

# Amazon-Web-Services

## Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate



#### NEW QUESTION 1

A company is building a deep learning model on Amazon SageMaker. The company uses a large amount of data as the training dataset. The company needs to optimize the model's hyperparameters to minimize the loss function on the validation dataset.

Which hyperparameter tuning strategy will accomplish this goal with the LEAST computation time?

- A. Hyperbaric!
- B. Grid search
- C. Bayesian optimization
- D. Random search

**Answer:** A

#### NEW QUESTION 2

A company is using an Amazon Redshift database as its single data source. Some of the data is sensitive.

A data scientist needs to use some of the sensitive data from the database. An ML engineer must give the data scientist access to the data without transforming the source data and without storing anonymized data in the database.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Configure dynamic data masking policies to control how sensitive data is shared with the data scientist at query time.
- B. Create a materialized view with masking logic on top of the database.
- C. Grant the necessary read permissions to the data scientist.
- D. Unload the Amazon Redshift data to Amazon S3. Use Amazon Athena to create schema-on-read with masking logic.
- E. Share the view with the data scientist.
- F. Unload the Amazon Redshift data to Amazon S3. Create an AWS Glue job to anonymize the data.
- G. Share the dataset with the data scientist.

**Answer:** A

#### NEW QUESTION 3

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 table.
- 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 bucket.
- 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 4

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 pre-made 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 5

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 text.
- 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 text.
- 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 text.
- 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 text.
- H. Use Amazon Bedrock with the Stable Diffusion model to summarize the text.

**Answer:** B

#### NEW QUESTION 6

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 is experimenting with consecutive training jobs.  
How can the company MINIMIZE infrastructure startup times for these jobs?

- A. Use Managed Spot Training.
- B. Use SageMaker managed warm pools.
- C. Use SageMaker Training Compiler.
- D. Use the SageMaker distributed data parallelism (SMDDP) library.

**Answer: B**

#### NEW QUESTION 7

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 8

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 CloudTrail
- F. Use an Amazon CloudWatch dashboard for monitoring
- 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 monitoring
- I. Set up a CloudWatch alarm to provide notification when the threshold is breached.

**Answer: D**

#### NEW QUESTION 9

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 job
- B. Publish the findings to an Amazon Simple Notification Service (Amazon SNS) topic
- 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 job
- E. Use the metrics 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 step
- 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 job
- I. Configure the rules to initiate the predefined actions.

**Answer: D**

#### 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.)

City (name) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

Type\_year (type of home and year the home was built) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

Size of the building (square feet or square meters) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

City (name) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

Type\_year (type of home and year the home was built) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

Size of the building (square feet or square meters) Select...

- Select...
- Feature splitting
- Logarithmic transformation
- One-hot encoding
- Standardized distribution

**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 11

A company has developed a new ML model. The company requires online model validation on 10% of the traffic before the company fully releases the model in production. The company uses an Amazon SageMaker endpoint behind an Application Load Balancer (ALB) to serve the model. Which solution will set up the required online validation with the LEAST operational overhead?

- A. Use production variants to add the new model to the existing SageMaker endpoint
- B. Set the variant weight to 0.1 for the new mode
- C. Monitor the number of invocations by using Amazon CloudWatch.
- D. Use production variants to add the new model to the existing SageMaker endpoint
- E. Set the variant weight to 1 for the new mode
- F. Monitor the number of invocations by using Amazon CloudWatch.
- G. Create a new SageMaker endpoint
- H. Use production variants to add the new model to the new endpoint
- I. Monitor the number of invocations by using Amazon CloudWatch.
- J. Configure the ALB to route 10% of the traffic to the new model at the existing SageMaker endpoint
- K. Monitor the number of invocations by using AWS CloudTrail.

**Answer: A**

#### NEW QUESTION 12

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 15

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 16

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 17

An ML engineer needs to use AWS CloudFormation to create an ML model that an Amazon SageMaker endpoint will host.

Which resource should the ML engineer declare in the CloudFormation template to meet this requirement?

- A. AWS::SageMaker::Model

- B. AWS::SageMaker::Endpoint
- C. AWS::SageMaker::NotebookInstance
- D. AWS::SageMaker::Pipeline

**Answer:** A

#### 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 bucket
- 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 upload
- 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.
- G. The data contains meaningful ordered features with sensitive information that should not be discarded
- H. An ML engineer must ensure that the sensitive data is masked before another team starts to build the model.
- I. Use Amazon Macie 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 23

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.

Which AWS service or feature can aggregate the data from the various data sources?

- A. Amazon EMR Spark jobs
- B. Amazon Kinesis Data Streams
- C. Amazon DynamoDB
- D. AWS Lake Formation

**Answer:** A

#### NEW QUESTION 27

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. K-means clustering
- D. Neural Topic Model (NTM)

**Answer:** B

#### NEW QUESTION 30

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 categories
- B. Add the tags to the model packages in the SageMaker Model Registry.
- C. Create a model group for each category
- 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 categories
- G. Move the existing model groups into the collections.

**Answer:** A

#### NEW QUESTION 34

A credit card company has a fraud detection model in production on an Amazon SageMaker endpoint. The company develops a new version of the model. The company needs to assess the new model's performance by using live data and without affecting production end users. Which solution will meet these requirements?

- A. Set up SageMaker Debugger and create a custom rule.
- B. Set up blue/green deployments with all-at-once traffic shifting.
- C. Set up blue/green deployments with canary traffic shifting.
- D. Set up shadow testing with a shadow variant of the new model.

**Answer:** D

**NEW QUESTION 36**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions 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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The MLA-C01 Practice Test Here](#)**