

MLA-C01 Exam Consultant - 100% Realistic Questions Pool



BONUS!!! Download part of TestValid MLA-C01 dumps for free: <https://drive.google.com/open?id=19W34FWUX9Hssmv-gOnJz5hAYq46VZs64>

The Amazon MLA-C01 pdf questions learning material provided to the customers from TestValid is in three different formats. The first format is PDF format which is printable and portable. It means it can be accessed from tablets, laptops, and smartphones to prepare for the AWS Certified Machine Learning Engineer - Associate exam. The Amazon MLA-C01 PdfFormat can be used offline, and candidates can even prepare for it in the classroom or library by printing questions or on their smart devices.

Amazon MLA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> • Deployment and Orchestration of ML Workflows: This section of the exam measures skills of Forensic Data Analysts and focuses on deploying machine learning models into production environments. It covers choosing the right infrastructure, managing containers, automating scaling, and orchestrating workflows through CI • CD pipelines. Candidates must be able to build and script environments that support consistent deployment and efficient retraining cycles in real-world fraud detection systems.
Topic 2	<ul style="list-style-type: none"> • Data Preparation for Machine Learning (ML): This section of the exam measures skills of Forensic Data Analysts and covers collecting, storing, and preparing data for machine learning. It focuses on understanding different data formats, ingestion methods, and AWS tools used to process and transform data. Candidates are expected to clean and engineer features, ensure data integrity, and address biases or compliance issues, which are crucial for preparing high-quality datasets in fraud analysis contexts.
Topic 3	<ul style="list-style-type: none"> • ML Solution Monitoring, Maintenance, and Security: This section of the exam measures skills of Fraud Examiners and assesses the ability to monitor machine learning models, manage infrastructure costs, and apply security best practices. It includes setting up model performance tracking, detecting drift, and using AWS tools for logging and alerts. Candidates are also tested on configuring access controls, auditing environments, and maintaining compliance in sensitive data environments like financial fraud detection.
Topic 4	<ul style="list-style-type: none"> • ML Model Development: This section of the exam measures skills of Fraud Examiners and covers choosing and training machine learning models to solve business problems such as fraud detection. It includes selecting algorithms, using built-in or custom models, tuning parameters, and evaluating performance with standard metrics. The domain emphasizes refining models to avoid overfitting and maintaining version control to support ongoing investigations and audit trails.

MLA-C01 Latest Exam Discount | MLA-C01 Certification Test Answers

TestValid is website that can take you access to the road of success. TestValid can provide the quickly passing Amazon certification MLA-C01 exam training materials for you, which enable you to grasp the knowledge of the certification exam within a short period of time, and pass Amazon Certification MLA-C01 Exam for only one-time.

Amazon AWS Certified Machine Learning Engineer - Associate Sample Questions (Q11-Q16):

NEW QUESTION # 11

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 permissions. Attach the role to each notebook instance that the team uses.
- B. Create a single IAM group. Add the data scientists to the group. Create an IAM role. Attach the AdministratorAccess AWS managed IAM policy to the role. Associate the role with the group. Associate the group with each notebook instance that the team uses.
- C. Create a single IAM group. Add the data scientists to the group. Associate the group with each notebook instance that the team uses.
- D. Create a single IAM user. Attach the AdministratorAccess AWS managed IAM policy to the user. Configure each notebook instance to use the IAM user.

Answer: A

NEW QUESTION # 12

A company uses an Amazon SageMaker AI model for real-time inference with auto scaling enabled. During peak usage, new instances launch before existing instances are fully ready, causing inefficiencies and delays.

Which solution will optimize the scaling process without affecting response times?

- A. Increase the cooldown period after scale-out activities.
- B. Decrease the scale-in cooldown period and increase the maximum instance count.
- C. Change to a multi-model endpoint configuration.
- D. Integrate Amazon API Gateway and AWS Lambda to manage invocations.

Answer: A

Explanation:

Amazon SageMaker auto scaling uses cooldown periods to control how frequently scaling activities occur.

When scale-out happens too quickly, new instances may receive traffic before they are fully initialized, leading to inefficiencies and latency.

AWS documentation recommends increasing the scale-out cooldown period to give newly launched instances sufficient time to initialize and become healthy before additional scaling events occur. This ensures stable performance during traffic spikes without impacting response times.

Multi-model endpoints address model hosting efficiency, not scaling timing. API Gateway and Lambda add unnecessary latency and complexity. Decreasing scale-in cooldown does not address scale-out issues.

Therefore, Option D is the correct and AWS-aligned solution.

NEW QUESTION # 13

A company has a custom extract, transform, and load (ETL) process that runs on premises. The ETL process is written in the R language and runs for an average of 6 hours. The company wants to migrate the process to run on AWS.

Which solution will meet these requirements?

- A. Use Amazon SageMaker AI script mode to build a Docker image. Run the ETL jobs by using SageMaker Notebook Jobs.
- B. Use AWS Glue to prepare and run the ETL jobs.

- C. Use an AWS Lambda function created from a container image to run the ETL jobs.
- D. Use Amazon SageMaker AI processing jobs with a custom Docker image stored in Amazon Elastic Container Registry (Amazon ECR).

Answer: D

Explanation:

The ETL process has two critical characteristics: it is long-running (6 hours) and written in R. AWS Lambda is unsuitable because it has a maximum execution time of 15 minutes. AWS Glue primarily supports Spark-based ETL and does not natively support custom R-based workloads.

AWS documentation recommends using Amazon SageMaker Processing Jobs for long-running, custom data processing workloads. Processing jobs allow users to run arbitrary code in custom Docker containers, making them ideal for migrating on-premises ETL jobs written in R.

By building a custom Docker image that includes the R runtime and required libraries and storing it in Amazon ECR, the company can run the ETL job at scale on managed infrastructure without rewriting the code.

SageMaker script mode is intended for training, not ETL. Therefore, SageMaker processing jobs with a custom container are the correct solution.

NEW QUESTION # 14

A company is developing a new online application to gather information from customers. An ML engineer has developed a new ML model that will determine a score for each customer. The model will use the score to determine which product to display to the customer. The ML engineer needs to minimize response-time latency for the model. How should the ML engineer deploy the application in Amazon SageMaker to meet these requirements?

- A. Configure a real-time inference endpoint.
- B. Configure batch transform.
- C. Configure an asynchronous inference endpoint.
- D. Configure a serverless inference endpoint.

Answer: A

Explanation:

To minimize response-time latency, the ML model should be deployed to a real-time inference endpoint in Amazon SageMaker. This provides low-latency predictions by keeping the model loaded and ready to handle incoming requests, which is critical for an online application serving customers in real time.

NEW QUESTION # 15

An ML engineer has an Amazon Comprehend custom model in Account A in the us-east-1 Region. The ML engineer needs to copy the model to Account # in the same Region.

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

- A. Create an AWS Site-to-Site VPN connection between Account A and Account # to transfer the model.
- B. Use AWS DataSync to replicate the model from Account A to Account B.
- C. Use Amazon S3 to make a copy of the model. Transfer the copy to Account B.
- D. Create a resource-based IAM policy. Use the Amazon Comprehend ImportModel API operation to copy the model to Account B.

Answer: D

Explanation:

Amazon Comprehend provides the ImportModel API operation, which allows you to copy a custom model between AWS accounts. By creating a resource-based IAM policy on the model in Account A, you can grant Account B the necessary permissions to access and import the model. This approach requires minimal development effort and is the AWS-recommended method for sharing custom models across accounts.

NEW QUESTION # 16

.....

