



Amazon-Web-Services

Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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 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 3

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 4

A company has a large collection of chat recordings from customer interactions after a product release. An ML engineer needs to create an ML model to analyze the chat data. The ML engineer needs to determine the success of the product by reviewing customer sentiments about the product.

Which action should the ML engineer take to complete the evaluation in the LEAST amount of time?

- A. Use Amazon Rekognition to analyze sentiments of the chat conversations.
- B. Train a Naive Bayes classifier to analyze sentiments of the chat conversations.
- C. Use Amazon Comprehend to analyze sentiments of the chat conversations.
- D. Use random forests to classify sentiments of the chat conversations.

Answer: C

NEW QUESTION 5

A financial company receives a high volume of real-time market data streams from an external provider. The streams consist of thousands of JSON records every second.

The company needs to implement a scalable solution on AWS to identify anomalous data points.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Ingest real-time data into Amazon Kinesis data stream
- B. Use the built-in RANDOM_CUT_FOREST function in Amazon Managed Service for Apache Flink to process the data streams and to detect data anomalies.
- C. Ingest real-time data into Amazon Kinesis data stream
- D. Deploy an Amazon SageMaker endpoint for real-time outlier detection
- E. Create an AWS Lambda function to detect anomalies
- F. Use the data streams to invoke the Lambda function.
- G. Ingest real-time data into Apache Kafka on Amazon EC2 instance
- H. Deploy an Amazon SageMaker endpoint for real-time outlier detection
- I. Create an AWS Lambda function to detect anomalies
- J. Use the data streams to invoke the Lambda function.
- K. Send real-time data to an Amazon Simple Queue Service (Amazon SQS) FIFO queue
- L. Create an AWS Lambda function to consume the queue message
- M. Program the Lambda function to start an AWS Glue extract, transform, and load (ETL) job for batch processing and anomaly detection.

Answer: A

NEW QUESTION 6

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 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

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

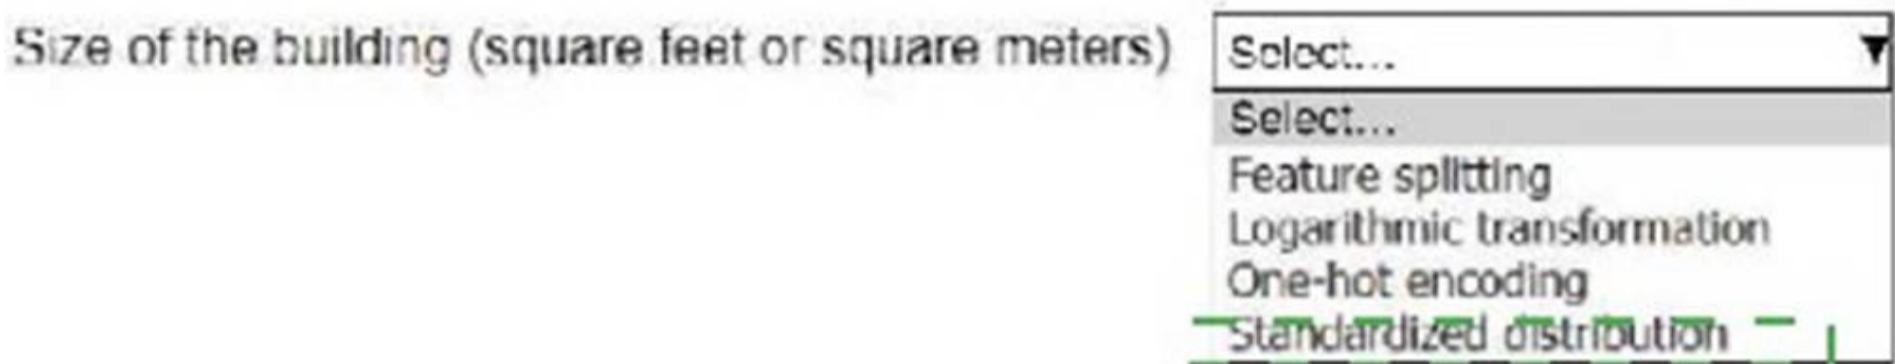
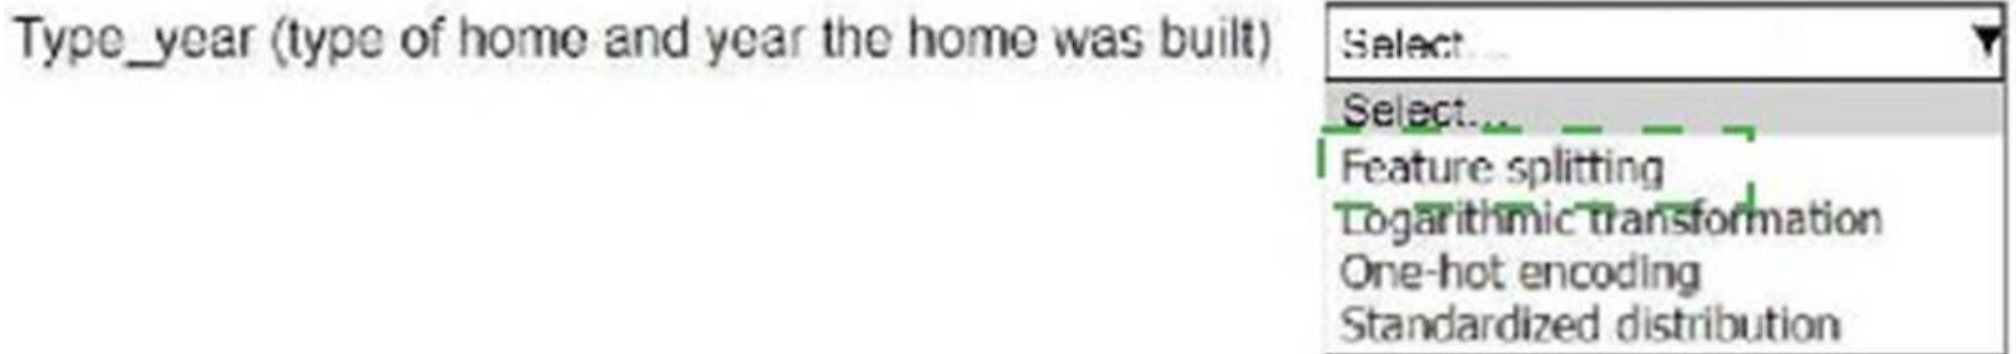
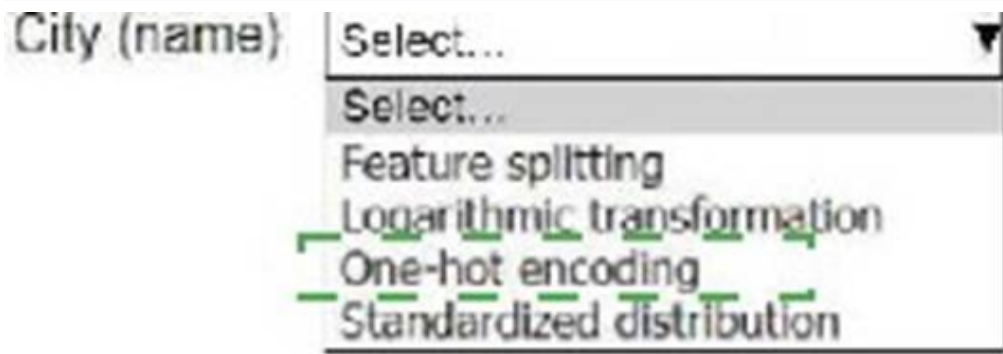
Size of the building (square feet or square meters)

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 9

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 domain
- 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 domain
- D. Configure AWS Budgets to send an alert when the threshold is reached.
- E. Add resource tagging by editing each user's IAM profile
- F. Configure AWS Cost Explorer to send an alert when the threshold is reached.
- G. Add resource tagging by editing each user's IAM profile
- H. Configure AWS Budgets to send an alert when the threshold is reached.

Answer: B

NEW QUESTION 10

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 Kendra
- H. Query Amazon Kendra to perform the semantic searches.
- I. Use an Amazon Textract asynchronous job to ingest the documents from the S3 bucket
- J. Query Amazon Textract to perform the semantic searches.

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 policies

E. Attach the policies to IAM users or IAM roles.

Answer: D

NEW QUESTION 12

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.

Before the ML engineer trains the model, the ML engineer must resolve the issue of the imbalanced data.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon Athena to identify patterns that contribute to the imbalance.
- B. Adjust the dataset accordingly.
- C. Use Amazon SageMaker Studio Classic built-in algorithms to process the imbalanced dataset.
- D. Use AWS Glue DataBrew built-in features to oversample the minority class.
- E. Use the Amazon SageMaker Data Wrangler balance data operation to oversample the minority class.

Answer: D

NEW QUESTION 15

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 18

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_classifier` 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 20

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 23

A company needs to create a central catalog for all the company's ML models. The models are in AWS accounts where the company developed the models initially. The models are hosted in Amazon Elastic Container Registry (Amazon ECR) repositories.

Which solution will meet these requirements?

- A. Configure ECR cross-account replication for each existing ECR repository.
- B. Ensure that each model is visible in each AWS account.
- C. Create a new AWS account with a new ECR repository as the central catalog.
- D. Configure ECR cross-account replication between the initial ECR repositories and the central catalog.
- E. Use the Amazon SageMaker Model Registry to create a model group for models hosted in Amazon ECR.
- F. Create a new AWS account.
- G. In the new account, use the SageMaker Model Registry as the central catalog.
- H. Attach a cross-account resource policy to each model group in the initial AWS accounts.
- I. Use an AWS Glue Data Catalog to store the model.
- J. Run an AWS Glue crawler to migrate the models from the ECR repositories to the Data Catalog.
- K. Configure cross-account access to the Data Catalog.

Answer: C

NEW QUESTION 26

An ML engineer trained an ML model on Amazon SageMaker to detect automobile accidents from closed-circuit TV footage. The ML engineer used SageMaker Data Wrangler to create a training dataset of images of accidents and non-accidents. The model performed well during training and validation. However, the model is underperforming in production because of variations in the quality of the images from various cameras. Which solution will improve the model's accuracy in the LEAST amount of time?

- A. Collect more images from all the camera
- B. Use Data Wrangler to prepare a new trainingdataset.
- C. Recreate the training dataset by using the Data Wrangler corrupt image transfor
- D. Specify the impulse noise option.
- E. Recreate the training dataset by using the Data Wrangler enhance image contrast transfor
- F. Specify the Gamma contrast option.
- G. Recreate the training dataset by using the Data Wrangler resize image transfor
- H. Crop all images to the same size.

Answer: B

NEW QUESTION 30

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

Text representation of basic units of data processed by LLMs

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

Token

High-dimensional vectors that contain the semantic meaning of text

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

Token

Enrichment of information from additional data sources to improve a generated response

Select...

Select...

Embedding

Retrieval Augmented Generation (RAG)

Temperature

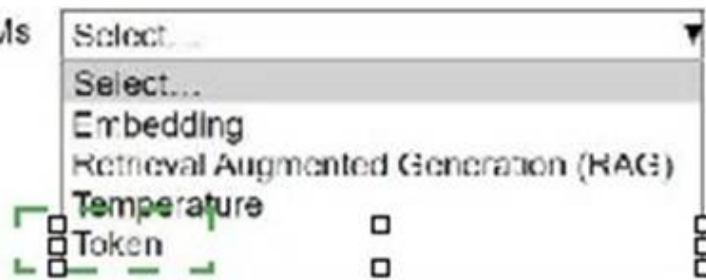
Token

- A. Mastered
- B. Not Mastered

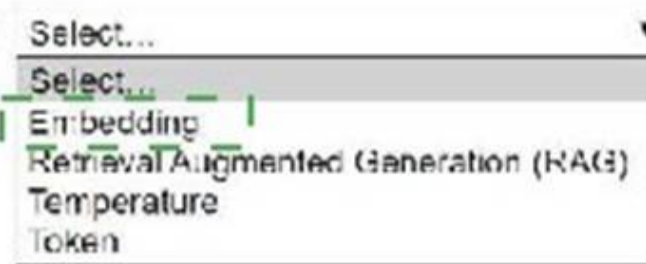
Answer: A

Explanation:

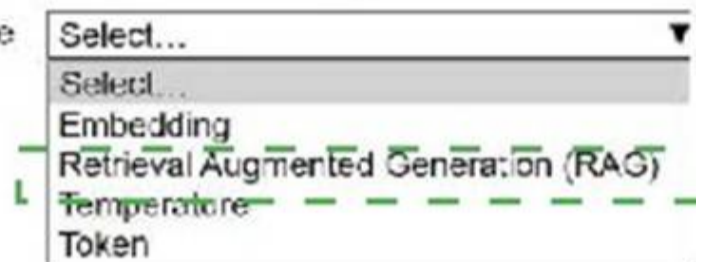
Text representation of basic units of data processed by LLMs



High-dimensional vectors that contain the semantic meaning of text



Enrichment of information from additional data sources to improve a generated response



NEW QUESTION 31

An ML engineer needs to use an ML model to predict the price of apartments in a specific location. Which metric should the ML engineer use to evaluate the model's performance?

- A. Accuracy
- B. Area Under the ROC Curve (AUC)
- C. F1 score
- D. Mean absolute error (MAE)

Answer: D

NEW QUESTION 34

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. Which solution will meet these requirements?
- 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 36

A company is planning to use Amazon Redshift ML in its primary AWS account. The source data is in an Amazon S3 bucket in a secondary account.

An ML engineer needs to set up an ML pipeline in the primary account to access the S3 bucket in the secondary account. The solution must not require public IPv4 addresses.

Which solution will meet these requirements?

- A. Provision a Redshift cluster and Amazon SageMaker Studio in a VPC with no public access enabled in the primary account
- B. Create a VPC peering connection between the accounts
- C. Update the VPC route tables to remove the route to 0.0.0.0/0.
- D. Provision a Redshift cluster and Amazon SageMaker Studio in a VPC with no public access enabled in the primary account
- E. Create an AWS Direct Connect connection and a transit gateway
- F. Associate the VPCs from both accounts with the transit gateway
- G. Update the VPC route tables to remove the route to 0.0.0.0/0.
- H. Provision a Redshift cluster and Amazon SageMaker Studio in a VPC in the primary account
- I. Create an AWS Site-to-Site VPN connection with two encrypted IPsec tunnels between the accounts
- J. Set up interface VPC endpoints for Amazon S3.
- K. Provision a Redshift cluster and Amazon SageMaker Studio in a VPC in the primary account
- L. Create an S3 gateway endpoint
- M. Update the S3 bucket policy to allow IAM principals from the primary account

N. Set up interface VPC endpoints for SageMaker and Amazon Redshift.

Answer: D

NEW QUESTION 38

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 41

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

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 Instance
- 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 47

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 51

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 53

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 56

.....

Relate Links

100% Pass Your MLA-C01 Exam with ExamBible Prep Materials

<https://www.exambible.com/MLA-C01-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>