

# Amazon-Web-Services

## Exam Questions MLA-C01

AWS Certified Machine Learning Engineer - Associate



#### NEW QUESTION 1

An ML engineer is using Amazon SageMaker to train a deep learning model that requires distributed training. After some training attempts, the ML engineer observes that the instances are not performing as expected. The ML engineer identifies communication overhead between the training instances. What should the ML engineer do to MINIMIZE the communication overhead between the instances?

- A. Place the instances in the same VPC subne
- B. Store the data in a different AWS Region from where the instances are deployed.
- C. Place the instances in the same VPC subnet but in different Availability Zone
- D. Store the data in a different AWS Region from where the instances are deployed.
- E. Place the instances in the same VPC subne
- F. Store the data in the same AWS Region and Availability Zone where the instances are deployed.
- G. Place the instances in the same VPC subne
- H. Store the data in the same AWS Region but in a different Availability Zone from where the instances are deployed.

**Answer: C**

#### NEW QUESTION 2

A company uses Amazon Athena to query a dataset in Amazon S3. The dataset has a target variable that the company wants to predict. The company needs to use the dataset in a solution to determine if a model can predict the target variable. Which solution will provide this information with the LEAST development effort?

- A. Create a new model by using Amazon SageMaker Autopilo
- B. Report the model's achieved performance.
- C. Implement custom scripts to perform data pre-processing, multiple linear regression, and performance evaluatio
- D. Run the scripts on Amazon EC2 instances.
- E. Configure Amazon Macie to analyze the dataset and to create a mode
- F. Report the model's achieved performance.
- G. Select a model from Amazon Bedroc
- H. Tune the model with the dat
- I. Report the model's achieved performance.

**Answer: A**

#### NEW QUESTION 3

A company is using an AWS Lambda function to monitor the metrics from an ML model. An ML engineer needs to implement a solution to send an email message when the metrics breach a threshold. Which solution will meet this requirement?

- A. Log the metrics from the Lambda function to AWS CloudTrai
- B. Configure a CloudTrail trail to send the email message.
- C. Log the metrics from the Lambda function to Amazon CloudFron
- D. Configure an Amazon CloudWatch alarm to send the email message.
- E. Log the metrics from the Lambda function to Amazon CloudWatc
- F. Configure a CloudWatch alarm to send the email message.
- G. Log the metrics from the Lambda function to Amazon CloudWatc
- H. Configure an Amazon CloudFront rule to send the email message.

**Answer: D**

#### NEW QUESTION 4

A company has an application that uses different APIs to generate embeddings for input text. The company needs to implement a solution to automatically rotate the API tokens every 3 months. Which solution will meet this requirement?

- A. Store the tokens in AWS Secrets Manage
- B. Create an AWS Lambda function to perform the rotation.
- C. Store the tokens in AWS Systems Manager Parameter Stor
- D. Create an AWS Lambda function to perform the rotation.
- E. Store the tokens in AWS Key Management Service (AWS KMS). Use an AWS managed key to perform the rotation.
- F. Store the tokens in AWS Key Management Service (AWS KMS). Use an AWS owned key to perform the rotation.

**Answer: A**

#### NEW QUESTION 5

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

Answer: D

#### NEW QUESTION 6

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 7

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 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 address
- B. Assign the security group to the domain.
- C. Create a network ACL inbound rule to deny traffic from the specific IP address
- D. Assign the rule to the default network ACL for the subnet where the domain is located.
- E. Create a shadow variant for the domain
- 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 address
- H. Assign the route table to the domain.

Answer: B

#### NEW QUESTION 9

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 metrics
- 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 metrics
- 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 10

A company has a conversational AI assistant that sends requests through Amazon Bedrock to an Anthropic Claude large language model (LLM). Users report that when they ask similar questions multiple times, they sometimes receive different answers. An ML engineer needs to improve the responses to be more consistent and less random.

Which solution will meet these requirements?

- A. Increase the temperature parameter and the top\_k parameter.
- B. Increase the temperature parameter
- C. Decrease the top\_k parameter.
- D. Decrease the temperature parameter
- E. Increase the top\_k parameter.
- F. Decrease the temperature parameter and the top\_k parameter.

**Answer: D**

#### NEW QUESTION 10

A company wants to reduce the cost of its containerized ML applications. The applications use ML models that run on Amazon EC2 instances, AWS Lambda functions, and an Amazon Elastic Container Service (Amazon ECS) cluster. The EC2 workloads and ECS workloads use Amazon Elastic Block Store (Amazon EBS) volumes to save predictions and artifacts.

An ML engineer must identify resources that are being used inefficiently. The ML engineer also must generate recommendations to reduce the cost of these resources.

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

- A. Create code to evaluate each instance's memory and compute usage.
- B. Add cost allocation tags to the resource
- C. Activate the tags in AWS Billing and Cost Management.
- D. Check AWS CloudTrail event history for the creation of the resources.
- E. Run AWS Compute Optimizer.

**Answer: D**

#### NEW QUESTION 11

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 15

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 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 must implement a manual approval-based workflow to ensure that only approved models can be deployed to production endpoints.

Which solution will meet this requirement?

- A. Use SageMaker Experiments to facilitate the approval process during model registration.
- B. Use SageMaker ML Lineage Tracking on the central model registry
- C. Create tracking entities for the approval process.
- D. Use SageMaker Model Monitor to evaluate the performance of the model and to manage the approval.
- E. Use SageMaker Pipeline
- F. When a model version is registered, use the AWS SDK to change the approval status to "Approved."

**Answer: D**

### NEW QUESTION 23

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 use the central model registry to manage different versions of models in the application.

Which action will meet this requirement with the LEAST operational overhead?

- A. Create a separate Amazon Elastic Container Registry (Amazon ECR) repository for each model.
- B. Use Amazon Elastic Container Registry (Amazon ECR) and unique tags for each model version.
- C. Use the SageMaker Model Registry and model groups to catalog the models.
- D. Use the SageMaker Model Registry and unique tags for each model version.

**Answer: C**

### NEW QUESTION 26

An ML engineer needs to use Amazon SageMaker to fine-tune a large language model (LLM) for text summarization. The ML engineer must follow a low-code no-code (LCNC) approach.

Which solution will meet these requirements?

- A. Use SageMaker Studio to fine-tune an LLM that is deployed on Amazon EC2 instances.
- B. Use SageMaker Autopilot to fine-tune an LLM that is deployed by a custom API endpoint.
- C. Use SageMaker Autopilot to fine-tune an LLM that is deployed on Amazon EC2 instances.
- D. Use SageMaker Autopilot to fine-tune an LLM that is deployed by SageMaker JumpStart.

**Answer: D**

### NEW QUESTION 30

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 33

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 38

A company uses Amazon SageMaker for its ML workloads. The company's ML engineer receives a 50 MB Apache Parquet data file to build a fraud detection model. The file includes several correlated columns that are not required.

What should the ML engineer do to drop the unnecessary columns in the file with the LEAST effort?

- A. Download the file to a local workstation.
- B. Perform one-hot encoding by using a custom Python script.
- C. Create an Apache Spark job that uses a custom processing script on Amazon EMR.
- D. Create a SageMaker processing job by calling the SageMaker Python SDK.
- E. Create a data flow in SageMaker Data Wrangler.
- F. Configure a transform step.

**Answer: D**

### NEW QUESTION 43

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 predictorjype hyperparameter to regressor.

**Answer: C**

#### NEW QUESTION 48

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 49

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 repositior
- 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 catalo
- 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 EC
- F. Create a new AWS accoun
- G. In the new account, use the SageMaker Model Registry as the central catalo
- 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 Catalo
- K. Configure cross- account access to the Data Catalog.

**Answer: C**

#### NEW QUESTION 52

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 53

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 56

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 58

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 59**

A company has AWS Glue data processing jobs that are orchestrated by an AWS Glue workflow. The AWS Glue jobs can run on a schedule or can be launched manually.

The company is developing pipelines in Amazon SageMaker Pipelines for ML model development. The pipelines will use the output of the AWS Glue jobs during the data processing phase of model development. An ML engineer needs to implement a solution that integrates the AWS Glue jobs with the pipelines.

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

- A. Use AWS Step Functions for orchestration of the pipelines and the AWS Glue jobs.
- B. Use processing steps in SageMaker Pipeline
- C. Configure inputs that point to the Amazon Resource Names (ARNs) of the AWS Glue jobs.
- D. Use Callback steps in SageMaker Pipelines to start the AWS Glue workflow and to stop the pipelines until the AWS Glue jobs finish running.
- E. Use Amazon EventBridge to invoke the pipelines and the AWS Glue jobs in the desired order.

**Answer: C**

**NEW QUESTION 64**

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 account
- 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 account
- 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 69**

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 71**

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 74

An ML engineer needs to deploy ML models to get inferences from large datasets in an asynchronous manner. The ML engineer also needs to implement scheduled monitoring of the data quality of the models. The ML engineer must receive alerts when changes in data quality occur. Which solution will meet these requirements?

- A. Deploy the models by using scheduled AWS Glue job
- B. Use Amazon CloudWatch alarms to monitor the data quality and to send alerts.
- C. Deploy the models by using scheduled AWS Batch job
- D. Use AWS CloudTrail to monitor the data quality and to send alerts.
- E. Deploy the models by using Amazon Elastic Container Service (Amazon ECS) on AWS Fargat
- F. Use Amazon EventBridge to monitor the data quality and to send alerts.
- G. Deploy the models by using Amazon SageMaker batch transfor
- H. Use SageMaker Model Monitor to monitor the data quality and to send alerts.

**Answer: D**

#### NEW QUESTION 79

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.

After the data is aggregated, the ML engineer must implement a solution to automatically detect anomalies in the data and to visualize the result.

Which solution will meet these requirements?

- A. Use Amazon Athena to automatically detect the anomalies and to visualize the result.
- B. Use Amazon Redshift Spectrum to automatically detect the anomalie
- C. Use Amazon QuickSight to visualize the result.
- D. Use Amazon SageMaker Data Wrangler to automatically detect the anomalies and to visualize the result.
- E. Use AWS Batch to automatically detect the anomalie
- F. Use Amazon QuickSight to visualize the result.

**Answer: C**

#### NEW QUESTION 83

An ML engineer receives datasets that contain missing values, duplicates, and extreme outliers. The ML engineer must consolidate these datasets into a single data frame and must prepare the data for ML.

Which solution will meet these requirements?

- A. Use Amazon SageMaker Data Wrangler to import the datasets and to consolidate them into a single data fram
- B. Use the cleansing and enrichment functionalities to prepare the data.
- C. Use Amazon SageMaker Ground Truth to import the datasets and to consolidate them into a single data fram
- D. Use the human-in-the-loop capability to prepare the data.
- E. Manually import and merge the dataset
- F. Consolidate the datasets into a single data fram
- G. Use Amazon Q Developer to generate code snippets that will prepare the data.
- H. Manually import and merge the dataset
- I. Consolidate the datasets into a single data fram
- J. Use Amazon SageMaker data labeling to prepare the data.

**Answer: A**

#### NEW QUESTION 88

.....

## 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](#)