

MLA-C01 Actual Questions - Braindumps MLA-C01 Downloads



BTW, DOWNLOAD part of Easy4Engine MLA-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1p5rhF3MKG_KzGMXIxWH7uszJh8SFAas5

we can promise that our MLA-C01 study materials will be the best study materials in the world with the high pass rate as 98% to 100%. All these achievements are due to the reason that our MLA-C01 exam questions have a high quality that is unique in the market. If you decide to buy our MLA-C01 training dumps, we can make sure that you will have the opportunity to enjoy the MLA-C01 practice engine from team of experts.

There are a lot of advantages of our APP online version. On one hand, the online version of our MLA-C01 exam questions can apply in all kinds of the electronic devices. In addition, the online version of our MLA-C01 training materials can work in an offline state. If you buy our products, you have the chance to use our study materials for preparing your exam when you are in an offline state. We believe that you will like the online version of our MLA-C01 Exam Questions.

>> **MLA-C01 Actual Questions** <<

Braindumps MLA-C01 Downloads | Reliable MLA-C01 Exam Simulations

We guarantee to you that the refund process is very simple and only if you provide us the screenshot or the scanning copy of your failure marks we will refund you in full immediately. If you have doubts or problems about our MLA-C01 exam torrent, please contact our online customer service or contact us by mails and we will reply and solve your problem as quickly as we can. We won't waste your money and your time and if you fail in the exam we will refund you in full immediately at one time. We provide the best MLA-C01 Questions torrent to you and don't hope to let you feel disappointed.

Amazon MLA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 3	<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 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.

Amazon AWS Certified Machine Learning Engineer - Associate Sample Questions (Q99-Q104):

NEW QUESTION # 99

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 must host the model on AWS Marketplace.
- **B. The model must be registered in the SageMaker Model Registry.**
- C. The ML engineer and the Canvas user must be in separate SageMaker domains.
- **D. The Canvas user must have permissions to access the S3 bucket where the model artifacts are stored.**
- E. The ML engineer must deploy the model to a SageMaker endpoint.

Answer: B,D

Explanation:

The SageMaker Canvas user needs permissions to access the Amazon S3 bucket where the model artifacts are stored to retrieve the model for use in Canvas.

Registering the model in the SageMaker Model Registry allows the model to be tracked and managed within the SageMaker ecosystem. This makes it accessible for tuning and deployment through SageMaker Canvas.

This combination ensures proper access control and integration within SageMaker, enabling the Canvas user to work with the model.

NEW QUESTION # 100

An ML engineering team is spread across multiple locations. When the lead ML engineer opens an Amazon SageMaker AI notebook, the ML engineer does not see the latest merged notebook made by other team members from a Git repository.

The lead ML engineer must see the latest SageMaker AI notebook updates.

Which solution will meet this requirement?

- A. Run the `!git branch` command.
- B. Run the `!git push origin master` command.
- **C. Run the `!git pull origin master` command.**
- D. Run the `!git commit` command.

Answer: C

NEW QUESTION # 101

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

Answer: C

NEW QUESTION # 102

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

Answer: C

Explanation:

Amazon Kinesis Data Firehose allows for near real-time data streaming. Setting the buffering hint to zero or a very small value minimizes the buffering delay and ensures that records are delivered to the destination (Amazon OpenSearch Service) as quickly as possible. Additionally, tuning the batch size in the PutRecordBatch operation can further optimize the data ingestion for sub-second latency. This approach minimizes latency while maintaining the operational simplicity of using Firehose.

NEW QUESTION # 103

A company is developing an ML model to predict customer satisfaction. The company needs to use survey feedback and the past satisfaction level of customers to predict the future satisfaction level of customers.

The dataset includes a column named Feedback that contains long text responses. The dataset also includes a column named Satisfaction Level that contains three distinct values for past customer satisfaction: High, Medium, and Low. The company must apply encoding methods to transform the data in each column.

Which solution will meet these requirements?

- A. Apply one-hot encoding to the Feedback column. Apply ordinal encoding to the Satisfaction Level column.
- B. Apply one-hot encoding to the Feedback column and the Satisfaction Level column.
- C. Apply label encoding to the Feedback column. Apply binary encoding to the Satisfaction Level column.
- **D. Apply tokenization to the Feedback column. Apply ordinal encoding to the Satisfaction Level column.**

Answer: D

Explanation:

Different data types require different encoding strategies. The Feedback column contains long, unstructured text responses. According to AWS ML documentation, text data must first be converted into tokens before it can be vectorized using techniques such as embeddings or bag-of-words. Tokenization is the correct preprocessing step for textual features.

The Satisfaction Level column is categorical but has a natural ordering (Low < Medium < High). AWS best practices recommend ordinal encoding for such ordered categorical variables because it preserves the inherent ranking information.

Option A is incorrect because one-hot encoding is not suitable for free-form text and would create an unmanageable number of features. Option B has the same issue for the Feedback column. Option C incorrectly applies label encoding to text and binary encoding to a three-class ordinal variable.

Therefore, tokenization for text data and ordinal encoding for satisfaction levels is the correct solution.

