



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 endpoint
- E. Analyze prediction differences on samples.

Answer: B

NEW QUESTION 2

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 3

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 endpoint
- F. Deploy multiple copies of the model to the endpoint
- 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 endpoint
- 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 4

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 5

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 bucke
- B. Query the objects from the table.
- C. Create a new S3 bucket for processed dat
- D. Set up S3 replication from the central S3 bucket to the new S3 bucke
- E. Use S3 Object Lambda to query the objects based on transaction date.
- F. Create a new S3 bucket for processed dat
- G. Use AWS Glue for Apache Spark to create a job to query the CSV objects based on transaction dat
- H. Configure the job to store the results in the new S3 bucke
- I. Query the objects from the new S3 bucket.
- J. Create a new S3 bucket for processed dat

- K. Use Amazon Data Firehose to transfer the data from the central S3 bucket to the new S3 bucket.
- L. Configure Firehose to run an AWS Lambda function to query the data based on transaction date.

Answer: A

NEW QUESTION 6

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 7

A company is planning to use Amazon SageMaker to make classification ratings that are based on images. The company has 6 of training data that is stored on an Amazon FSx for NetApp ONTAP system virtual machine (SVM). The SVM is in the same VPC as SageMaker. An ML engineer must make the training data accessible for ML models that are in the SageMaker environment. Which solution will meet these requirements?

- A. Mount the FSx for ONTAP file system as a volume to the SageMaker Instance.
- B. Create an Amazon S3 bucket
- C. Use Mountpoint for Amazon S3 to link the S3 bucket to the FSx for ONTAP file system.
- D. Create a catalog connection from SageMaker Data Wrangler to the FSx for ONTAP file system.
- E. Create a direct connection from SageMaker Data Wrangler to the FSx for ONTAP file system.

Answer: A

NEW QUESTION 8

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 9

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

Answer: A

NEW QUESTION 10

A company has trained an ML model in Amazon SageMaker. The company needs to host the model to provide inferences in a production environment. The model must be highly available and must respond with minimum latency. The size of each request will be between 1 KB and 3 MB. The model will receive unpredictable bursts of requests during the day. The inferences must adapt proportionally to the changes in demand. How should the company deploy the model into production to meet these requirements?

- A. Create a SageMaker real-time inference endpoint
- B. Configure auto scaling
- C. Configure the endpoint to present the existing model.

- D. Deploy the model on an Amazon Elastic Container Service (Amazon ECS) cluster.
- E. Use ECS scheduled scaling that is based on the CPU of the ECS cluster.
- F. Install SageMaker Operator on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- G. Deploy the model in Amazon EKS.
- H. Set horizontal pod auto scaling to scale replicas based on the memory metric.
- I. Use Spot Instances with a Spot Fleet behind an Application Load Balancer (ALB) for inference.
- J. Use the ALBRequestCountPerTarget metric as the metric for auto scaling.

Answer: A

NEW QUESTION 10

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 13

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 15

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 20

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 25

An ML engineer is developing a fraud detection model by using the Amazon SageMaker XGBoost algorithm. The model classifies transactions as either fraudulent or legitimate. During testing, the model excels at identifying fraud in the training dataset. However, the model is inefficient at identifying fraud in new and unseen transactions. What should the ML engineer do to improve the fraud detection for new transactions?

- A. Increase the learning rate.
- B. Remove some irrelevant features from the training dataset.
- C. Increase the value of the max_depth hyperparameter.
- D. Decrease the value of the max_depth hyperparameter.

Answer: D

NEW QUESTION 26

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 28

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 bucket
- 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 bucket
- 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 bucket
- 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 29

An ML engineer has trained a neural network by using stochastic gradient descent (SGD). The neural network performs poorly on the test set. The values for training loss and validation loss remain high and show an oscillating pattern. The values decrease for a few epochs and then increase for a few epochs before repeating the same cycle.

What should the ML engineer do to improve the training process?

- A. Introduce early stopping.
- B. Increase the size of the test set.
- C. Increase the learning rate.
- D. Decrease the learning rate.

Answer: D

NEW QUESTION 34

A company is using ML to predict the presence of a specific weed in a farmer's field. The company is using the Amazon SageMaker linear learner built-in algorithm with a value of multiclass_dasser for the predictor_type hyperparameter.

What should the company do to MINIMIZE false positives?

- A. Set the value of the weight decay hyperparameter to zero.
- B. Increase the number of training epochs.
- C. Increase the value of the target_precision hyperparameter.
- D. Change the value of the predictor_type hyperparameter to regressor.

Answer: C

NEW QUESTION 37

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 38

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 40

A company has implemented a data ingestion pipeline for sales transactions from its ecommerce website. The company uses Amazon Data Firehose to ingest data into Amazon OpenSearch Service. The buffer interval of the Firehose stream is set for 60 seconds. An OpenSearch linear model generates real-time sales forecasts based on the data and presents the data in an OpenSearch dashboard.

The company needs to optimize the data ingestion pipeline to support sub-second latency for the real-time dashboard.

Which change to the architecture will meet these requirements?

- A. Use zero buffering in the Firehose stream
- B. Tune the batch size that is used in the PutRecordBatch operation.
- C. Replace the Firehose stream with an AWS DataSync task
- D. Configure the task with enhanced fan-out consumers.
- E. Increase the buffer interval of the Firehose stream from 60 seconds to 120 seconds.
- F. Replace the Firehose stream with an Amazon Simple Queue Service (Amazon SQS) queue.

Answer: A

NEW QUESTION 42

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 45

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 47

A company wants to improve the sustainability of its ML operations.

Which actions will reduce the energy usage and computational resources that are associated with the company's training jobs? (Choose two.)

- A. Use Amazon SageMaker Debugger to stop training jobs when non-converging conditions are detected.
- B. Use Amazon SageMaker Ground Truth for data labeling.
- C. Deploy models by using AWS Lambda functions.
- D. Use AWS Trainium instances for training.
- E. Use PyTorch or TensorFlow with the distributed training option.

Answer: AD

NEW QUESTION 50

.....

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

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