

MLA-C01 Reliable Exam Simulations, MLA-C01 PDF Dumps Files



Amazon MLA-C01 AWS Certified Machine Learning Engineer - Associate

**Questions & Answers PDF
(Demo Version – Limited Content)**

For More Information – Visit link below:

<https://p2pexam.com/>

Visit us at: <https://p2pexam.com/mla-c01>

What's more, part of that PDFBraindumps MLA-C01 dumps now are free: <https://drive.google.com/open?id=16n7beNsK4w48VgAiWXs3egV6mHoKONPa>

Nowadays, there are more and more people realize the importance of MLA-C01, because more and more enterprise more and more attention it. If someone pass the MLA-C01 exam and own relevant certificates that mean he had good grasp of this field of knowledge, that is to say, he will be popular and valued by more enterprise. In order to help most candidates who want to Pass MLA-C01 Exam, so we compiled such a study materials to make exam simply. Our MLA-C01 guide torrent has gone through strict analysis and summary according to the past exam papers and the popular trend in the industry and are revised and updated according to the change of the syllabus and the latest development conditions in the theory and the practice.

PDFBraindumps is a trusted platform that has been helping AWS Certified Machine Learning Engineer - Associate MLA-C01 candidates for many years. Over this long time period, countless candidates have passed their AWS Certified Machine Learning Engineer - Associate MLA-C01 Exam and they all got help from AWS Certified Machine Learning Engineer - Associate practice questions and easily pass the final exam.

>> **MLA-C01 Reliable Exam Simulations** <<

Amazon MLA-C01 PDF Dumps Files - Valid MLA-C01 Test Sims

Our MLA-C01 study guide has three formats which can meet your different needs, PDF version, software version and online version. If you choose the PDF version, you can download our MLA-C01 study material and print it for studying everywhere. If a new version comes out, we will send you a new link to your E-mail box and you can download it again. With our software version of

MLA-C01 Exam Material, you can practice in an environment just like the real examination. And our APP version of MLA-C01 practice guide can be available with all kinds of electronic devices.

Amazon AWS Certified Machine Learning Engineer - Associate Sample Questions (Q72-Q77):

NEW QUESTION # 72

An advertising company uses AWS Lake Formation to manage a data lake. The data lake contains structured data and unstructured data. The company's ML engineers are assigned to specific advertisement campaigns.

The ML engineers must interact with the data through Amazon Athena and by browsing the data directly in an Amazon S3 bucket.

The ML engineers must have access to only the resources that are specific to their assigned advertisement campaigns.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use Lake Formation to authorize AWS Glue to access the S3 bucket. Configure Lake Formation tags to map ML engineers to their campaigns.
- B. Configure IAM policies on an AWS Glue Data Catalog to restrict access to Athena based on the ML engineers' campaigns.
- C. Store users and campaign information in an Amazon DynamoDB table. Configure DynamoDB Streams to invoke an AWS Lambda function to update S3 bucket policies.
- D. Configure S3 bucket policies to restrict access to the S3 bucket based on the ML engineers' campaigns.

Answer: A

NEW QUESTION # 73

A company runs an Amazon SageMaker AI domain in a public subnet of a newly created VPC. The network is configured properly, and ML engineers can access the SageMaker AI 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 shadow variant for the domain. Configure SageMaker Inference Recommender to send traffic from the specific IP address to the shadow endpoint.
- B. Create a network ACL inbound rule to deny traffic from the specific IP address. Assign the rule to the default network ACL for the subnet where the domain is located.
- C. Create a VPC route table to deny inbound traffic from the specific IP address. Assign the route table to the domain.
- D. Create a security group inbound rule to deny traffic from the specific IP address. Assign the security group to the domain.

Answer: B

Explanation:

In AWS networking, security groups are stateful and allow-only, meaning they cannot explicitly deny traffic.

As a result, Option A is invalid. Network ACLs (NACLs), on the other hand, are stateless and support both allow and deny rules, making them the correct mechanism for blocking traffic from specific IP addresses.

Because the SageMaker AI domain is deployed in a public subnet, inbound traffic reaches the subnet before it reaches the resource. AWS documentation states that NACLs are evaluated at the subnet level and are ideal for implementing IP-based blocking rules. Route tables control routing paths, not traffic filtering, so Option D is incorrect. Option C is unrelated to network security and does not block traffic.

AWS best practices clearly recommend using network ACL deny rules when an explicit block is required for a specific IP address at the subnet boundary.

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

NEW QUESTION # 74

A company wants to improve its customer retention ML model. The current model has 85% accuracy and a new model shows 87% accuracy in testing. The company wants to validate the new model's performance in production.

Which solution will meet these requirements?

- A. Run both models in parallel for 4 weeks. Analyze offline predictions weekly by using historical customer data analysis.
- B. Deploy the new model for 4 weeks across all production traffic. Monitor performance metrics and validate improvements.

- C. Run A/B testing on both models for 4 weeks. Route 20% of traffic to the new model. Monitor customer retention rates across both variants.
- D. Implement alternating deployments for 4 weeks between the current model and the new model. Track performance metrics for comparison.

Answer: C

Explanation:

AWS ML best practices recommend A/B testing to validate model improvements in production while minimizing risk. By routing a controlled portion of live traffic (for example, 20%) to the new model and keeping the majority of traffic on the existing model, the company can directly compare real-world performance using the same data distribution.

This approach allows statistically meaningful comparison of business metrics such as customer retention, rather than relying solely on offline accuracy. It also limits potential negative impact if the new model underperforms in production.

Deploying the new model to 100% of traffic (Option A) introduces unnecessary risk. Offline analysis (Option C) does not reflect live user behavior. Alternating deployments (Option D) introduces confounding factors such as time-based effects.

Therefore, A/B testing is the correct solution.

NEW QUESTION # 75

An ML engineer has deployed an Amazon SageMaker model to a serverless endpoint in production. The model is invoked by the InvokeEndpoint API operation.

The model's latency in production is higher than the baseline latency in the test environment. The ML engineer thinks that the increase in latency is because of model startup time.

What should the ML engineer do to confirm or deny this hypothesis?

- A. Schedule a SageMaker Model Monitor job. Observe metrics about model quality.
- B. Enable Amazon CloudWatch metrics. Observe the ModelSetupTime metric in the SageMaker namespace.
- C. Schedule a SageMaker Model Monitor job with Amazon CloudWatch metrics enabled.
- D. Enable Amazon CloudWatch metrics. Observe the ModelLoadingWaitTime metric in the SageMaker namespace.

Answer: D

NEW QUESTION # 76

An ML engineer is using Amazon SageMaker Canvas to build a custom ML model from an imported dataset.

The model must make continuous numeric predictions based on 10 years of data.

Which metric should the ML engineer use to evaluate the model's performance?

- A. Accuracy
- B. Root Mean Square Error (RMSE)
- C. InferenceLatency
- D. Area Under the ROC Curve (AUC)

Answer: B

Explanation:

This is a regression problem, where the target variable is continuous and numeric. AWS documentation clearly states that classification metrics such as accuracy and AUC are not appropriate for regression models.

Root Mean Square Error (RMSE) measures the square root of the average squared differences between predicted and actual values. RMSE penalizes larger errors more heavily, making it especially useful when large prediction errors are costly or undesirable.

SageMaker Canvas automatically selects regression metrics such as RMSE and MAE when building regression models. RMSE is widely used for time-based and numeric prediction problems, especially when evaluating long historical datasets.

Inference latency measures system performance, not model accuracy.

Therefore, Option D is the correct and AWS-verified answer.

NEW QUESTION # 77

.....

The content of our MLA-C01 practice engine is chosen so carefully that all the questions for the MLA-C01 exam are contained. And our MLA-C01 study materials have three formats which help you to read, test and study anytime, anywhere. This means with

- [illegible]

www.stes.tyc.edu.tw, Disposable vapes

BONUS!!! Download part of PDFBrain dumps MLA-C01 dumps for free: <https://drive.google.com/open?id=16n7beNsK4w48VgAiWXs3egV6mHoKONPa>