MLA-C01 Latest Exam Discount & Official MLA-C01 Practice Test

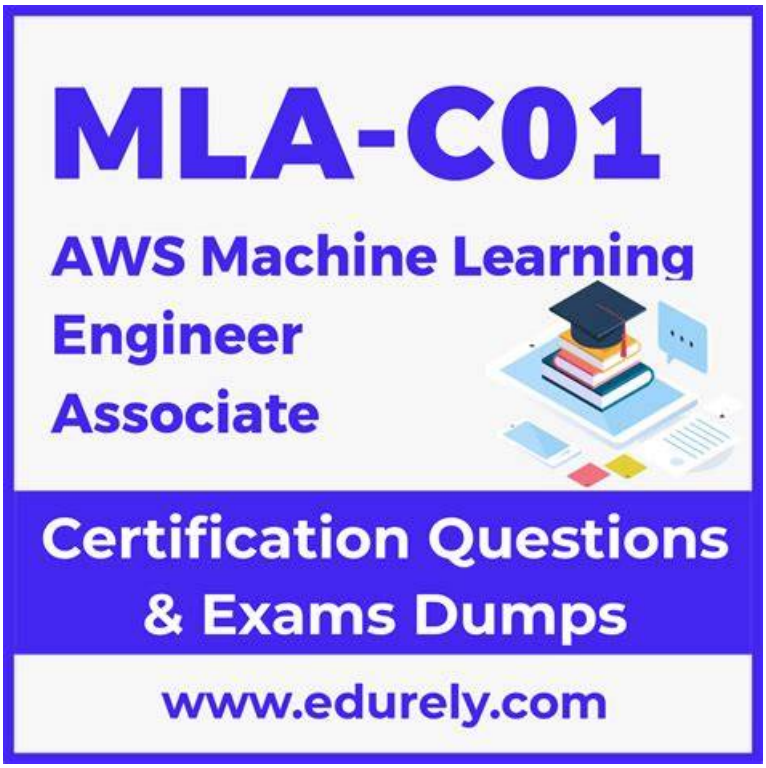

What's more, part of that ExamBoosts MLA-C01 dumps now are free: https://drive.google.com/open?id=159LggdIId7e0jxyL2fCdQoHWkGVUQ9kX

You should keep in mind to pass the MLA-C01 certification exam is not an easy task. It is a challenging job. If you want to pass the MLA-C01 exam then you have to put in some extra effort, time, and investment then you will be confident to pass the AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam. With the complete and comprehensive MLA-C01 exam dumps preparation you can pass the AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam with good scores. The ExamBoosts MLA-C01 Questions can be helpful in this regard. You must try this.

Amazon MLA-C01 Exam Overview:

Certification Vendor:	Amazon Web Services (AWS)
Exam Name:	AWS Certified Machine Learning Engineer – Associate (MLA-C01)
Exam Number:	MLA-C01
Passing Score:	720/1000
Exam Format:	Multiple choice, Multiple response, Ordering, Matching
Available Languages:	English, Japanese, Korean, Simplified Chinese
Related Certifications:	AWS Certified Data Engineer – Associate AWS Certified Solutions Architect – Associate AWS Certified DevOps Engineer – Professional AWS Certified AI Practitioner
Exam Duration:	130 minutes
Exam Price:	USD 150
Certificate Validity Period:	3 years
Real Exam Qty:	65 scored questions + 15 unscored questions
Recommended Training:	Amazon SageMaker Documentation AWS Skill Builder - ML Engineer Associate Exam Prep
Exam Registration:	AWS Certification Registration

Sample Questions:	Amazon MLA-C01 Sample Questions
Exam Way:	Online proctored or test center exam
Pre Condition:	Recommended: ~1 year experience with Amazon SageMaker and AWS-based ML or data engineering roles
Official Syllabus URL:	https://aws.amazon.com/certification/certified-machine-learning-engineer-associate/

Official MLA-C01 Practice Test | Exam MLA-C01 Prep

We would like to provide our customers with different kinds of MLA-C01 practice guide to learn, and help them accumulate knowledge and enhance their ability. Besides, we guarantee that the MLA-C01 exam questions of all our users can be answered by professional personal in the shortest time with our MLA-C01 Study Dumps. One more to mention, we can help you make full use of your sporadic time to absorb knowledge and information.

Amazon MLA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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 2	• 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	• 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.
Topic 4	• 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.

Amazon AWS Certified Machine Learning Engineer - Associate Sample Questions (Q12-Q17):

NEW QUESTION # 12
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 Autopilot 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 SageMaker JumpStart.
- C. Use SageMaker Studio 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 a custom API endpoint.

Answer: B

Explanation:
SageMaker JumpStart provides access to pre-trained models, including large language models (LLMs), which can be easily deployed and fine-tuned with a low-code/no-code (LCNC) approach. Using SageMaker Autopilot with JumpStart simplifies the fine-tuning process by automating model optimization and reducing the need for extensive coding, making it the ideal solution for this requirement.

NEW QUESTION # 13
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

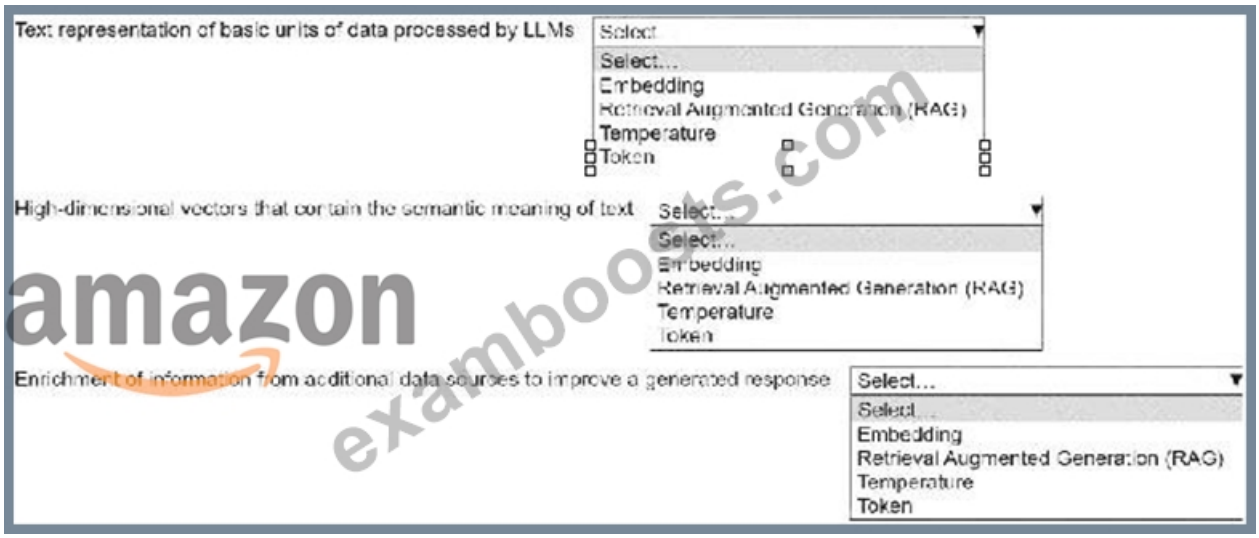

Answer:

Explanation:

Explanation:

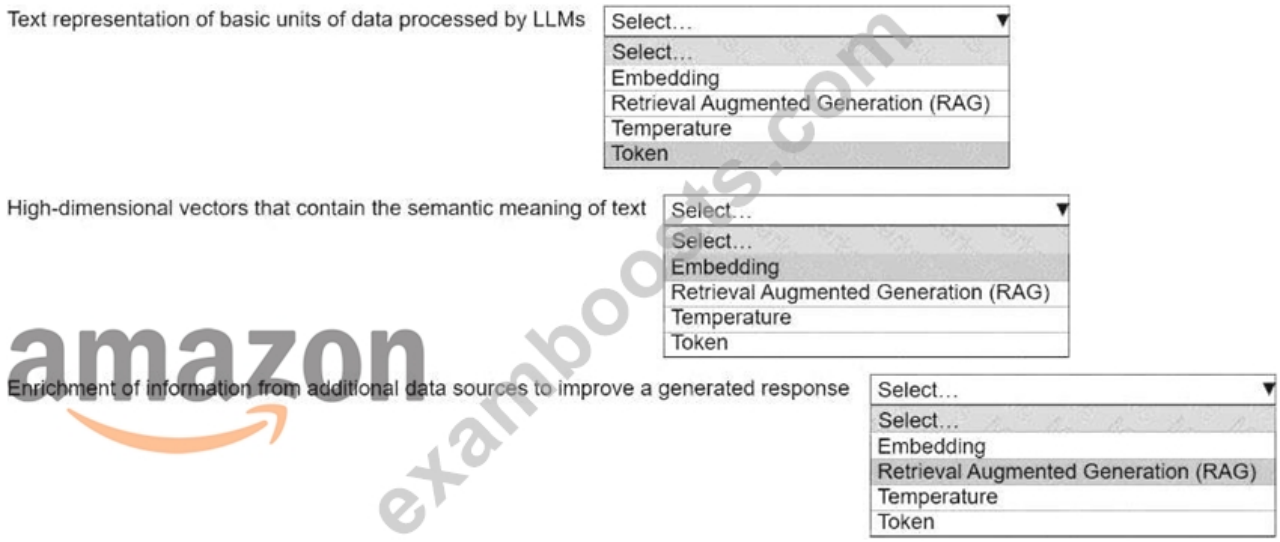

Text representation of basic units of data processed by LLMs: Token
High-dimensional vectors that contain the semantic meaning of text: Embedding Enrichment of information from additional data sources to improve a generated response: Retrieval Augmented Generation (RAG) Comprehensive Detailed Explanation Token:
Description: A token represents the smallest unit of text (e.g., a word or part of a word) that an LLM processes. For example, " running " might be split into two tokens: " run " and " ing. " Why? Tokens are the fundamental building blocks for LLM input and output processing, ensuring that the model can understand and generate text efficiently.
Embedding:
Description: High-dimensional vectors that encode the semantic meaning of text. These vectors are representations of words, sentences, or even paragraphs in a way that reflects their relationships and meaning.
Why? Embeddings are essential for enabling similarity search, clustering, or any task requiring semantic understanding. They allow the model to " understand " text contextually.
Retrieval Augmented Generation (RAG):
Description: A technique where information is enriched or retrieved from external data sources (e.g., knowledge bases or document stores) to improve the accuracy and relevance of a model ' s generated responses.
Why? RAG enhances the generative capabilities of LLMs by grounding their responses in factual and up-to- date information, reducing hallucinations in generated text.
By matching these terms to their respective descriptions, the ML engineer can effectively leverage these concepts to build robust and contextually aware generative AI applications on Amazon Bedrock.

NEW QUESTION # 14
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 Redshift Spectrum to automatically detect the anomalies. Use Amazon QuickSight to visualize the result.
- B. Use AWS Batch to automatically detect the anomalies. Use Amazon QuickSight to visualize the result.
- C. Use Amazon SageMaker Data Wrangler to automatically detect the anomalies and to visualize the result.
- D. Use Amazon Athena to automatically detect the anomalies and to visualize the result.

Answer: C

Explanation:
Amazon SageMaker Data Wrangler is a comprehensive tool that streamlines the process of data preparation and offers built-in capabilities for anomaly detection and visualization.
Key Features of SageMaker Data Wrangler:
* Data Importation: Connects seamlessly to various data sources, including Amazon S3 and on- premises databases, facilitating the aggregation of transaction logs, customer profiles, and MySQL tables.
* Anomaly Detection: Provides built-in analyses to detect anomalies in time series data, enabling the identification of outliers that may

indicate fraudulent activities.
* Visualization: Offers a suite of visualization tools, such as histograms and scatter plots, to help understand data distributions and relationships, which are crucial for feature engineering and model development.
Implementation Steps:
* Data Aggregation:
* Import data from Amazon S3 and on-premises MySQL databases into SageMaker Data Wrangler.
* Utilize Data Wrangler's data flow interface to combine and preprocess datasets, ensuring a unified dataset for analysis.
* Anomaly Detection:
* Apply the anomaly detection analysis feature to identify outliers in the dataset.
* Configure parameters such as the anomaly threshold to fine-tune the detection sensitivity.
* Visualization:
* Use built-in visualization tools to create charts and graphs that depict data distributions and highlight anomalies.
* Interpret these visualizations to gain insights into potential fraud patterns and feature interdependencies.
Advantages of Using SageMaker Data Wrangler:
* Integrated Workflow: Combines data preparation, anomaly detection, and visualization within a single interface, streamlining the ML development process.
* Operational Efficiency: Reduces the need for multiple tools and complex integrations, thereby minimizing operational overhead.
* Scalability: Handles large datasets efficiently, making it suitable for extensive transaction logs and customer profiles.
By leveraging SageMaker Data Wrangler, the ML engineer can effectively detect anomalies and visualize results, facilitating the development of a robust fraud detection model.
References:
* Analyze and Visualize - Amazon SageMaker
* Transform Data - Amazon SageMaker

NEW QUESTION # 15
A company wants to deploy an Amazon SageMaker AI model that can queue requests. The model needs to handle payloads of up to 1 GB that take up to 1 hour to process. The model must return an inference for each request. The model also must scale down when no requests are available to process.
Which inference option will meet these requirements?

- A. Batch transform
- B. Asynchronous inference
- C. Real-time inference
- D. Serverless inference

Answer: B

Explanation:
Amazon SageMaker Asynchronous Inference is specifically designed for long-running inference requests and large payloads. It supports payload sizes up to 1 GB and processing times of up to 1 hour, while automatically queuing requests.
Asynchronous inference stores results in Amazon S3 and allows clients to retrieve inference outputs after processing completes. It also supports auto scaling down to zero when there are no incoming requests, reducing cost.
Batch transform is intended for offline, bulk inference and does not return per-request results in an asynchronous request-response pattern. Serverless and real-time inference have strict payload size and timeout limits that do not support 1-hour processing.
Therefore, asynchronous inference is the only SageMaker inference option that meets all stated requirements.

NEW QUESTION # 16
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. Create a SageMaker processing job by calling the SageMaker Python SDK.
- B. Create an Apache Spark job that uses a custom processing script on Amazon EMR.
- C. Download the file to a local workstation. Perform one-hot encoding by using a custom Python script.
- D. Create a data flow in SageMaker Data Wrangler. Configure a transform step.

Answer: D

Explanation:

SageMaker Data Wrangler provides a no-code/low-code interface for preparing and transforming data, including dropping unnecessary columns. By creating a data flow and configuring a transform step, the ML engineer can easily remove correlated or unneeded columns from the Parquet file with minimal effort. This approach avoids the need for custom coding or managing additional infrastructure.

NEW QUESTION # 17

.....