# Amazon-Web-Services

## Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate

**NEW QUESTION 1**
A company is using Amazon SageMaker and millions of files to train an ML model. Each file is several megabytes in size. The files are stored in an Amazon S3 bucket. The company needs to improve training performance.
Which solution will meet these requirements in the LEAST amount of time?

A. Transfer the data to a new S3 bucket that provides S3 Express One Zone storag
B. Adjust the training job to use the new S3 bucket.
C. Create an Amazon FSx for Lustre file syste
D. Link the file system to the existing S3 bucke
E. Adjust the training job to read from the file system.
F. Create an Amazon Elastic File System (Amazon EFS) file syste
G. Transfer the existing data to the file syste
H. Adjust the training job to read from the file system.
I. Create an Amazon ElastiCache (Redis OSS) cluste
J. Link the Redis OSS cluster to the existing S3 bucke
K. Stream the data from the Redis OSS cluster directly to the training job.

**Answer:** B

**NEW QUESTION 2**
An ML engineer needs to implement a solution to host a trained ML model. The rate of requests to the model will be inconsistent throughout the day.
The ML engineer needs a scalable solution that minimizes costs when the model is not in use. The solution also must maintain the model's capacity to respond to requests during times of peak usage.
Which solution will meet these requirements?

A. Create AWS Lambda functions that have fixed concurrency to host the mode
B. Configure the Lambda functions to automatically scale based on the number of requests to the model.
C. Deploy the model on an Amazon Elastic Container Service (Amazon ECS) cluster that uses AWS Fargat
D. Set a static number of tasks to handle requests during times of peak usage.
E. Deploy the model to an Amazon SageMaker endpoin
F. Deploy multiple copies of the model to the endpoin
G. Create an Application Load Balancer to route traffic between the different copies of the model at the endpoint.
H. Deploy the model to an Amazon SageMaker endpoin
I. Create SageMaker endpoint auto scaling policies that are based on Amazon CloudWatch metrics to adjust the number of instances dynamically.

**Answer:** D

**NEW QUESTION 3**
A company has historical data that shows whether customers needed long-term support from company staff. The company needs to develop an ML model to predict whether new customers will require long-term support.
Which modeling approach should the company use to meet this requirement?

A. Anomaly detection
B. Linear regression
C. Logistic regression
D. Semantic segmentation

**Answer:** C

**NEW QUESTION 4**
An advertising company uses AWS Lake Formation to manage a data lake. The data lake contains structured data and unstructured data. The company's ML engineers are assigned to specific advertisement campaigns.
The ML engineers must interact with the data through Amazon Athena and by browsing the data directly in an Amazon S3 bucket. The ML engineers must have access to only the resources that are specific to their assigned advertisement campaigns.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Configure IAM policies on an AWS Glue Data Catalog to restrict access to Athena based on the ML engineers' campaigns.
B. Store users and campaign information in an Amazon DynamoDB tabl
C. Configure DynamoDB Streams to invoke an AWS Lambda function to update S3 bucket policies.
D. Use Lake Formation to authorize AWS Glue to access the S3 bucke
E. Configure Lake Formation tags to map ML engineers to their campaigns.
F. Configure S3 bucket policies to restrict access to the S3 bucket based on the ML engineers' campaigns.

**Answer:** C

**NEW QUESTION 5**
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 6**
A company has an ML model that needs to run one time each night to predict stock values. The model input is 3 MB of data that is collected during the current day. The model produces the predictions for the next day. The prediction process takes less than 1 minute to finish running.
How should the company deploy the model on Amazon SageMaker to meet these requirements?

A. Use a multi-model serverless endpoin
B. Enable caching.
C. Use an asynchronous inference endpoin
D. Set the InitialInstanceCount parameter to 0.
E. Use a real-time endpoin
F. Configure an auto scaling policy to scale the model to 0 when the model is not in use.
G. Use a serverless inference endpoin
H. Set the MaxConcurrency parameter to 1.

**Answer:** D

**NEW QUESTION 7**
A company has deployed an XGBoost prediction model in production to predict if a customer is likely to cancel a subscription. The company uses Amazon SageMaker Model Monitor to detect deviations in the F1 score.
During a baseline analysis of model quality, the company recorded a threshold for the F1 score. After several months of no change, the model's F1 score decreases significantly.
What could be the reason for the reduced F1 score?

A. Concept drift occurred in the underlying customer data that was used for predictions.
B. The model was not sufficiently complex to capture all the patterns in the original baseline data.
C. The original baseline data had a data quality issue of missing values.
D. Incorrect ground truth labels were provided to Model Monitor during the calculation of the baseline.

**Answer:** A

**NEW QUESTION 8**
A company runs an Amazon SageMaker domain in a public subnet of a newly created VPC. The network is configured properly, and ML engineers can access the SageMaker domain.
Recently, the company discovered suspicious traffic to the domain from a specific IP address. The company needs to block traffic from the specific IP address.
Which update to the network configuration will meet this requirement?

A. Create a security group inbound rule to deny traffic from the specific IP addres
B. Assign the security group to the domain.
C. Create a network ACL inbound rule to deny traffic from the specific IP addres
D. Assign the rule to the default network Ad for the subnet where the domain is located.
E. Create a shadow variant for the domai
F. Configure SageMaker Inference Recommender to send traffic from the specific IP address to the shadow endpoint.
G. Create a VPC route table to deny inbound traffic from the specific IP addres
H. Assign the route table to the domain.

**Answer:** B

**NEW QUESTION 9**
A company is creating an application that will recommend products for customers to purchase. The application will make API calls to Amazon Q Business. The company must ensure that responses from Amazon Q Business do not include the name of the company's main competitor.
Which solution will meet this requirement?

A. Configure the competitor's name as a blocked phrase in Amazon Q Business.
B. Configure an Amazon Q Business retriever to exclude the competitor??s name.
C. Configure an Amazon Kendra retriever for Amazon Q Business to build indexes that exclude the competitor's name.
D. Configure document attribute boosting in Amazon Q Business to deprioritize the competitor's name.

**Answer:** A

**NEW QUESTION 10**
HOTSPOT
An ML engineer is working on an ML model to predict the prices of similarly sized homes. The model will base predictions on several features The ML engineer will use the following feature engineering techniques to estimate the prices of the homes:
• Feature splitting
• Logarithmic transformation
• One-hot encoding
• Standardized distribution
Select the correct feature engineering techniques for the following list of features. Each feature engineering technique should be selected one time or not at all (Select three.)

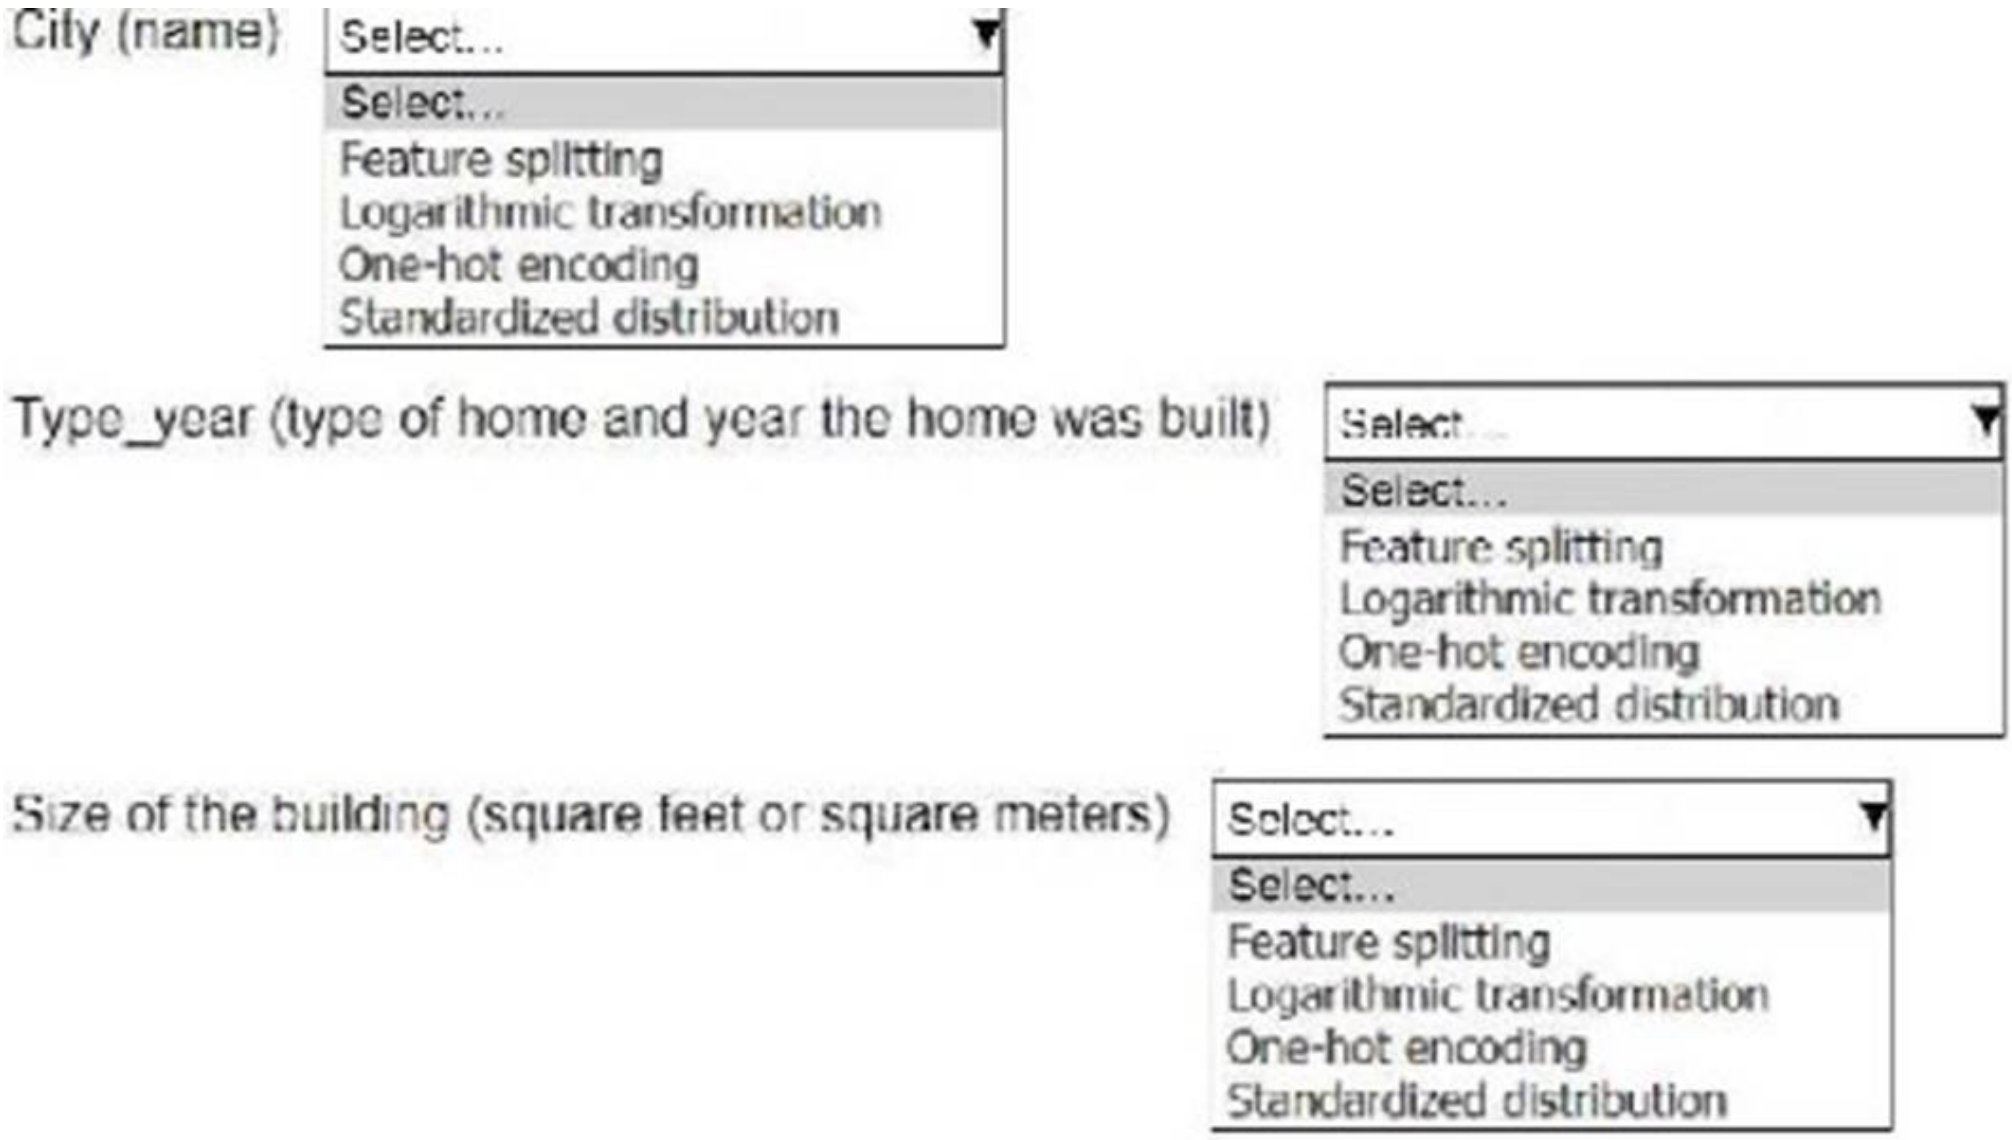


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

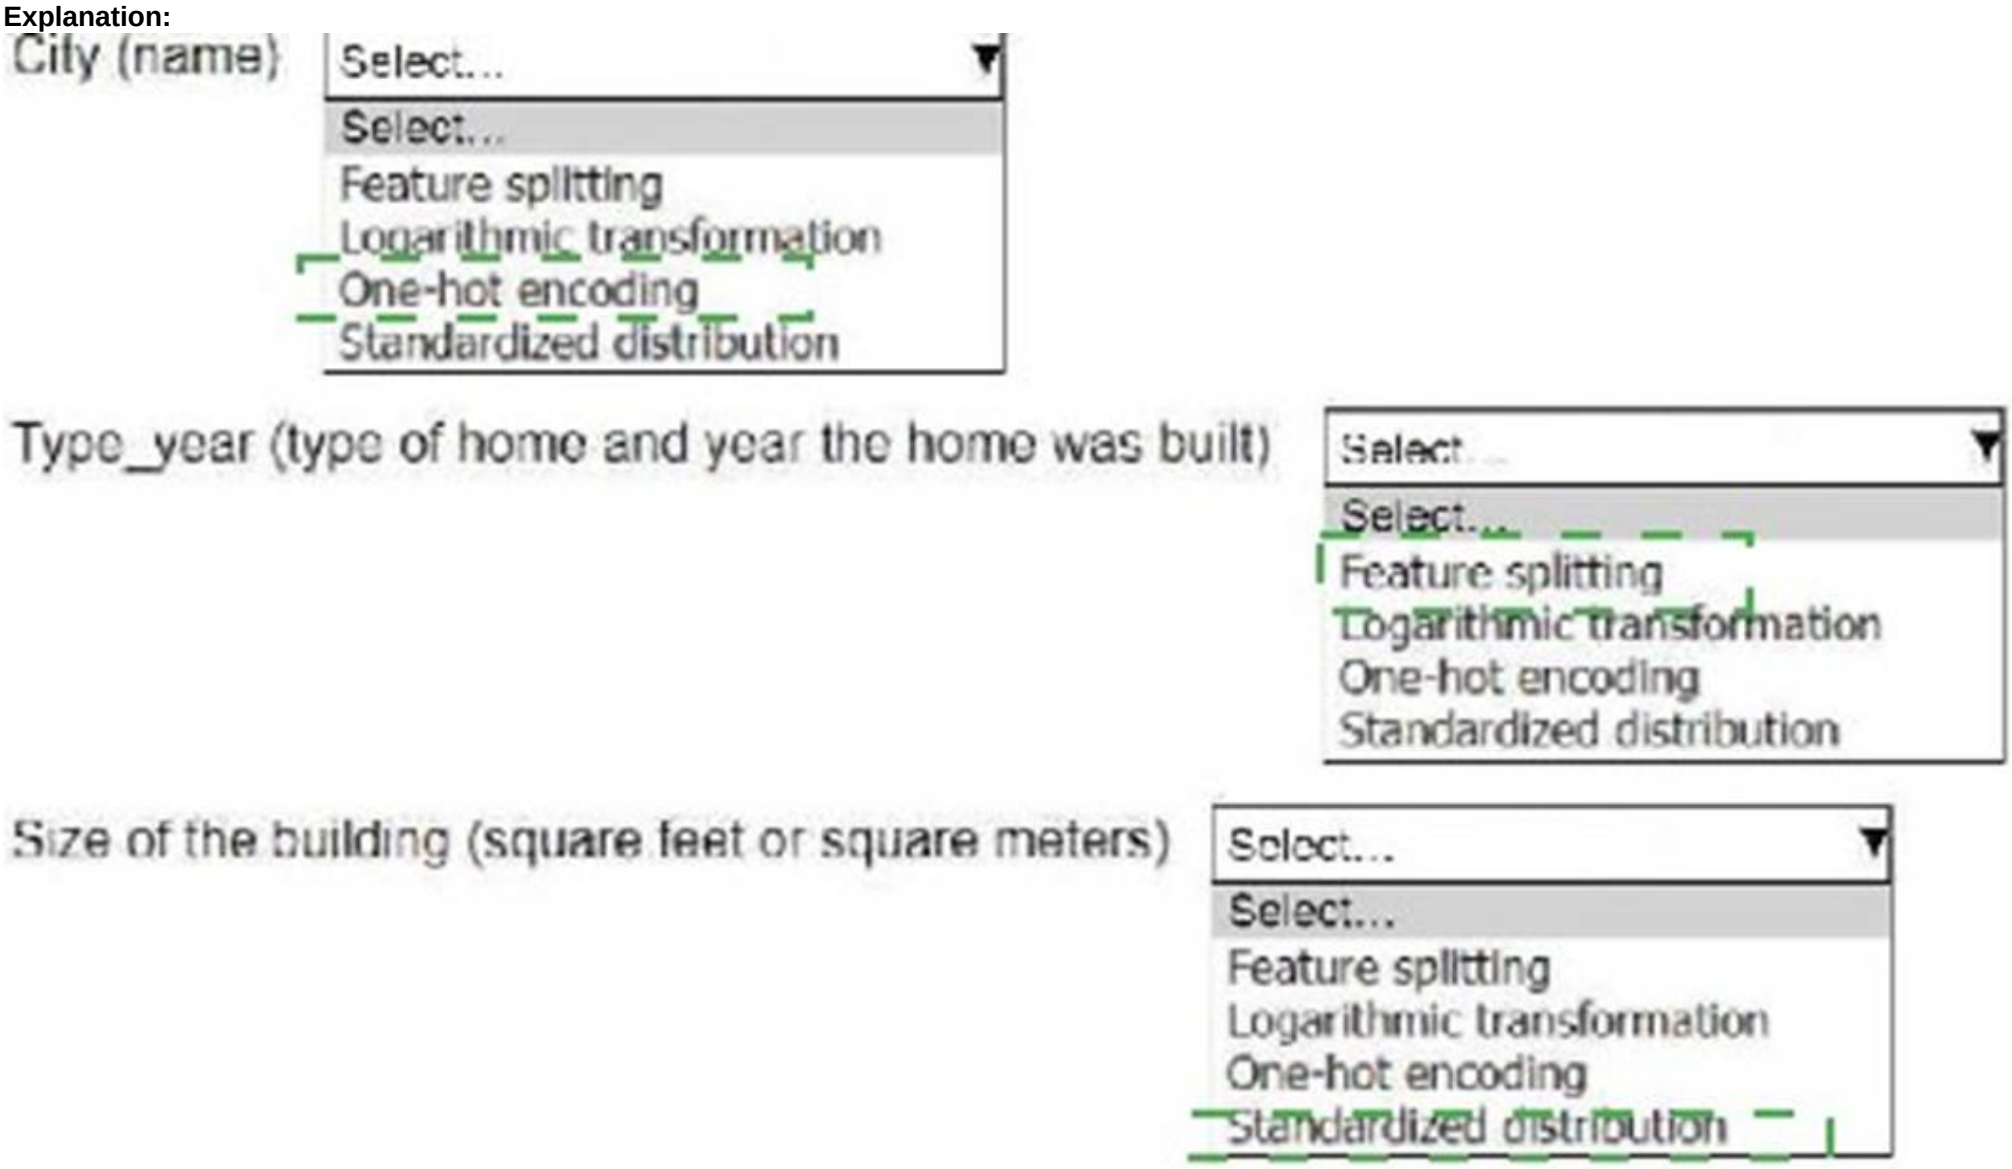


**NEW QUESTION 10**
A company is using Amazon SageMaker to create ML models. The company's data scientists need fine-grained control of the ML workflows that they orchestrate. The data scientists also need the ability to visualize SageMaker jobs and workflows as a directed acyclic graph (DAG). The data scientists must keep a running history of model discovery experiments and must establish model governance for auditing and compliance verifications.
Which solution will meet these requirements?

A. Use AWS CodePipeline and its integration with SageMaker Studio to manage the entire ML workflow
B. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.
C. Use AWS CodePipeline and its integration with SageMaker Experiments to manage the entire ML workflow
D. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.
E. Use SageMaker Pipelines and its integration with SageMaker Studio to manage the entire ML workflow
F. Use SageMaker ML Lineage Tracking for the running history of experiments and for auditing and compliance verifications.
G. Use SageMaker Pipelines and its integration with SageMaker Experiments to manage the entire ML workflow
H. Use SageMaker Experiments for the running history of experiments and for auditing and compliance verifications.

**Answer:** C

**NEW QUESTION 13**
A company has deployed an ML model that detects fraudulent credit card transactions in real time in a banking application. The model uses Amazon SageMaker Asynchronous Inference. Consumers are reporting delays in receiving the inference results.
An ML engineer needs to implement a solution to improve the inference performance. The solution also must provide a notification when a deviation in model quality occurs.
Which solution will meet these requirements?

A. Use SageMaker real-time inference for inferenc
B. Use SageMaker Model Monitor for notifications about model quality.
C. Use SageMaker batch transform for inferenc
D. Use SageMaker Model Monitor for notifications about model quality.
E. Use SageMaker Serverless Inference for inferenc
F. Use SageMaker Inference Recommender for notifications about model quality.
G. Keep using SageMaker Asynchronous Inference for inferenc
H. Use SageMaker Inference Recommender for notifications about model quality.

**Answer:** A

**NEW QUESTION 16**
A company needs to run a batch data-processing job on Amazon EC2 instances. The job will run during the weekend and will take 90 minutes to finish running. The processing can
handle interruptions. The company will run the job every weekend for the next 6 months.
Which EC2 instance purchasing option will meet these requirements MOST cost- effectively?

A. Spot Instances
B. Reserved Instances
C. On-Demand Instances
D. Dedicated Instances

**Answer:** A

**NEW QUESTION 21**
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 25**
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

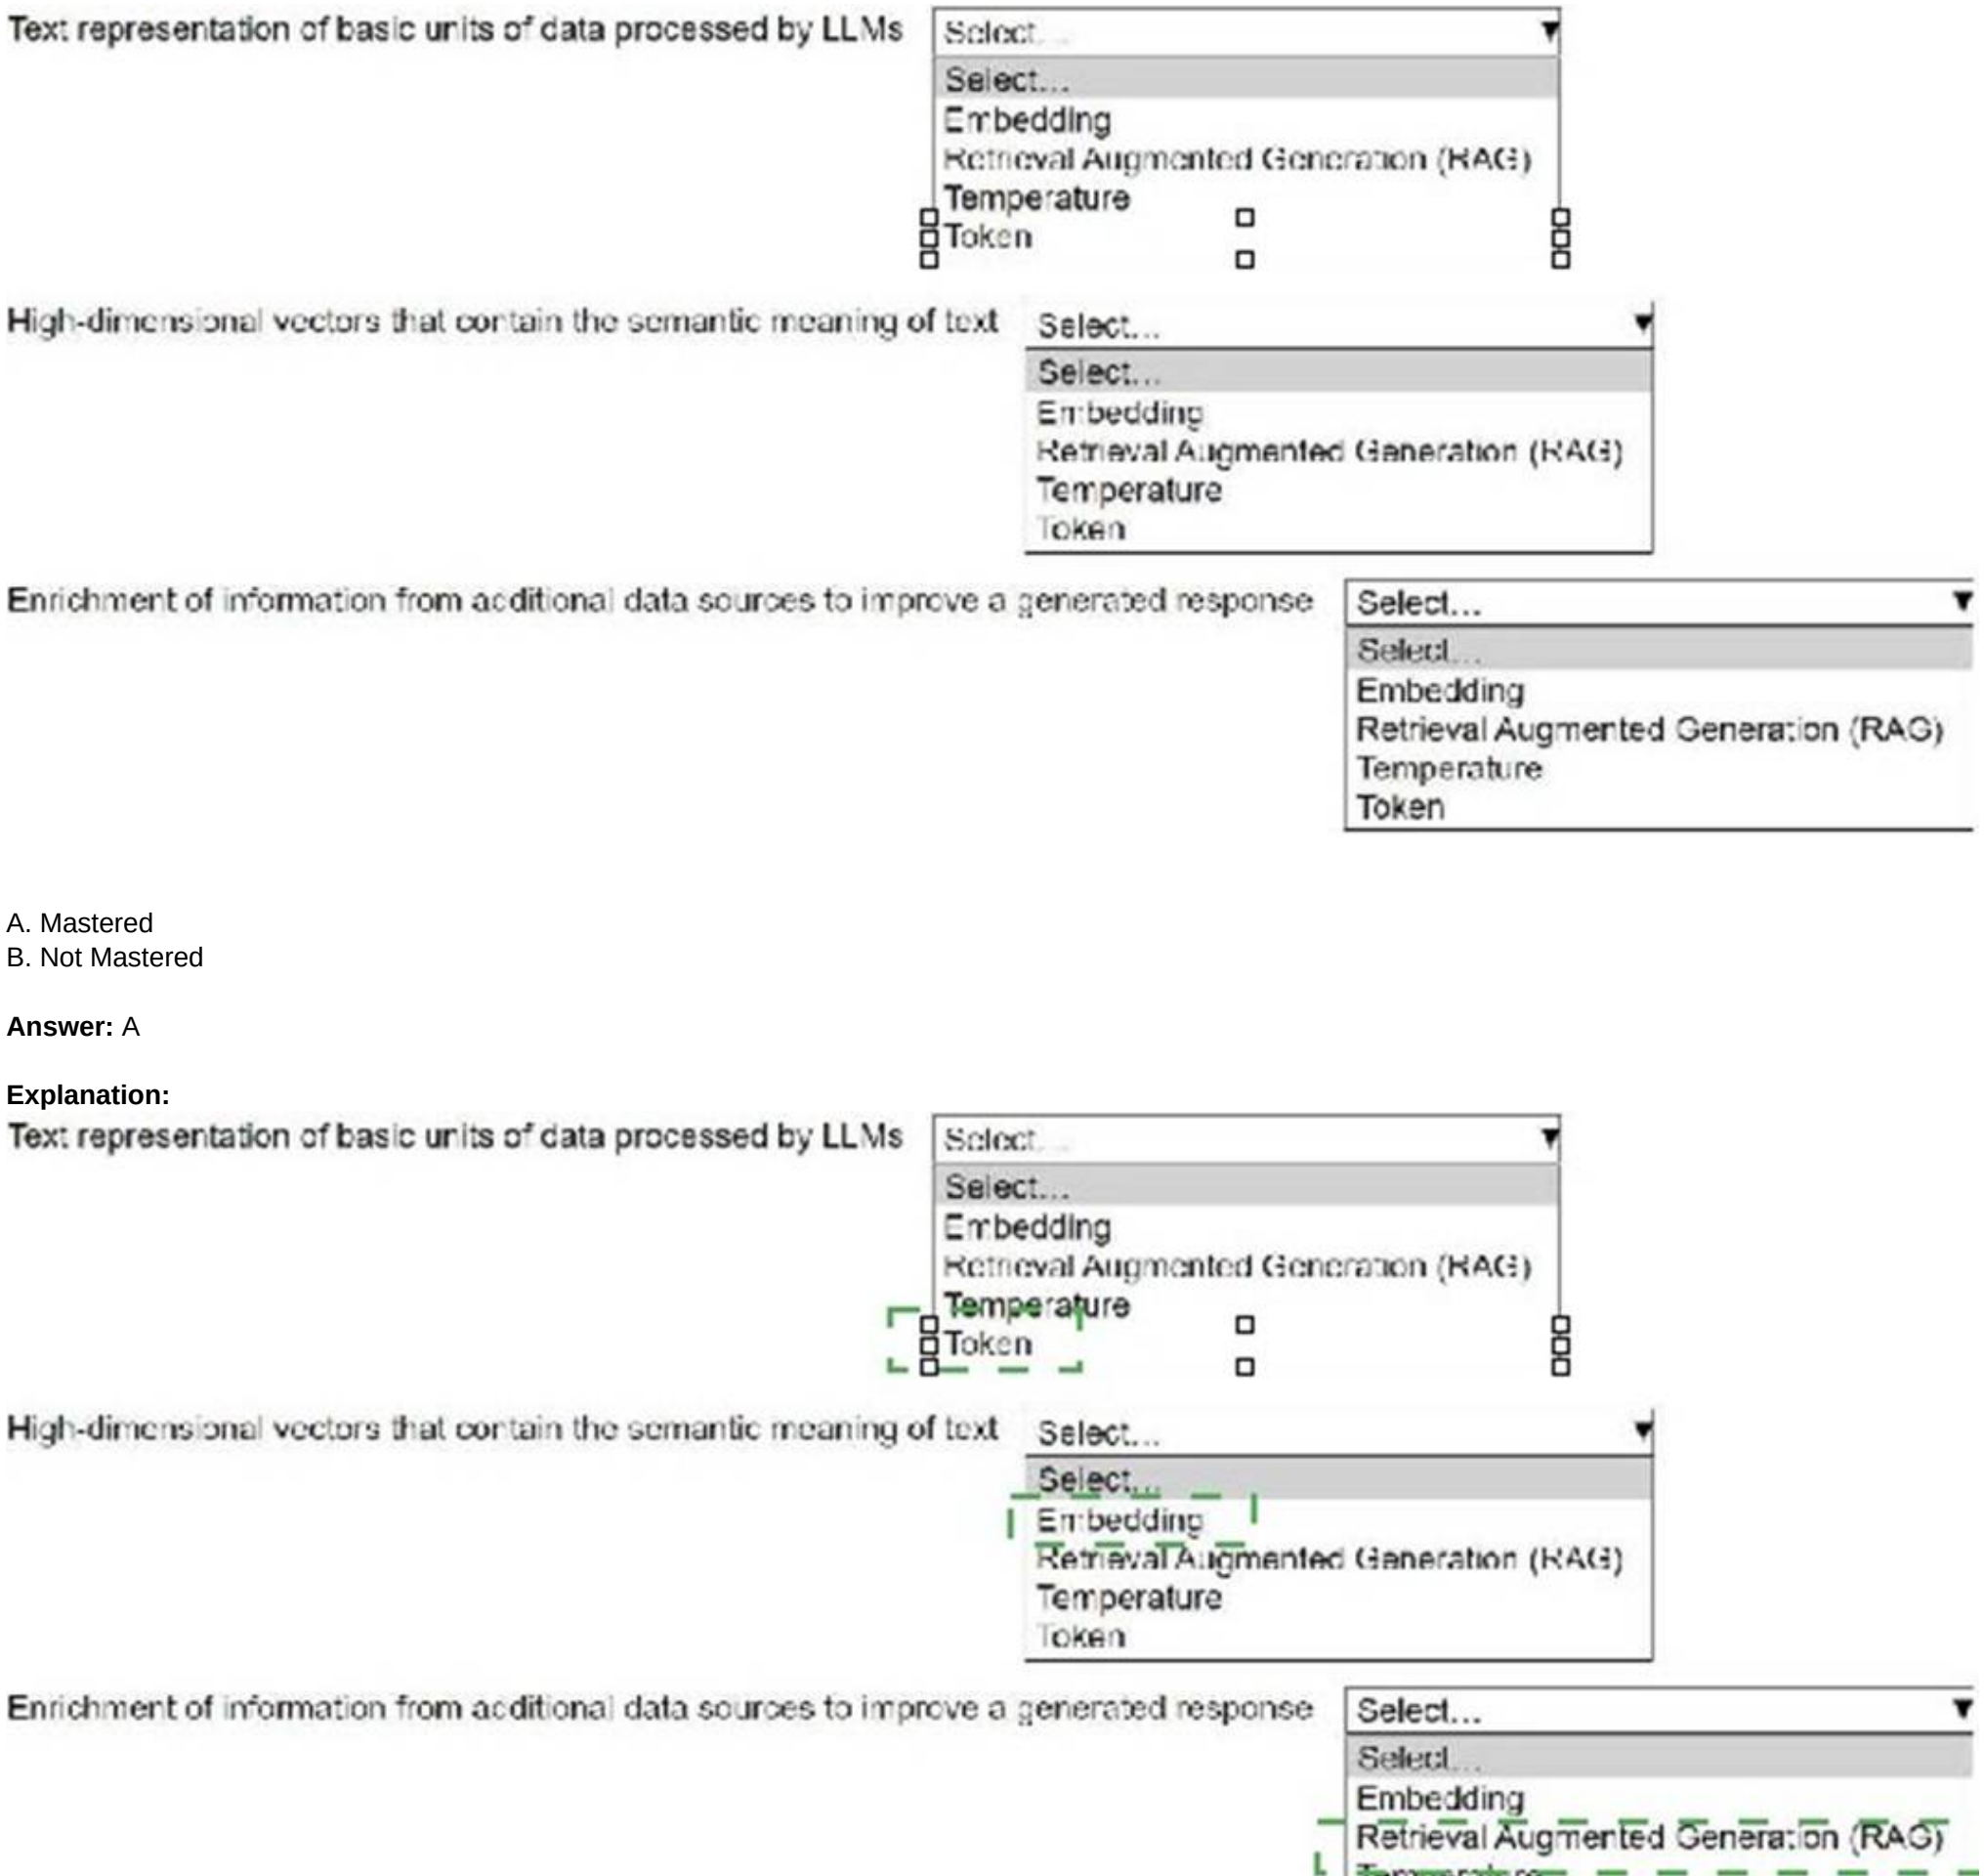


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

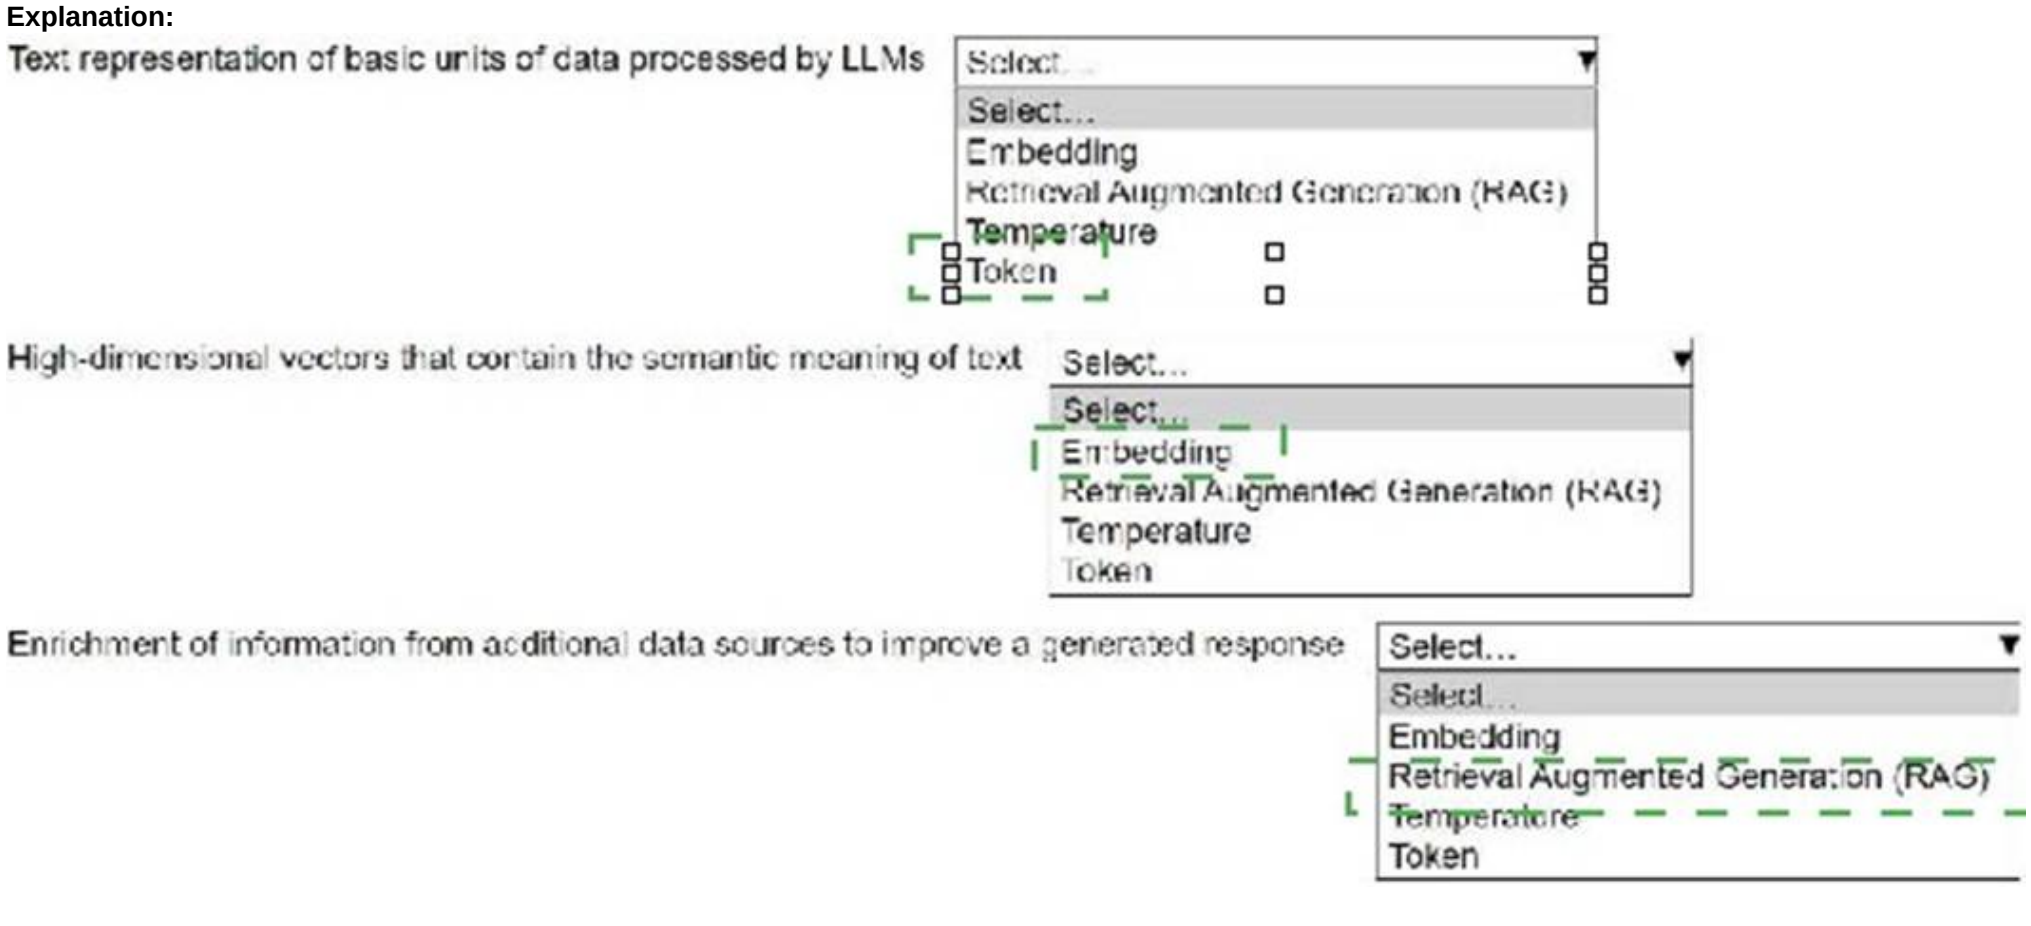


**NEW QUESTION 30**
A company needs to host a custom ML model to perform forecast analysis. The forecast analysis will occur with predictable and sustained load during the same 2-hour period every day.
Multiple invocations during the analysis period will require quick responses. The company needs AWS to manage the underlying infrastructure and any auto scaling activities.
Which solution will meet these requirements?

A. Schedule an Amazon SageMaker batch transform job by using AWS Lambda.
B. Configure an Auto Scaling group of Amazon EC2 instances to use scheduled scaling.
C. Use Amazon SageMaker Serverless Inference with provisioned concurrency.
D. Run the model on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster on Amazon EC2 with pod auto scaling.

**Answer:** C

**NEW QUESTION 34**
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 ML engineer needs to use an Amazon SageMaker built-in algorithm to train the model. Which algorithm should the ML engineer use to meet this requirement?

A. LightGBM
B. Linear learner
C. -means clustering
D. Neural Topic Model (NTM)

**Answer:** B

**NEW QUESTION 39**
An ML engineer needs to use an Amazon EMR cluster to process large volumes of data in batches. Any data loss is unacceptable.
Which instance purchasing option will meet these requirements MOST cost-effectively?

A. Run the primary node, core nodes, and task nodes on On-Demand Instances.
B. Run the primary node, core nodes, and task nodes on Spot Instances.
C. Run the primary node on an On-Demand Instanc
D. Run the core nodes and task nodes on Spot Instances.
E. Run the primary node and core nodes on On-Demand Instance
F. Run the task nodes on Spot Instances.

**Answer:** D

**NEW QUESTION 41**
A company that has hundreds of data scientists is using Amazon SageMaker to create ML models. The models are in model groups in the SageMaker Model Registry.
The data scientists are grouped into three categories: computer vision, natural language processing (NLP), and speech recognition. An ML engineer needs to implement a solution to organize the existing models into these groups to improve model discoverability at scale. The solution must not affect the integrity of the model artifacts and their existing groupings.
Which solution will meet these requirements?

A. Create a custom tag for each of the three categorie
B. Add the tags to the model packages in the SageMaker Model Registry.
C. Create a model group for each categor
D. Move the existing models into these category model groups.
E. Use SageMaker ML Lineage Tracking to automatically identify and tag which model groups should contain the models.
F. Create a Model Registry collection for each of the three categorie
G. Move the existing model groups into the collections.

**Answer:** A

**NEW QUESTION 43**
.....

# 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