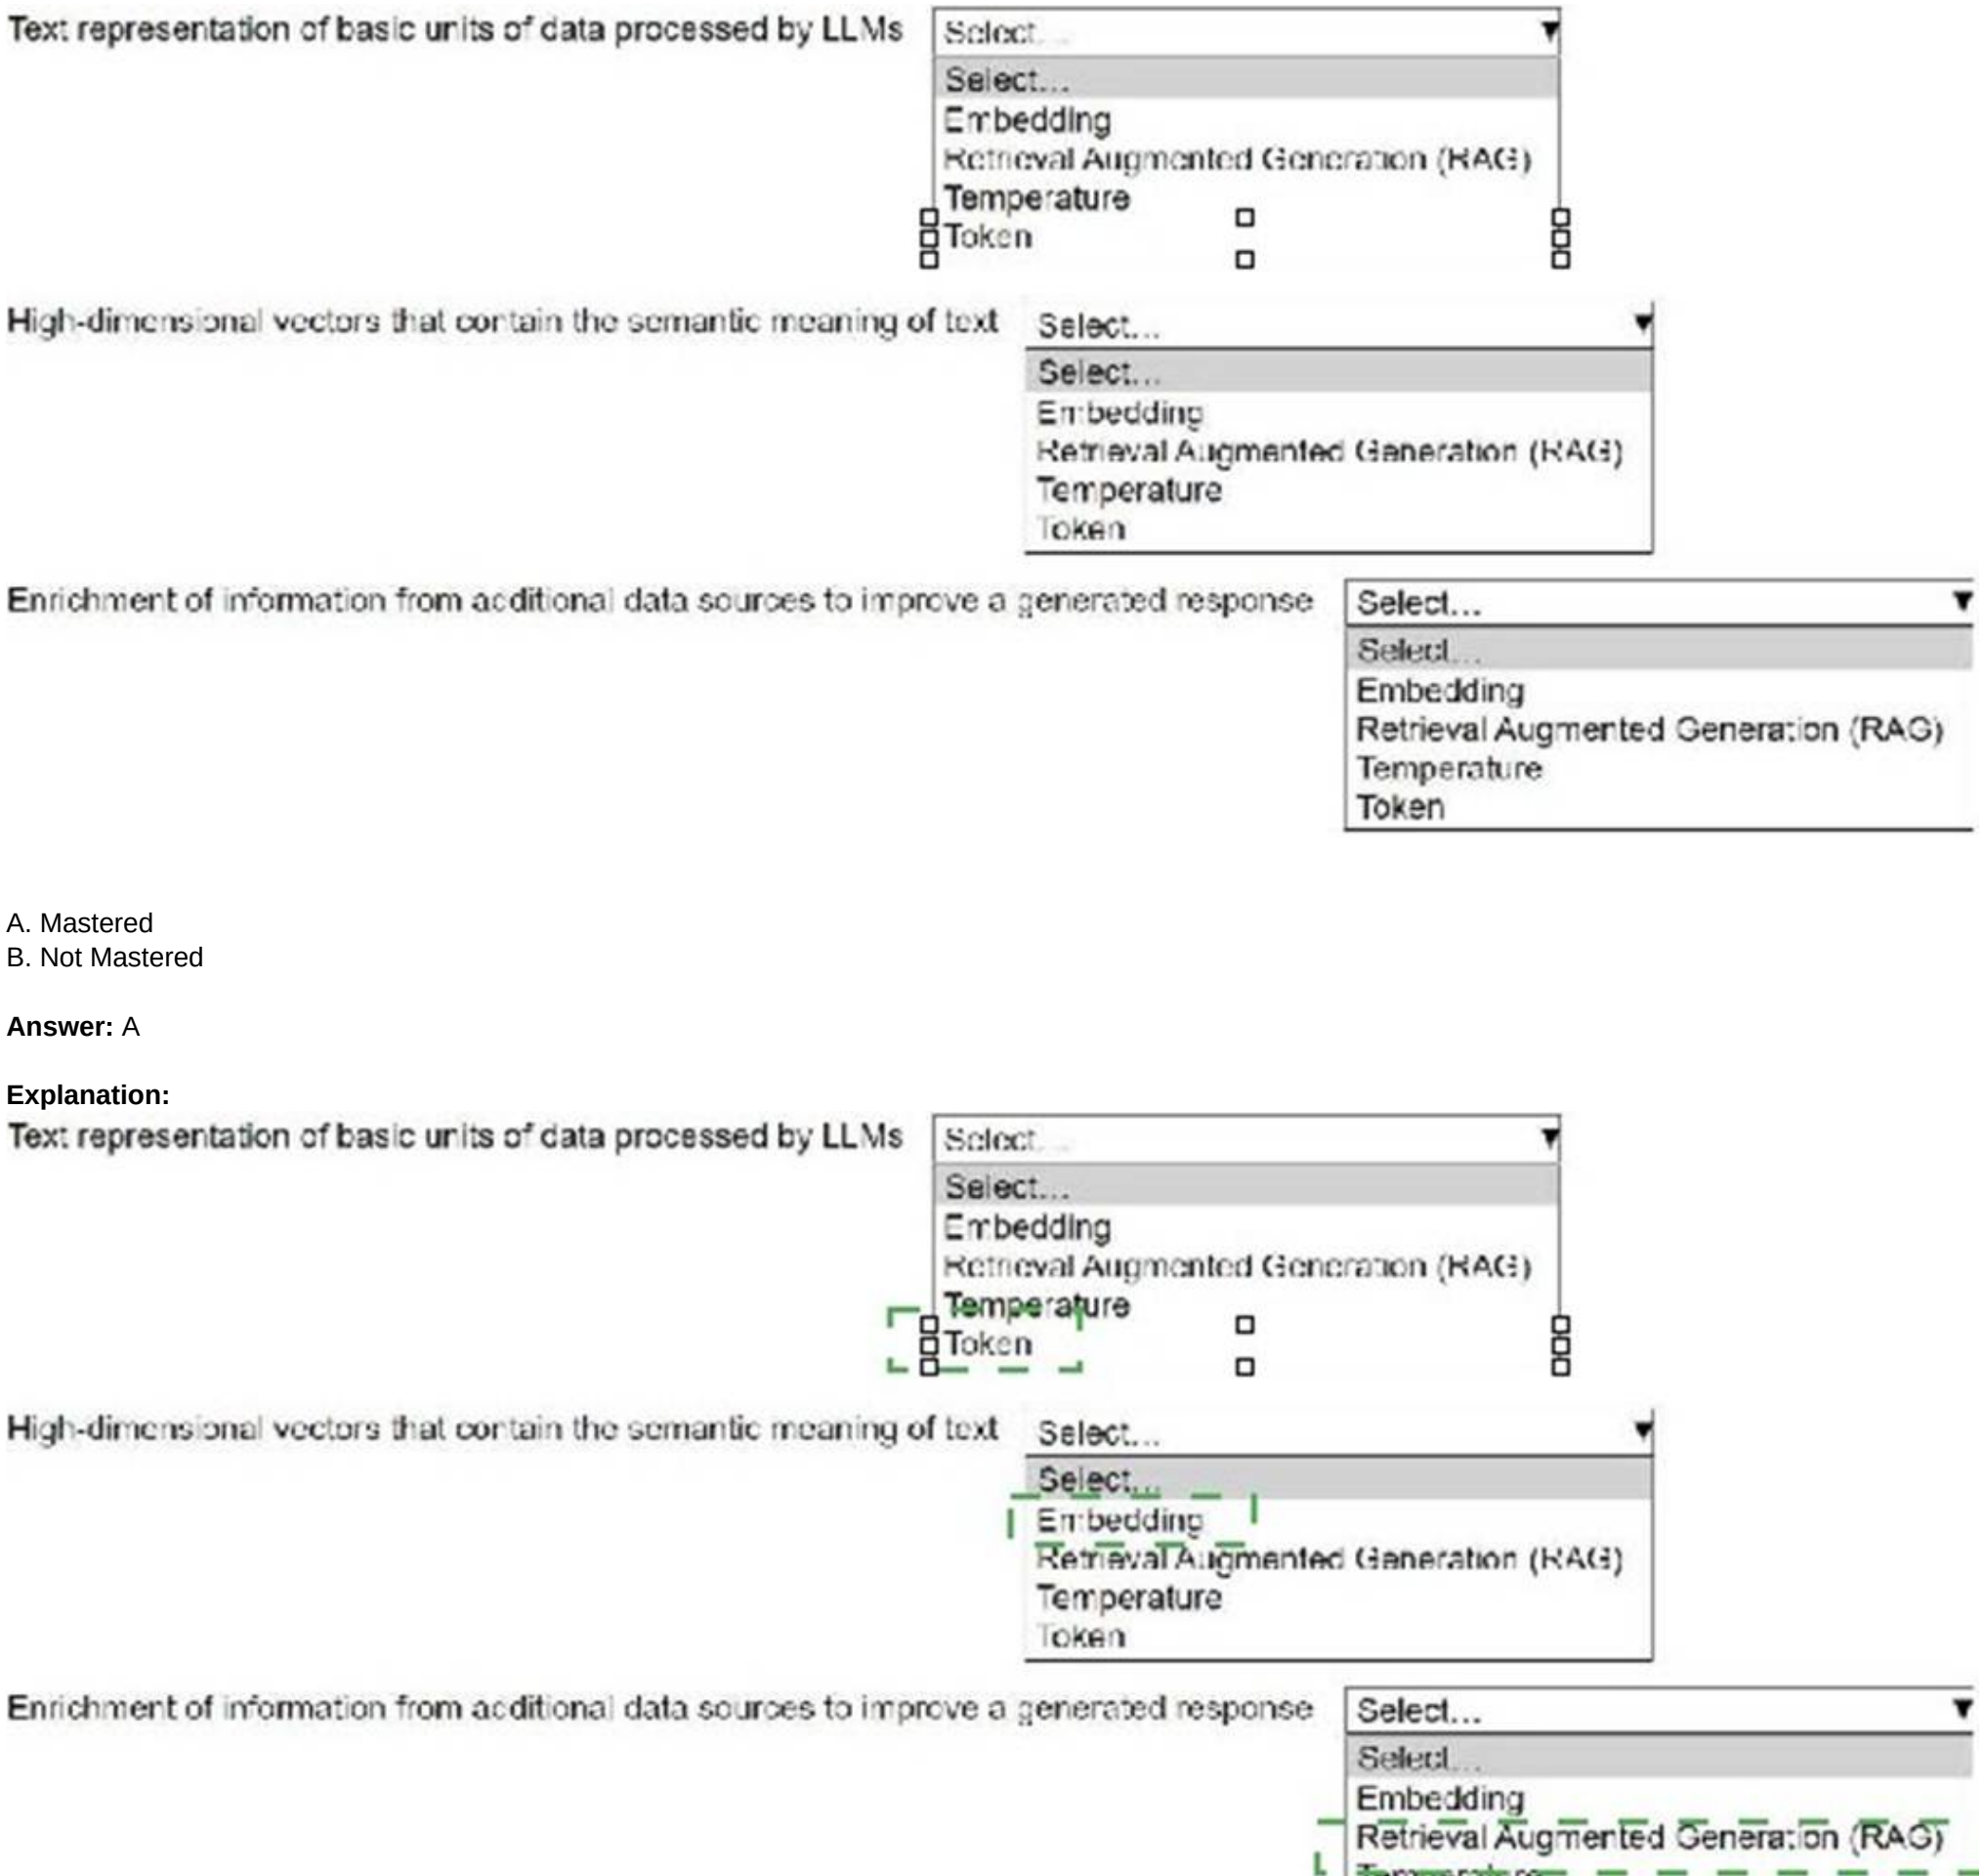

A. Mastered
B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 14
An ML engineer normalized training data by using min-max normalization in AWS Glue DataBrew. The ML engineer must normalize the production inference data in the same way as the training data before passing the production inference data to the model for predictions.
Which solution will meet this requirement?

A. Apply statistics from a well-known dataset to normalize the production samples.
B. Keep the min-max normalization statistics from the training se
C. Use these values to normalize the production samples.
D. Calculate a new set of min-max normalization statistics from a batch of production sample
E. Use these values to normalize all the production samples.
F. Calculate a new set of min-max normalization statistics from each production sampl
G. Use these values to normalize all the production samples.

Answer: B

NEW QUESTION 19
A company regularly receives new training data from the vendor of an ML model. The vendor delivers cleaned and prepared data to the company's Amazon S3 bucket every 3-4 days.
The company has an Amazon SageMaker pipeline to retrain the model. An ML engineer needs to implement a solution to run the pipeline when new data is uploaded to the S3 bucket.
Which solution will meet these requirements with the LEAST operational effort?

A. Create an S3 Lifecycle rule to transfer the data to the SageMaker training instance and to initiate training.
B. Create an AWS Lambda function that scans the S3 bucke
C. Program the Lambda function to initiate the pipeline when new data is uploaded.
D. Create an Amazon EventBridge rule that has an event pattern that matches the S3 uploa
E. Configure the pipeline as the target of the rule.
F. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the pipeline when new data is uploaded.A company wants to develop an ML model by using tabular data from its customer
G. The data contains meaningful ordered features with sensitive information that should not be discarde
H. An ML engineer must ensure that the sensitive data is masked before another team starts to build the model.Which solution will meet these requirements?
I. Use Amazon Made to categorize the sensitive data.
J. Prepare the data by using AWS Glue DataBrew.
K. Run an AWS Batch job to change the sensitive data to random values.
L. Run an Amazon EMR job to change the sensitive data to random values.

Answer: B

NEW QUESTION 20
Case study
An ML engineer is developing a fraud detection model on AWS. The training dataset includes transaction logs, customer profiles, and tables from an on-premises MySQL database. The transaction logs and customer profiles are stored in Amazon S3.
The dataset has a class imbalance that affects the learning of the model's algorithm. Additionally, many of the features have interdependencies. The algorithm is not capturing all the desired underlying patterns in the data.
After the data is aggregated, the ML engineer must implement a solution to automatically detect anomalies in the data and to visualize the result.
Which solution will meet these requirements?

A. Use Amazon Athena to automatically detect the anomalies and to visualize the result.
B. Use Amazon Redshift Spectrum to automatically detect the anomalie
C. Use Amazon QuickSight to visualize the result.
D. Use Amazon SageMaker Data Wrangler to automatically detect the anomalies and to visualize the result.
E. Use AWS Batch to automatically detect the anomalie
F. Use Amazon QuickSight to visualize the result.

Answer: C

NEW QUESTION 21
Case study
An ML engineer is developing a fraud detection model on AWS. The training dataset includes transaction logs, customer profiles, and tables from an on-premises MySQL database. The transaction logs and customer profiles are stored in Amazon S3.
The dataset has a class imbalance that affects the learning of the model's algorithm. Additionally, many of the features have interdependencies. The algorithm is not capturing all the desired underlying patterns in the data.
The training dataset includes categorical data and numerical data. The ML engineer must prepare the training dataset to maximize the accuracy of the model.
Which action will meet this requirement with the LEAST operational overhead?

A. Use AWS Glue to transform the categorical data into numerical data.
B. Use AWS Glue to transform the numerical data into categorical data.
C. Use Amazon SageMaker Data Wrangler to transform the categorical data into numerical data.
D. Use Amazon SageMaker Data Wrangler to transform the numerical data into categorical data.

Answer: C

NEW QUESTION 26
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

MLA-C01 Practice Exam Features:

* MLA-C01 Questions and Answers Updated Frequently

* MLA-C01 Practice Questions Verified by Expert Senior Certified Staff

* MLA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* MLA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year