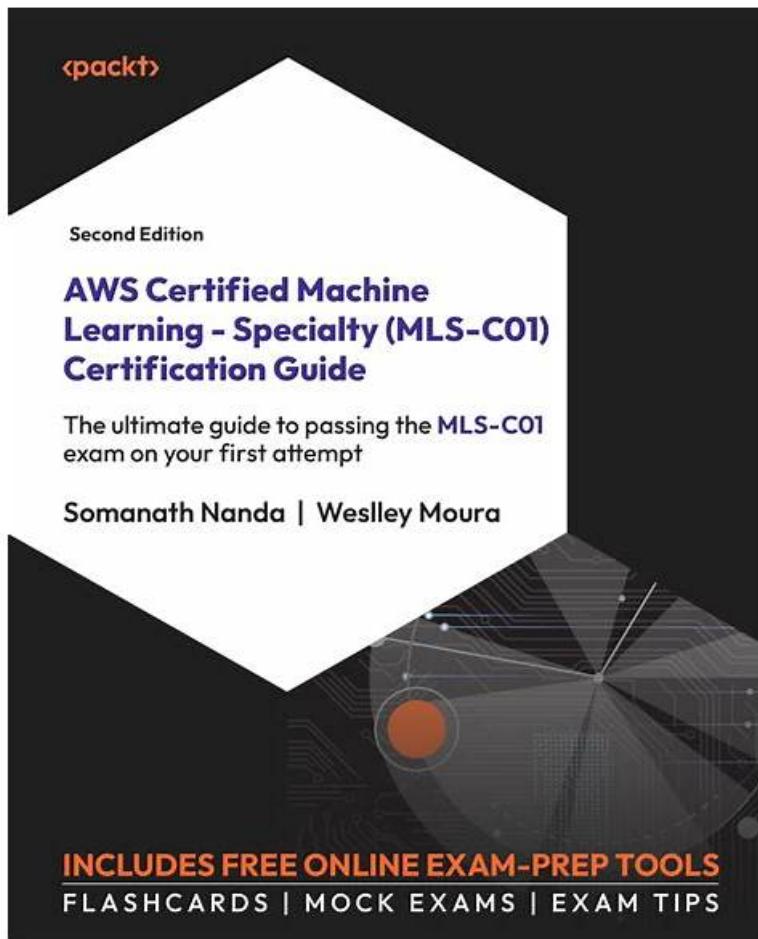


Free PDF Quiz MLS-C01 - AWS Certified Machine Learning - Specialty–Professional Dump Collection



2026 Latest BraindumpsIT MLS-C01 PDF Dumps and MLS-C01 Exam Engine Free Share: https://drive.google.com/open?id=1JWkbmN-F5g_PUtMJV0a6IdieM345G951

Among global market, AWS Certified Specialty guide question is not taking up such a large share with high reputation for nothing. And we are the leading practice materials in this dynamic market. To facilitate your review process, all questions and answers of our MLS-C01 test question is closely related with the real exam by our experts who constantly keep the updating of products to ensure the accuracy of questions, so all MLS-C01 guide question is 100 percent assured. We make MLS-C01 exam prep from exam candidate perspective, and offer high quality practice materials with reasonable prices but various benefits. The more times you choose us, the more discounts you may get. To make your whole experience more comfortable, we also provide considerate whole package services once you make decisions of our MLS-C01 Test Question. If you have any questions related to our MLS-C01 exam prep, pose them and our employees will help you as soon as possible.

Understanding functional and technical aspects of AWS Certified Machine Learning - Specialty Data Engineering

The following will be discussed in **AMAZON MLS-C01 exam dumps**:

- Create data repositories for machine learning
- Identify and implement a data-ingestion solution
- Identify and implement a data-transformation solution

Amazon MLS-C01 exam covers a wide range of topics, including data preparation, feature engineering, model training, evaluation and deployment, and machine learning implementation and operations. MLS-C01 Exam also covers AWS services such as Amazon SageMaker, AWS Deep Learning AMIs, Amazon S3, Amazon EMR, and Amazon Redshift. MLS-C01 exam is designed to assess

the candidate's ability to use these services to build, train, and deploy ML models.

>> Dump MLS-C01 Collection <<

Free PDF Quiz Amazon - MLS-C01 –High Pass-Rate Dump Collection

Just install the AWS Certified Machine Learning - Specialty (MLS-C01) PDF dumps file on your desktop computer, laptop, tab, or even on your smartphone and start AWS Certified Machine Learning - Specialty (MLS-C01) exam preparation anytime and anywhere. Whereas the other two AWS Certified Machine Learning - Specialty (MLS-C01) exam questions formats are concerned both are the easy-to-use and compatible Mock MLS-C01 Exam that will give you a real-time environment for quick Amazon Exams preparation. Now choose the right Amazon MLS-C01 exam questions format and start this career advancement journey.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q189-Q194):

NEW QUESTION # 189

A company has an ecommerce website with a product recommendation engine built in TensorFlow. The recommendation engine endpoint is hosted by Amazon SageMaker. Three compute-optimized instances support the expected peak load of the website. Response times on the product recommendation page are increasing at the beginning of each month. Some users are encountering errors. The website receives the majority of its traffic between 8 AM and 6 PM on weekdays in a single time zone.

Which of the following options are the MOST effective in solving the issue while keeping costs to a minimum? (Choose two.)

- A. Reconfigure the endpoint to use burstable instances.
- B. Configure the endpoint to use Amazon Elastic Inference (EI) accelerators.
- C. Configure the endpoint to automatically scale with the Invocations Per Instance metric.
- D. Create a new endpoint configuration with two production variants.
- E. Deploy a second instance pool to support a blue/green deployment of models.

Answer: B,C

Explanation:

The solution A and C are the most effective in solving the issue while keeping costs to a minimum. The solution A and C involve the following steps:

Configure the endpoint to use Amazon Elastic Inference (EI) accelerators. This will enable the company to reduce the cost and latency of running TensorFlow inference on SageMaker. Amazon EI provides GPU- powered acceleration for deep learning models without requiring the use of GPU instances. Amazon EI can attach to any SageMaker instance type and provide the right amount of acceleration based on the workload1.

Configure the endpoint to automatically scale with the Invocations Per Instance metric. This will enable the company to adjust the number of instances based on the demand and traffic patterns of the website. The Invocations Per Instance metric measures the average number of requests that each instance processes over a period of time. By using this metric, the company can scale out the endpoint when the load increases and scale in when the load decreases. This can improve the response time and availability of the product recommendation engine2.

The other options are not suitable because:

Option B: Creating a new endpoint configuration with two production variants will not solve the issue of increasing response time and errors. Production variants are used to split the traffic between different models or versions of the same model. They can be useful for testing, updating, or A/B testing models. However, they do not provide any scaling or acceleration benefits for the inference workload3.

Option D: Deploying a second instance pool to support a blue/green deployment of models will not solve the issue of increasing response time and errors. Blue/green deployment is a technique for updating models without downtime or disruption. It involves creating a new endpoint configuration with a different instance pool and model version, and then shifting the traffic from the old endpoint to the new endpoint gradually. However, this technique does not provide any scaling or acceleration benefits for the inference workload4.

Option E: Reconfiguring the endpoint to use burstable instances will not solve the issue of increasing response time and errors.

Burstable instances are instances that provide a baseline level of CPU performance with the ability to burst above the baseline when needed. They can be useful for workloads that have moderate CPU utilization and occasional spikes. However, they are not suitable for workloads that have high and consistent CPU utilization, such as the product recommendation engine. Moreover, burstable instances may incur additional charges when they exceed their CPU credits5.

1: Amazon Elastic Inference

2: How to Scale Amazon SageMaker Endpoints

- 3: Deploying Models to Amazon SageMaker Hosting Services
- 4: Updating Models in Amazon SageMaker Hosting Services
- 5: Burstable Performance Instances

NEW QUESTION # 190

A data scientist is using the Amazon SageMaker Neural Topic Model (NTM) algorithm to build a model that recommends tags from blog posts. The raw blog post data is stored in an Amazon S3 bucket in JSON format. During model evaluation, the data scientist discovered that the model recommends certain stopwords such as "a," "an," and "the" as tags to certain blog posts, along with a few rare words that are present only in certain blog entries. After a few iterations of tag review with the content team, the data scientist notices that the rare words are unusual but feasible. The data scientist also must ensure that the tag recommendations of the generated model do not include the stopwords.

What should the data scientist do to meet these requirements?

- A. Use the SageMaker built-in Object Detection algorithm instead of the NTM algorithm for the training job to process the blog post data.
- B. Remove the stop words from the blog post data by using the Count Vectorizer function in the scikit-learn library. Replace the blog post data in the S3 bucket with the results of the vectorizer.
- C. Run the SageMaker built-in principal component analysis (PCA) algorithm with the blog post data from the S3 bucket as the data source. Replace the blog post data in the S3 bucket with the results of the training job.
- D. Use the Amazon Comprehend entity recognition API operations. Remove the detected words from the blog post data. Replace the blog post data source in the S3 bucket.

Answer: B

Explanation:

The data scientist should remove the stop words from the blog post data by using the Count Vectorizer function in the scikit-learn library, and replace the blog post data in the S3 bucket with the results of the vectorizer. This is because:

The Count Vectorizer function is a tool that can convert a collection of text documents to a matrix of token counts 1. It also enables the pre-processing of text data prior to generating the vector representation, such as removing accents, converting to lowercase, and filtering out stop words 1. By using this function, the data scientist can remove the stop words such as "a," "an," and "the" from the blog post data, and obtain a numerical representation of the text that can be used as input for the NTM algorithm.

The NTM algorithm is a neural network-based topic modeling technique that can learn latent topics from a corpus of documents 2. It can be used to recommend tags from blog posts by finding the most probable topics for each document, and ranking the words associated with each topic 3. However, the NTM algorithm does not perform any text pre-processing by itself, so it relies on the quality of the input data. Therefore, the data scientist should replace the blog post data in the S3 bucket with the results of the vectorizer, to ensure that the NTM algorithm does not include the stop words in the tag recommendations.

The other options are not suitable for the following reasons:

Option A is not relevant because the Amazon Comprehend entity recognition API operations are used to detect and extract named entities from text, such as people, places, organizations, dates, etc4. This is not the same as removing stop words, which are common words that do not carry much meaning or information. Moreover, removing the detected entities from the blog post data may reduce the quality and diversity of the tag recommendations, as some entities may be relevant and useful as tags.

Option B is not optimal because the SageMaker built-in principal component analysis (PCA) algorithm is used to reduce the dimensionality of a dataset by finding the most important features that capture the maximum amount of variance in the data 5. This is not the same as removing stop words, which are words that have low variance and high frequency in the data. Moreover, replacing the blog post data in the S3 bucket with the results of the PCA algorithm may not be compatible with the input format expected by the NTM algorithm, which requires a bag-of-words representation of the text 2.

Option C is not suitable because the SageMaker built-in Object Detection algorithm is used to detect and localize objects in images 6. This is not related to the task of recommending tags from blog posts, which are text documents. Moreover, using the Object Detection algorithm instead of the NTM algorithm would require a different type of input data (images instead of text), and a different type of output data (bounding boxes and labels instead of topics and words).

References:

[Neural Topic Model \(NTM\) Algorithm](#)

[Introduction to the Amazon SageMaker Neural Topic Model](#)

[Amazon Comprehend - Entity Recognition](#)

[sklearn.feature_extraction.text.CountVectorizer](#)

[Principal Component Analysis \(PCA\) Algorithm](#)

[Object Detection Algorithm](#)

NEW QUESTION # 191

A Machine Learning Specialist is working for a credit card processing company and receives an unbalanced dataset containing credit card transactions. It contains 99,000 valid transactions and 1,000 fraudulent transactions. The Specialist is asked to score a model that was run against the dataset. The Specialist has been advised that identifying valid transactions is equally as important as identifying fraudulent transactions. What metric is BEST suited to score the model?

- A. Recall
- B. Precision
- C. Root Mean Square Error (RMSE)
- D. **Area Under the ROC Curve (AUC)**

Answer: D

Explanation:

Area Under the ROC Curve (AUC) is a metric that is best suited to score the model for the given scenario.

AUC is a measure of the performance of a binary classifier, such as a model that predicts whether a credit card transaction is valid or fraudulent. AUC is calculated based on the Receiver Operating Characteristic (ROC) curve, which is a plot that shows the trade-off between the true positive rate (TPR) and the false positive rate (FPR) of the classifier as the decision threshold is varied. The TPR, also known as recall or sensitivity, is the proportion of actual positive cases (fraudulent transactions) that are correctly predicted as positive by the classifier. The FPR, also known as the fall-out, is the proportion of actual negative cases (valid transactions) that are incorrectly predicted as positive by the classifier. The ROC curve illustrates how well the classifier can distinguish between the two classes, regardless of the class distribution or the error costs. A perfect classifier would have a TPR of 1 and an FPR of 0 for all thresholds, resulting in a ROC curve that goes from the bottom left to the top left and then to the top right of the plot. A random classifier would have a TPR and an FPR that are equal for all thresholds, resulting in a ROC curve that goes from the bottom left to the top right of the plot along the diagonal line. AUC is the area under the ROC curve, and it ranges from 0 to 1. A higher AUC indicates a better classifier, as it means that the classifier has a higher TPR and a lower FPR for all thresholds. AUC is a useful metric for imbalanced classification problems, such as the credit card transaction dataset, because it is insensitive to the class imbalance and the error costs. AUC can capture the overall performance of the classifier across all possible scenarios, and it can be used to compare different classifiers based on their ROC curves.

The other options are not as suitable as AUC for the given scenario for the following reasons:

* Precision: Precision is the proportion of predicted positive cases (fraudulent transactions) that are actually positive. Precision is a useful metric when the cost of a false positive is high, such as in spam detection or medical diagnosis. However, precision is not a good metric for imbalanced classification problems, because it can be misleadingly high when the positive class is rare. For example, a classifier that predicts all transactions as valid would have a precision of 0, but a very high accuracy of 99%.

Precision is also dependent on the decision threshold and the error costs, which may vary for different scenarios.

* Recall: Recall is the same as the TPR, and it is the proportion of actual positive cases (fraudulent transactions) that are correctly predicted as positive by the classifier. Recall is a useful metric when the cost of a false negative is high, such as in fraud detection or cancer diagnosis. However, recall is not a good metric for imbalanced classification problems, because it can be misleadingly low when the positive class is rare. For example, a classifier that predicts all transactions as fraudulent would have a recall of 1, but a very low accuracy of 1%. Recall is also dependent on the decision threshold and the error costs, which may vary for different scenarios.

* Root Mean Square Error (RMSE): RMSE is a metric that measures the average difference between the predicted and the actual values. RMSE is a useful metric for regression problems, where the goal is to predict a continuous value, such as the price of a house or the temperature of a city. However, RMSE is not a good metric for classification problems, where the goal is to predict a discrete value, such as the class label of a transaction. RMSE is not meaningful for classification problems, because it does not capture the accuracy or the error costs of the predictions.

ROC Curve and AUC

How and When to Use ROC Curves and Precision-Recall Curves for Classification in Python Precision-Recall Root Mean Squared Error

NEW QUESTION # 192

A Machine Learning Specialist deployed a model that provides product recommendations on a company's website. Initially, the model was performing very well and resulted in customers buying more products on average. However, within the past few months, the Specialist has noticed that the effect of product recommendations has diminished and customers are starting to return to their original habits of spending less. The Specialist is unsure of what happened, as the model has not changed from its initial deployment over a year ago. Which method should the Specialist try to improve model performance?

- A. The model should be periodically retrained from scratch using the original data while adding a regularization term to handle product inventory changes
- B. **The model should be periodically retrained using the original training data plus new data as product inventory changes**

- C. The model's hyperparameters should be periodically updated to prevent drift
- D. The model needs to be completely re-engineered because it is unable to handle product inventory changes

Answer: B

Explanation:

The problem that the Machine Learning Specialist is facing is likely due to concept drift, which is a phenomenon where the statistical properties of the target variable change over time, making the model less accurate and relevant. Concept drift can occur due to various reasons, such as changes in customer preferences, market trends, product inventory, seasonality, etc. In this case, the product recommendations model may have become outdated as the product inventory changed over time, making the recommendations less appealing to the customers. To address this issue, the model should be periodically retrained using the original training data plus new data as product inventory changes. This way, the model can learn from the latest data and adapt to the changing customer behavior and preferences. Retraining the model from scratch using the original data while adding a regularization term may not be sufficient, as it does not account for the new data. Updating the model's hyperparameters may not help either, as it does not address the underlying data distribution change. Re-engineering the model completely may not be necessary, as the model may still be valid and useful with periodic retraining.

References:

Concept Drift - Amazon SageMaker

Detecting and Handling Concept Drift - Amazon SageMaker

Machine Learning Concepts - Amazon Machine Learning

NEW QUESTION # 193

A data scientist needs to identify fraudulent user accounts for a company's ecommerce platform. The company wants the ability to determine if a newly created account is associated with a previously known fraudulent user. The data scientist is using AWS Glue to cleanse the company's application logs during ingestion.

Which strategy will allow the data scientist to identify fraudulent accounts?

- A. Execute the built-in FindDuplicates Amazon Athena query.
- B. **Create a FindMatches machine learning transform in AWS Glue.**
- C. Create an AWS Glue crawler to infer duplicate accounts in the source data.
- D. Search for duplicate accounts in the AWS Glue Data Catalog.

Answer: B

NEW QUESTION # 194

.....

BraindumpsIT aims to assist its clients in making them capable of passing the Amazon MLS-C01 certification exam with flying colors. It fulfills its mission by giving them an entirely free AWS Certified Machine Learning - Specialty (MLS-C01) demo of the dumps. Thus, this demonstration will enable them to scrutinize the quality of the Amazon MLS-C01 study material.

MLS-C01 Sure Pass: https://www.braindumpsit.com/MLS-C01_real-exam.html

- Your Investment with www.prep4sures.top Amazon MLS-C01 Exam Questions is Secured Immediately open “www.prep4sures.top” and search for ▶ MLS-C01 ▲ to obtain a free download MLS-C01 Valid Test Sims
- Quiz Amazon - MLS-C01 - AWS Certified Machine Learning - Specialty –The Best Dump Collection Open ▷ www.pdfvce.com ▲ and search for ▷ MLS-C01 to download exam materials for free MLS-C01 Reliable Exam Tips
- MLS-C01 New Exam Materials MLS-C01 Reliable Exam Tips MLS-C01 Exam Simulator Online Download ▷ MLS-C01 for free by simply searching on 「 www.prep4away.com 」 MLS-C01 Latest Mock Exam
- Pass Guaranteed Quiz 2026 MLS-C01: The Best Dump AWS Certified Machine Learning - Specialty Collection Search for MLS-C01 and download it for free immediately on www.pdfvce.com MLS-C01 Reliable Exam Tips
- Quiz 2026 Reliable Amazon Dump MLS-C01 Collection Search for 《 MLS-C01 》 and download it for free on 《 www.prepawaypdf.com 》 website Exam Dumps MLS-C01 Collection
- MLS-C01 Test Dumps Demo Exam Dumps MLS-C01 Collection MLS-C01 Reliable Test Voucher Search for ▷ MLS-C01 ▲ and download exam materials for free through ▷ www.pdfvce.com New MLS-C01 Test Price
- Pass Guaranteed Quiz 2026 MLS-C01: The Best Dump AWS Certified Machine Learning - Specialty Collection Open www.vce4dumps.com and search for “ MLS-C01 ” to download exam materials for free Practice MLS-C01 Exam

BTW, DOWNLOAD part of BraindumpsIT MLS-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1JWkbmNF5g_PUTMJV0a61dieM345G951