MLA-C01 Dumps

AWS Certified Machine Learning Engineer - Associate

https://www.certleader.com/MLA-C01-dumps.html

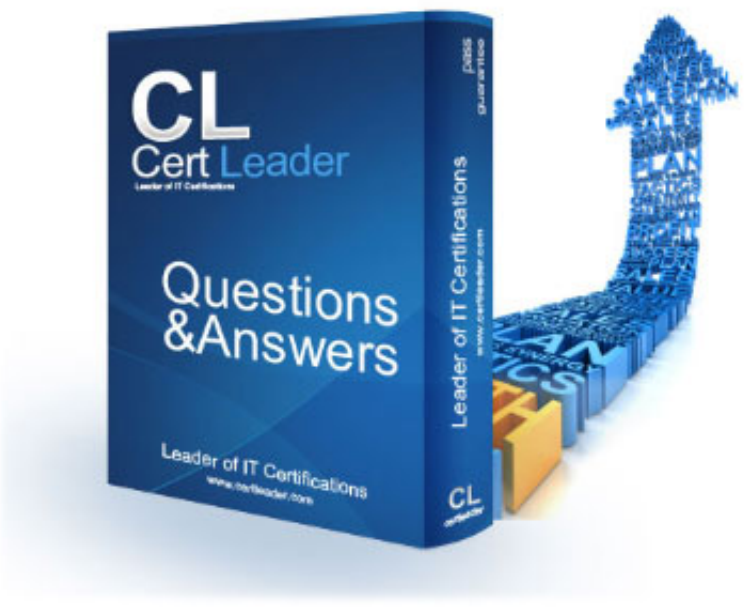

NEW QUESTION 1
An ML engineer needs to implement a solution to host a trained ML model. The rate of requests to the model will be inconsistent throughout the day.
The ML engineer needs a scalable solution that minimizes costs when the model is not in use. The solution also must maintain the model's capacity 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 2
An ML engineer needs to use data with Amazon SageMaker Canvas to train an ML model. The data is stored in Amazon S3 and is complex in structure. The ML engineer must use a file format that minimizes processing time for the data.
Which file format will meet these requirements?

A. CSV files compressed with Snappy
B. JSON objects in JSONL format
C. JSON files compressed with gzip
D. Apache Parquet files

Answer: D

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
HOTSPOT
A company stores historical data in .csv files in Amazon S3. Only some of the rows and columns in the .csv files are populated. The columns are not labeled. An ML engineer needs to prepare and store the data so that the company can use the data to train ML models.
Select and order the correct steps from the following list to perform this task. Each step should be selected one time or not at all. (Select and order three.)
• Create an Amazon SageMaker batch transform job for data cleaning and feature engineering.
• Store the resulting data back in Amazon S3.
• Use Amazon Athena to infer the schemas and available columns.
• Use AWS Glue crawlers to infer the schemas and available columns.
• Use AWS Glue DataBrew for data cleaning and feature engineering.

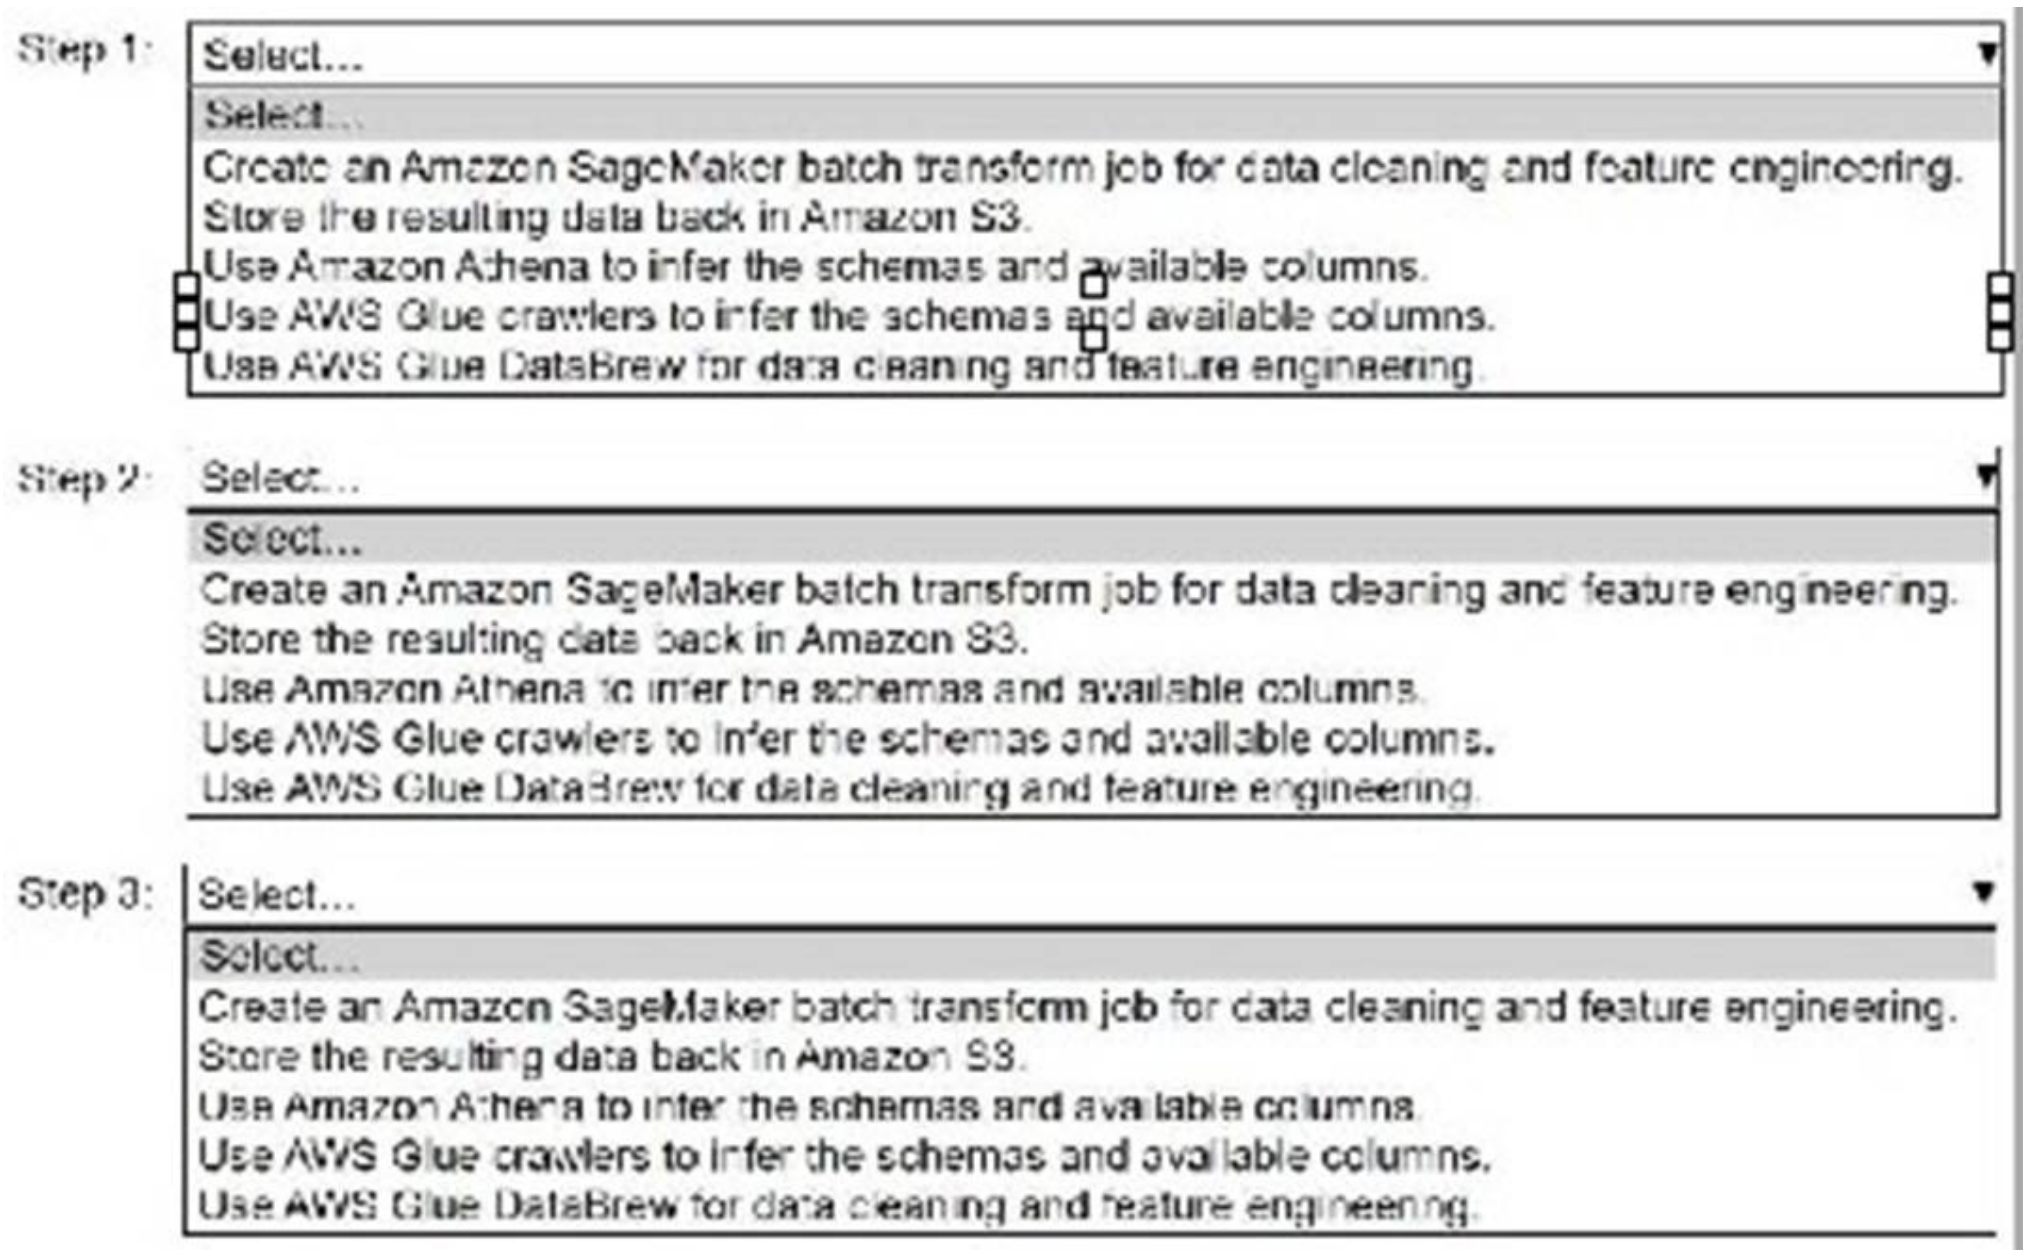

A. Mastered
B. Not Mastered

Answer: A

Explanation:

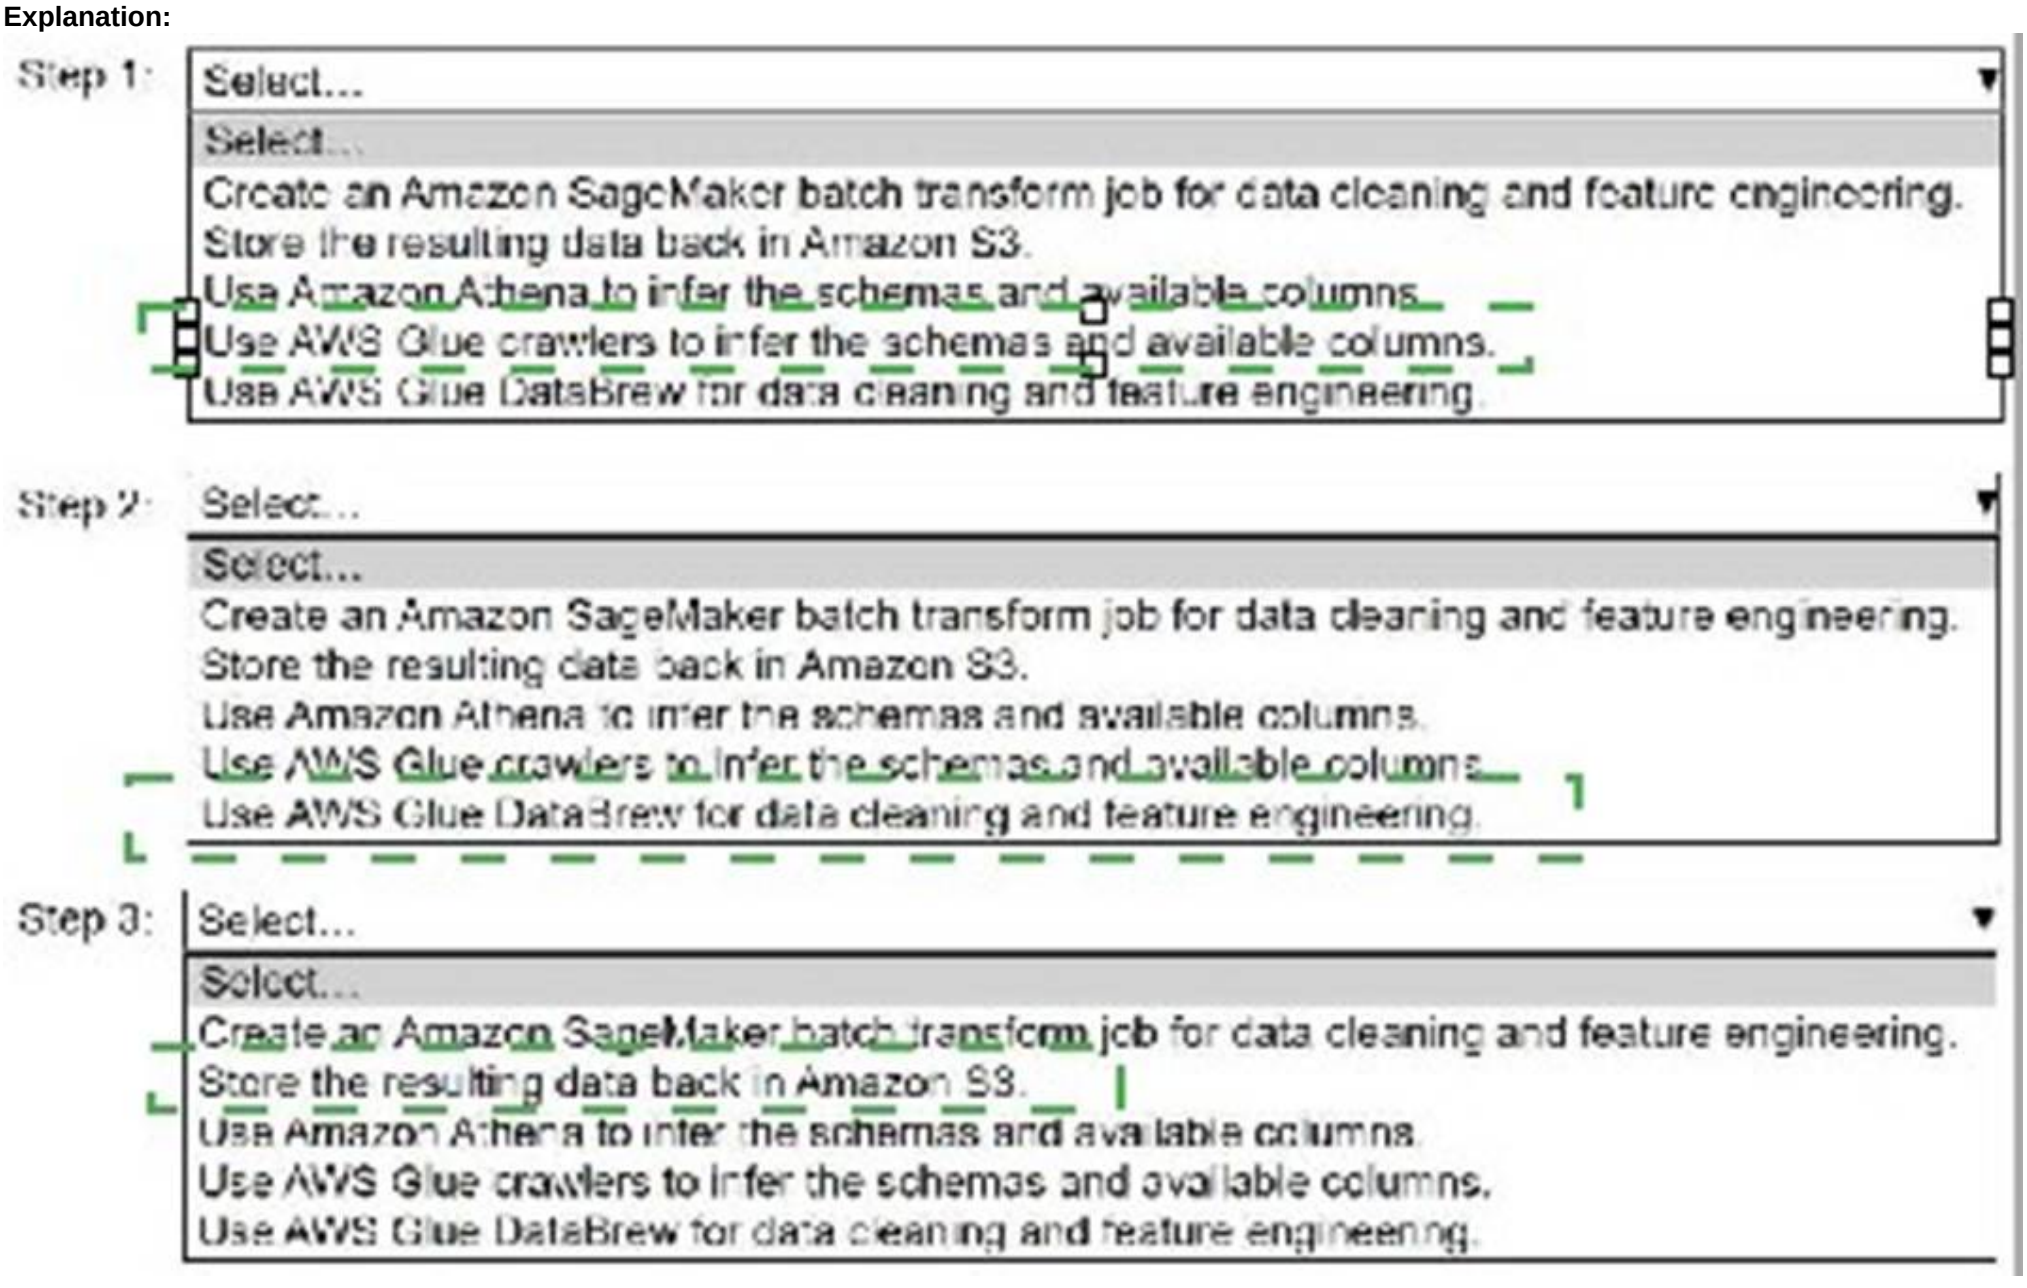

NEW QUESTION 5

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 6
An ML engineer is using a training job to fine-tune a deep learning model in Amazon SageMaker Studio. The ML engineer previously used the same pre-trained model with a similar dataset. The ML engineer expects vanishing gradient, underutilized GPU, and overfitting problems.
The ML engineer needs to implement a solution to detect these issues and to react in predefined ways when the issues occur. The solution also must provide comprehensive real-time metrics during the training.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use TensorBoard to monitor the training jo
B. Publish the findings to an Amazon Simple Notification Service (Amazon SNS) topi
C. Create an AWS Lambda function to consume the findings and to initiate the predefined actions.
D. Use Amazon CloudWatch default metrics to gain insights about the training jo
E. Use themetrics to invoke an AWS Lambda function to initiate the predefined actions.
F. Expand the metrics in Amazon CloudWatch to include the gradients in each training ste
G. Use the metrics to invoke an AWS Lambda function to initiate the predefined actions.
H. Use SageMaker Debugger built-in rules to monitor the training jo
I. Configure the rules to initiate the predefined actions.

Answer: D

NEW QUESTION 7
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.)

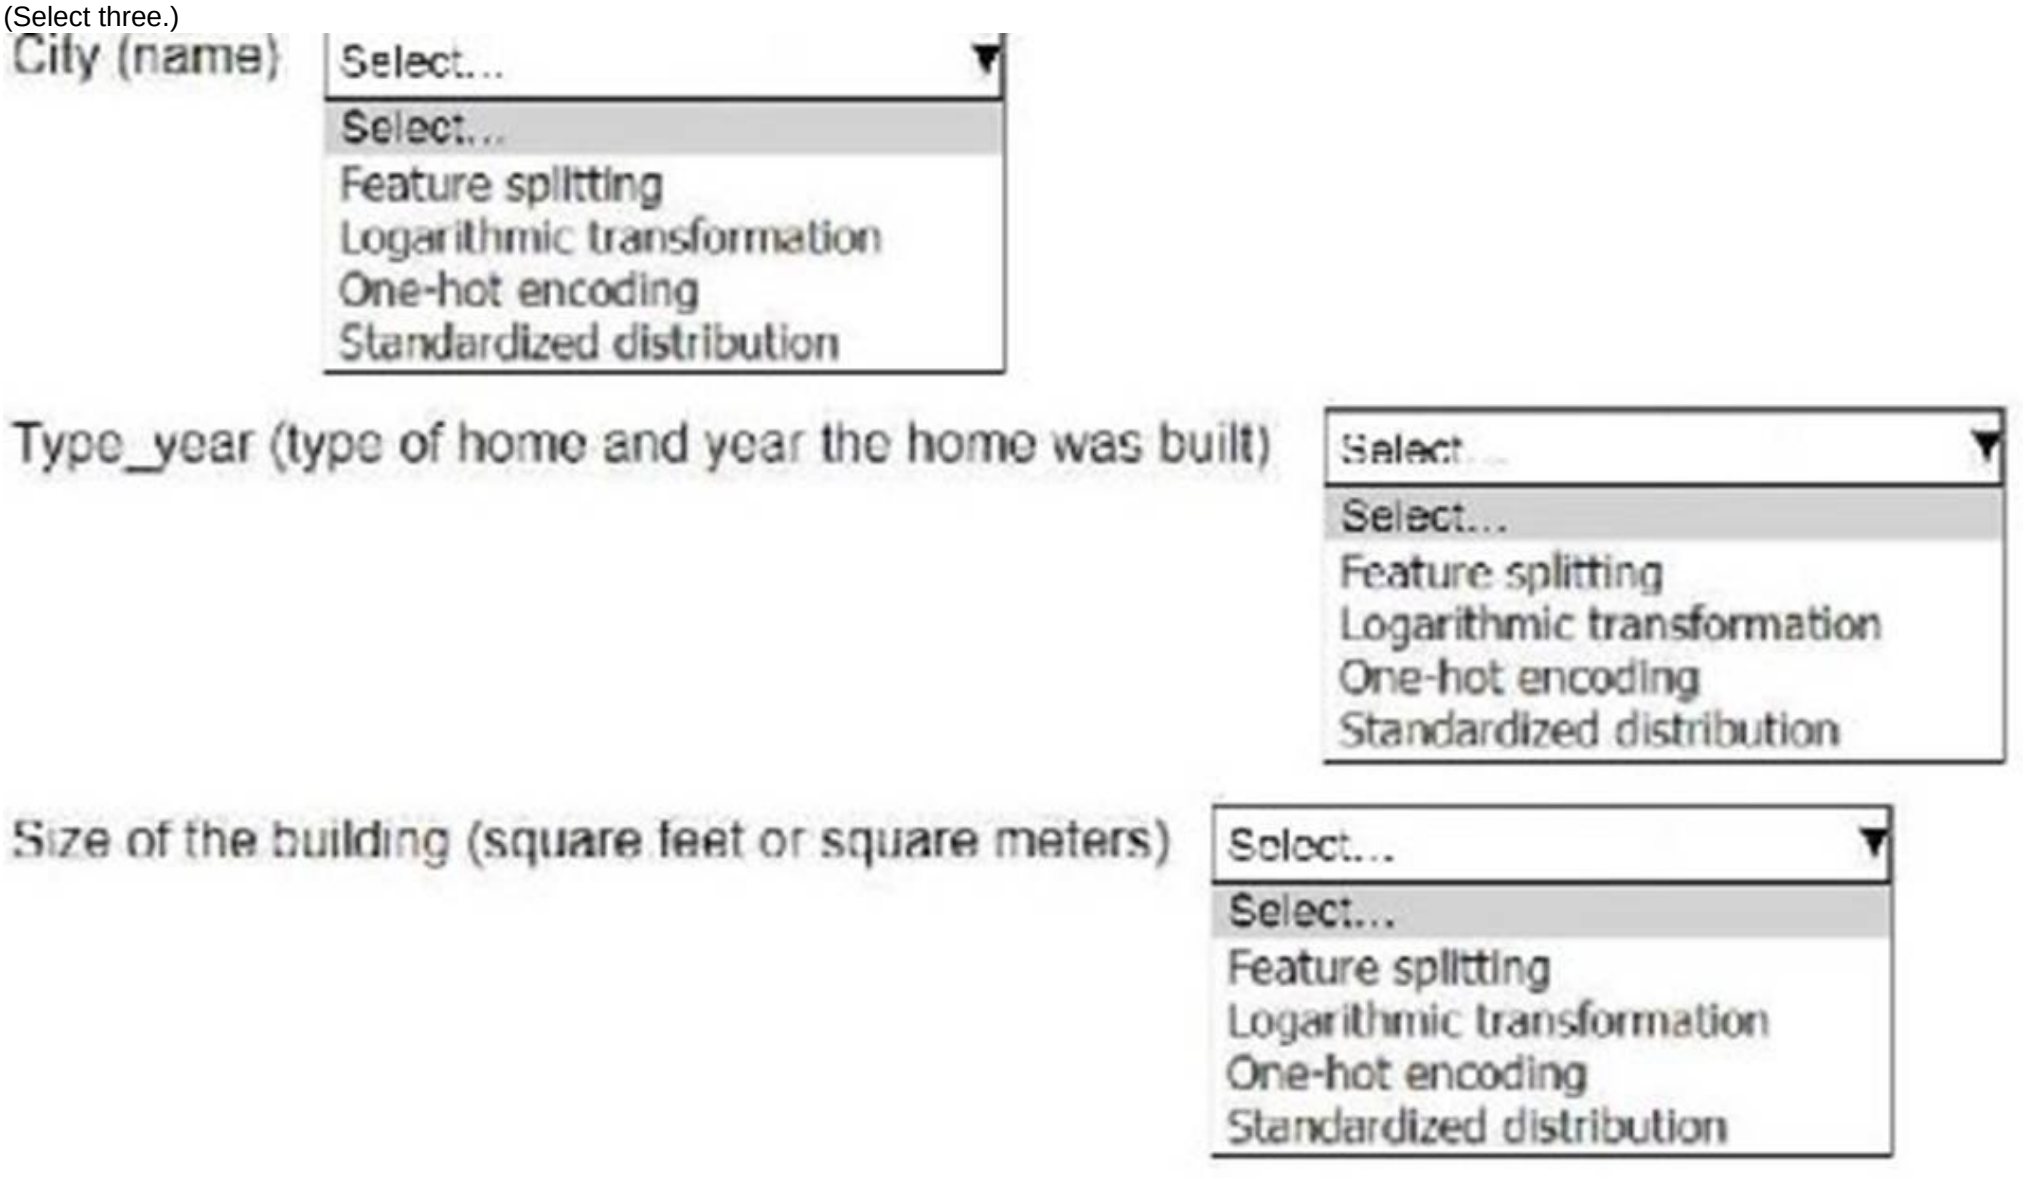

A. Mastered
B. Not Mastered

Answer: A

Explanation:

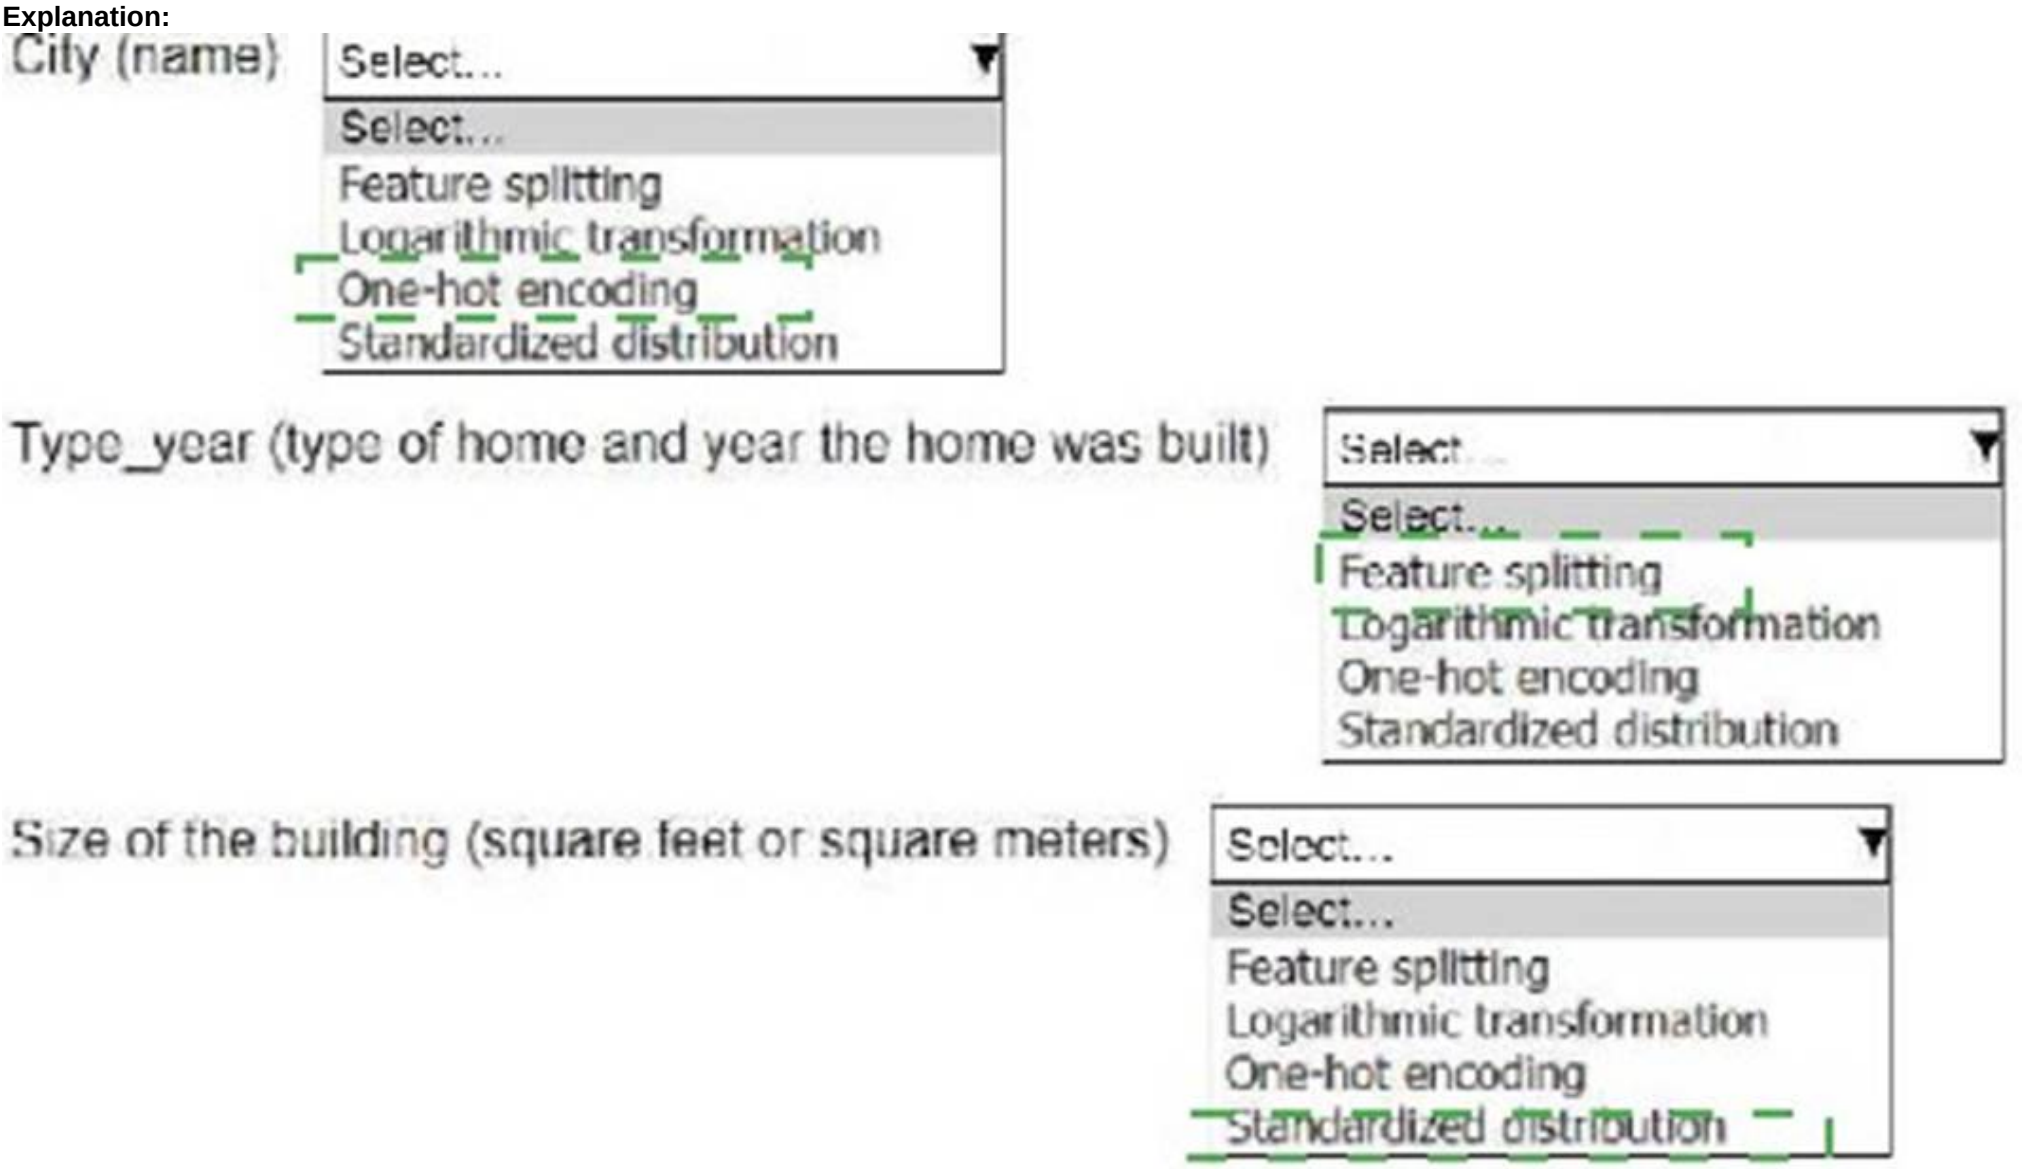

NEW QUESTION 8
A company uses Amazon SageMaker Studio to develop an ML model. The company has a single SageMaker Studio domain. An ML engineer needs to implement a solution that provides an automated alert when SageMaker compute costs reach a specific threshold.
Which solution will meet these requirements?

A. Add resource tagging by editing the SageMaker user profile in the SageMaker domai
B. Configure AWS Cost Explorer to send an alert when the threshold is reached.
C. Add resource tagging by editing the SageMaker user profile in the SageMaker domai
D. Configure AWS Budgets to send an alert when the threshold is reached.
E. Add resource tagging by editing each user's IAM profil
F. Configure AWS Cost Explorer to send an alert when the threshold is reached.
G. Add resource tagging by editing each user's IAM profil
H. Configure AWS Budgets to send an alert when the threshold is reached.

Answer: B

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
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 10
A company needs to give its ML engineers appropriate access to training data. The ML engineers must access training data from only their own business group. The ML engineers must not be allowed to access training data from other business groups.
The company uses a single AWS account and stores all the training data in Amazon S3 buckets. All ML model training occurs in Amazon SageMaker.
Which solution will provide the ML engineers with the appropriate access?

A. Enable S3 bucket versioning.
B. Configure S3 Object Lock settings for each user.
C. Add cross-origin resource sharing (CORS) policies to the S3 buckets.
D. Create IAM policie
E. Attach the policies to IAM users or IAM roles.

Answer: D

NEW QUESTION 11
FILL IN THE BLANK
A company stores time-series data about user clicks in an Amazon S3 bucket. The raw data consists of millions of rows of user activity every day. ML engineers access the data to develop their ML models.
The ML engineers need to generate daily reports and analyze click trends over the past 3 days by using Amazon Athena. The company must retain the data for 30 days before archiving the data.
Which solution will provide the HIGHEST performance for data retrieval?

A. Keep all the time-series data without partitioning in the S3 bucke
B. Manually move data that is older than 30 days to separate S3 buckets.
C. Create AWS Lambda functions to copy the time-series data into separate S3 bucket
D. Apply S3 Lifecycle policies to archive data that is older than 30 days to S3 Glacier Flexible Retrieval.
E. Organize the time-series data into partitions by date prefix in the S3 bucke
F. Apply S3 Lifecycle policies to archive partitions that are older than 30 days to S3 Glacier Flexible Retrieval.
G. Put each day's time-series data into its own S3 bucke
H. Use S3 Lifecycle policies to archive S3 buckets that hold data that is older than 30 days to S3 Glacier Flexible Retrieval.

Answer: C

NEW QUESTION 13
A company has a binary classification model in production. An ML engineer needs to develop a new version of the model.
The new model version must maximize correct predictions of positive labels and negative labels. The ML engineer must use a metric to recalibrate the model to meet these requirements.
Which metric should the ML engineer use for the model recalibration?

A. Accuracy
B. Precision
C. Recall
D. Specificity

Answer: A

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 18
Case Study
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. Configure the application to invoke an AWS Lambda function that runs a SageMaker Clarify job.
B. Invoke an AWS Lambda function to pull the sagemaker-model-monitor-analyzer built-in SageMaker image.
C. Use AWS Glue Data Quality to monitor bias.
D. Use SageMaker notebooks to compare the bias.

Answer: A

NEW QUESTION 19
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 22
An ML engineer has developed a binary classification model outside of Amazon SageMaker. The ML engineer needs to make the model accessible to a SageMaker Canvas user for additional tuning.
The model artifacts are stored in an Amazon S3 bucket. The ML engineer and the Canvas user are part of the same SageMaker domain.
Which combination of requirements must be met so that the ML engineer can share the model with the Canvas user? (Choose two.)

A. The ML engineer and the Canvas user must be in separate SageMaker domains.
B. The Canvas user must have permissions to access the S3 bucket where the model artifacts are stored.
C. The model must be registered in the SageMaker Model Registry.
D. The ML engineer must host the model on AWS Marketplace.
E. The ML engineer must deploy the model to a SageMaker endpoint.

Answer: BC

NEW QUESTION 25
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your MLA-C01 Exam with Our Prep Materials Via below:

https://www.certleader.com/MLA-C01-dumps.html