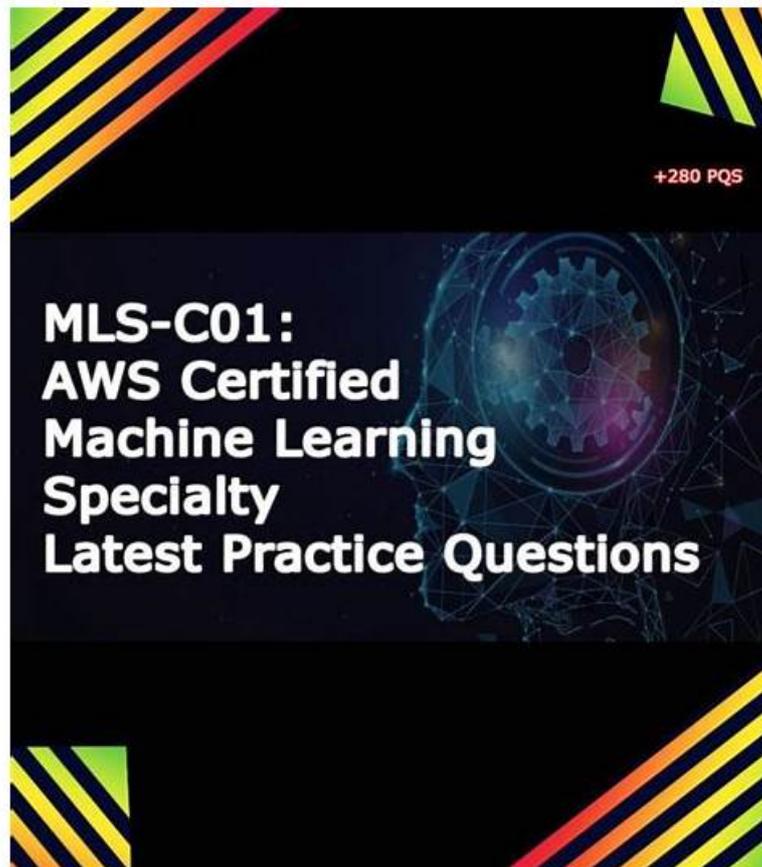


Amazon Latest AWS-Certified-Machine-Learning-Specialty Practice Questions: AWS Certified Machine Learning - Specialty - PrepAwayPDF Instant Download



BTW, DOWNLOAD part of PrepAwayPDF AWS-Certified-Machine-Learning-Specialty dumps from Cloud Storage: <https://drive.google.com/open?id=1B0k80cE1B15a5sgY3L7DFIKLfxuwUauy>

PrepAwayPDF provides you with free demos of its AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty exam product. You can try a free demo to eliminate any confusion regarding the authenticity of our AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty PDF and practice tests (web-based & desktop software). It is also our policy to facilitate you with AWS-Certified-Machine-Learning-Specialty free actual dumps updates in case of new AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty test changes within three months of your shopping. Contact us any time, if you need any guidance about our Amazon AWS-Certified-Machine-Learning-Specialty exam product. There is only one way to get all these amazing AWS-Certified-Machine-Learning-Specialty exam dumps offers and that is purchasing our product today.

The AWS Certified Machine Learning - Specialty certification is a valuable credential for professionals who want to advance their careers in the field of ML. Certified individuals have a competitive edge in the job market, as they demonstrate their ability to design and implement cutting-edge ML solutions on the AWS platform. Moreover, the certification is recognized by industry leaders and organizations, which further enhances its value and credibility. Overall, the AWS Certified Machine Learning - Specialty exam is a challenging but rewarding certification that can help individuals prove their expertise in the field of ML and advance their careers.

>> **Latest AWS-Certified-Machine-Learning-Specialty Practice Questions** <<

Test AWS-Certified-Machine-Learning-Specialty Cram | AWS-Certified-Machine-Learning-Specialty Real Question

This format enables you to assess your AWS-Certified-Machine-Learning-Specialty test preparation with a Amazon AWS-

Certified-Machine-Learning-Specialty certification exam. You can also customize your time and the kinds of Amazon AWS-Certified-Machine-Learning-Specialty Exam Questions of the Amazon AWS-Certified-Machine-Learning-Specialty practice test. PrepAwayPDF has formulated AWS-Certified-Machine-Learning-Specialty PDF questions for the convenience of Amazon AWS-Certified-Machine-Learning-Specialty test takers.

To earn the AWS Certified Machine Learning - Specialty certification, candidates must have a strong understanding of machine learning algorithms, data preprocessing, and feature engineering. They should also have experience working with AWS services such as Amazon SageMaker, AWS Glue, and AWS Kinesis. Additionally, candidates should be familiar with deep learning frameworks such as TensorFlow, Keras, and PyTorch. AWS-Certified-Machine-Learning-Specialty Exam covers a range of topics including machine learning algorithms, data modeling and evaluation, and deployment strategies. Passing the exam demonstrates that an individual has the skills and knowledge necessary to implement machine learning solutions on AWS.

Career Opportunities

Machine Learning is no doubt one of the hottest topics within the Information Technology sector. Therefore, the Amazon AWS Certified Machine Learning – Specialty certification is simply the key to become a highly regarded certified professional in the field. Those professionals who obtain this certificate can boost their career to a higher level and get a decent salary. They can opt for different job roles, such as a Solutions Architect, a Technical Curriculum Developer, an Electrical Safety Program Manager, a Systems Development Engineer, a Software Development Manager, a Global Ergonomics Engineer, and many more. The average salary can range from \$30,000 to \$160,000 per year.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q83-Q88):

NEW QUESTION # 83

A Data Scientist needs to create a serverless ingestion and analytics solution for high-velocity, real-time streaming data. The ingestion process must buffer and convert incoming records from JSON to a query- optimized, columnar format without data loss. The output datastore must be highly available, and Analysts must be able to run SQL queries against the data and connect to existing business intelligence dashboards.

Which solution should the Data Scientist build to satisfy the requirements?

- A. Use Amazon Kinesis Data Analytics to ingest the streaming data and perform real-time SQL queries to convert the records to Apache Parquet before delivering to Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.
- B. Write each JSON record to a staging location in Amazon S3. Use the S3 Put event to trigger an AWS Lambda function that transforms the data into Apache Parquet or ORC format and inserts it into an Amazon RDS PostgreSQL database. Have the Analysts query and run dashboards from the RDS database.
- C. Write each JSON record to a staging location in Amazon S3. Use the S3 Put event to trigger an AWS Lambda function that transforms the data into Apache Parquet or ORC format and writes the data to a processed data location in Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena, and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.
- **D. Create a schema in the AWS Glue Data Catalog of the incoming data format. Use an Amazon Kinesis Data Firehose delivery stream to stream the data and transform the data to Apache Parquet or ORC format using the AWS Glue Data Catalog before delivering to Amazon S3. Have the Analysts query the data directly from Amazon S3 using Amazon Athena, and connect to BI tools using the Athena Java Database Connectivity (JDBC) connector.**

Answer: D

NEW QUESTION # 84

A machine learning (ML) specialist must develop a classification model for a financial services company. A domain expert provides the dataset, which is tabular with 10,000 rows and 1,020 features. During exploratory data analysis, the specialist finds no missing values and a small percentage of duplicate rows. There are correlation scores of > 0.9 for 200 feature pairs. The mean value of each feature is similar to its 50th percentile.

Which feature engineering strategy should the ML specialist use with Amazon SageMaker?

- A. Drop the features with low correlation scores by using a Jupyter notebook.
- **B. Apply dimensionality reduction by using the principal component analysis (PCA) algorithm.**
- C. Apply anomaly detection by using the Random Cut Forest (RCF) algorithm.
- D. Concatenate the features with high correlation scores by using a Jupyter notebook.

Answer: B

NEW QUESTION # 85

A machine learning specialist works for a fruit processing company and needs to build a system that categorizes apples into three types. The specialist has collected a dataset that contains 150 images for each type of apple and applied transfer learning on a neural network that was pretrained on ImageNet with this dataset.

The company requires at least 85% accuracy to make use of the model.

After an exhaustive grid search, the optimal hyperparameters produced the following:

68% accuracy on the training set

67% accuracy on the validation set

What can the machine learning specialist do to improve the system's accuracy?

- A. Upload the model to an Amazon SageMaker notebook instance and use the Amazon SageMaker HPO feature to optimize the model's hyperparameters.
- B. Use a neural network model with more layers that are pretrained on ImageNet and apply transfer learning to increase the variance.
- C. Train a new model using the current neural network architecture.
- **D. Add more data to the training set and retrain the model using transfer learning to reduce the bias.**

Answer: D

Explanation:

The problem described in the question is a case of underfitting, where the neural network model performs poorly on both the training and validation sets. This means that the model has not learned the features of the data well enough and has high bias. To solve this issue, the machine learning specialist should consider the following change:

* Add more data to the training set and retrain the model using transfer learning to reduce the bias:

Adding more data to the training set can help the model learn more patterns and variations in the data and improve its performance.

Transfer learning can also help the model leverage the knowledge from the pre-trained network and adapt it to the new data. This can reduce the bias and increase the accuracy of the model.

Transfer learning for TensorFlow image classification models in Amazon SageMaker Transfer learning for custom labels using a TensorFlow container and "bring your own algorithm" in Amazon SageMaker Machine Learning Concepts - AWS Training and Certification

NEW QUESTION # 86

A company is building a new version of a recommendation engine. Machine learning (ML) specialists need to keep adding new data from users to improve personalized recommendations. The ML specialists gather data from the users' interactions on the platform and from sources such as external websites and social media.

The pipeline cleans, transforms, enriches, and compresses terabytes of data daily, and this data is stored in Amazon S3. A set of Python scripts was coded to do the job and is stored in a large Amazon EC2 instance.

The whole process takes more than 20 hours to finish, with each script taking at least an hour. The company wants to move the scripts out of Amazon EC2 into a more managed solution that will eliminate the need to maintain servers.

Which approach will address all of these requirements with the LEAST development effort?

- A. Load the data into Amazon DynamoDB. Convert the scripts to an AWS Lambda function. Execute the pipeline by triggering Lambda executions. Store the results in Amazon S3.
- B. Create a set of individual AWS Lambda functions to execute each of the scripts. Build a step function by using the AWS Step Functions Data Science SDK. Store the results in Amazon S3.
- **C. Create an AWS Glue job. Convert the scripts to PySpark. Execute the pipeline. Store the results in Amazon S3.**
- D. Load the data into an Amazon Redshift cluster. Execute the pipeline by using SQL. Store the results in Amazon S3.

Answer: C

Explanation:

Explanation

The best approach to address all of the requirements with the least development effort is to create an AWS Glue job, convert the scripts to PySpark, execute the pipeline, and store the results in Amazon S3. This is because:

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load data for analytics 1. AWS Glue can run Python and Scala scripts to process data from various sources, such as Amazon S3, Amazon DynamoDB, Amazon Redshift, and more 2. AWS Glue also provides a serverless Apache Spark environment to run ETL jobs, eliminating the

need to provision and manage servers 3.

PySpark is the Python API for Apache Spark, a unified analytics engine for large-scale data processing 4. PySpark can perform various data transformations and manipulations on structured and unstructured data, such as cleaning, enriching, and compressing 5. PySpark can also leverage the distributed computing power of Spark to handle terabytes of data efficiently and scalably 6.

By creating an AWS Glue job and converting the scripts to PySpark, the company can move the scripts out of Amazon EC2 into a more managed solution that will eliminate the need to maintain servers. The company can also reduce the development effort by using the AWS Glue console, AWS SDK, or AWS CLI to create and run the job 7. Moreover, the company can use the AWS Glue Data Catalog to store and manage the metadata of the data sources and targets 8.

The other options are not as suitable as option C for the following reasons:

Option A is not optimal because loading the data into an Amazon Redshift cluster and executing the pipeline by using SQL will incur additional costs and complexity for the company. Amazon Redshift is a fully managed data warehouse service that enables fast and scalable analysis of structured data .

However, it is not designed for ETL purposes, such as cleaning, transforming, enriching, and compressing data. Moreover, using SQL to perform these tasks may not be as expressive and flexible as using Python scripts. Furthermore, the company will have to provision and configure the Amazon Redshift cluster, and load and unload the data from Amazon S3, which will increase the development effort and time.

Option B is not feasible because loading the data into Amazon DynamoDB and converting the scripts to an AWS Lambda function will not work for the company's use case. Amazon DynamoDB is a fully managed key-value and document database service that provides fast and consistent performance at any scale . However, it is not suitable for storing and processing terabytes of data daily, as it has limits on the size and throughput of each table and item . Moreover, using AWS Lambda to execute the pipeline will not be efficient or cost-effective, as Lambda has limits on the memory, CPU, and execution time of each function . Therefore, using Amazon DynamoDB and AWS Lambda will not meet the company's requirements for processing large amounts of data quickly and reliably.

Option D is not relevant because creating a set of individual AWS Lambda functions to execute each of the scripts and building a step function by using the AWS Step Functions Data Science SDK will not address the main issue of moving the scripts out of Amazon EC2. AWS Step Functions is a fully managed service that lets you coordinate multiple AWS services into serverless workflows . The AWS Step Functions Data Science SDK is an open source library that allows data scientists to easily create workflows that process and publish machine learning models using Amazon SageMaker and AWS Step Functions . However, these services and tools are not designed for ETL purposes, such as cleaning, transforming, enriching, and compressing data. Moreover, as mentioned in option B, using AWS Lambda to execute the scripts will not be efficient or cost-effective for the company's use case.

References:

What Is AWS Glue?

AWS Glue Components

AWS Glue Serverless Spark ETL

PySpark - Overview

PySpark - RDD

PySpark - SparkContext

Adding Jobs in AWS Glue

Populating the AWS Glue Data Catalog

[What Is Amazon Redshift?]

[What Is Amazon DynamoDB?]

[Service, Account, and Table Quotas in DynamoDB]

[AWS Lambda quotas]

[What Is AWS Step Functions?]

[AWS Step Functions Data Science SDK for Python]

NEW QUESTION # 87

An e-commerce company sends a weekly email newsletter to all of its customers. Management has hired a team of writers to create additional targeted content. A data scientist needs to identify five customer segments based on age, income, and location. The customers' current segmentation is unknown. The data scientist previously built an XGBoost model to predict the likelihood of a customer responding to an email based on age, income, and location.

Why does the XGBoost model NOT meet the current requirements, and how can this be fixed?

- A. The XGBoost model provides a true/false binary output. Increase the number of classes the XGBoost model predicts to five classes to predict a segment.
- B. The XGBoost model provides a true/false binary output. Apply principal component analysis (PCA) with five feature dimensions to predict a segment.
- C. The XGBoost model is a supervised machine learning algorithm. Train a k-Nearest-Neighbors (kNN) model with $K = 5$

on the same dataset to predict a segment.

- **D. The XGBoost model is a supervised machine learning algorithm. Train a k-means model with K = 5 on the same dataset to predict a segment.**

Answer: D

Explanation:

The XGBoost model is a supervised machine learning algorithm, which means it requires labeled data to learn from. The customers' current segmentation is unknown, so there is no label to train the XGBoost model on. Moreover, the XGBoost model is designed for classification or regression tasks, not for clustering.

Clustering is a type of unsupervised machine learning, which means it does not require labeled data.

Clustering algorithms try to find natural groups or clusters in the data based on their similarity or distance. A common clustering algorithm is k-means, which partitions the data into K clusters, where each data point belongs to the cluster with the nearest mean. To meet the current requirements, the data scientist should train a k-means model with K = 5 on the same dataset to predict a segment for each customer. This way, the data scientist can identify five customer segments based on age, income, and location, without needing any labels.

References:

- * What is XGBoost? - Amazon SageMaker
- * What is Clustering? - Amazon SageMaker
- * K-Means Algorithm - Amazon SageMaker

NEW QUESTION # 88

.....

Test AWS-Certified-Machine-Learning-Specialty Cram: <https://www.prepawaypdf.com/Amazon/AWS-Certified-Machine-Learning-Specialty-practice-exam-dumps.html>

- Updated Amazon AWS-Certified-Machine-Learning-Specialty Exam Questions with Accurate Answers in PDF Open **【 www.prepawaypdf.com 】** enter **>** AWS-Certified-Machine-Learning-Specialty and obtain a free download AWS-Certified-Machine-Learning-Specialty Test Tutorials
- Newest Amazon - AWS-Certified-Machine-Learning-Specialty - Latest AWS Certified Machine Learning - Specialty Practice Questions Enter **【 www.pdfvce.com 】** and search for [AWS-Certified-Machine-Learning-Specialty] to download for free Test AWS-Certified-Machine-Learning-Specialty Testking
- Effective Latest AWS-Certified-Machine-Learning-Specialty Practice Questions - Guaranteed Amazon AWS-Certified-Machine-Learning-Specialty Exam Success with Authoritative Test AWS-Certified-Machine-Learning-Specialty Cram Download AWS-Certified-Machine-Learning-Specialty for free by simply searching on **➡ www.vce4dumps.com** AWS-Certified-Machine-Learning-Specialty Certification Book Torrent
- Latest AWS-Certified-Machine-Learning-Specialty Practice Questions - Amazon AWS Certified Machine Learning - Specialty Realistic Latest Practice Questions Pass Guaranteed Search for AWS-Certified-Machine-Learning-Specialty and download exam materials for free through **>** www.pdfvce.com New AWS-Certified-Machine-Learning-Specialty Mock Test
- Test AWS-Certified-Machine-Learning-Specialty Collection Pdf Pass4sure AWS-Certified-Machine-Learning-Specialty Dumps Pdf Test AWS-Certified-Machine-Learning-Specialty Testking Search on “www.examcollectionpass.com” for **⇒** AWS-Certified-Machine-Learning-Specialty **⇐** to obtain exam materials for free download AWS-Certified-Machine-Learning-Specialty Official Study Guide
- AWS-Certified-Machine-Learning-Specialty New Dumps Book Test AWS-Certified-Machine-Learning-Specialty Testking New AWS-Certified-Machine-Learning-Specialty Braindumps Copy URL **➡** www.pdfvce.com open and search for AWS-Certified-Machine-Learning-Specialty to download for free AWS-Certified-Machine-Learning-Specialty Test Tutorials
- New AWS-Certified-Machine-Learning-Specialty Test Fee Latest AWS-Certified-Machine-Learning-Specialty Test Materials Test AWS-Certified-Machine-Learning-Specialty Collection Pdf Easily obtain AWS-Certified-Machine-Learning-Specialty for free download through **⇒** www.easy4engine.com **⇐** New AWS-Certified-Machine-Learning-Specialty Test Fee
- Effective Latest AWS-Certified-Machine-Learning-Specialty Practice Questions - Guaranteed Amazon AWS-Certified-Machine-Learning-Specialty Exam Success with Authoritative Test AWS-Certified-Machine-Learning-Specialty Cram Search for **▶** AWS-Certified-Machine-Learning-Specialty **◀** and easily obtain a free download on **➡** www.pdfvce.com AWS-Certified-Machine-Learning-Specialty New Dumps Book
- Newest Amazon - AWS-Certified-Machine-Learning-Specialty - Latest AWS Certified Machine Learning - Specialty Practice Questions Download **▶** AWS-Certified-Machine-Learning-Specialty **◀** for free by simply entering **【** www.vce4dumps.com **】** website New AWS-Certified-Machine-Learning-Specialty Test Fee

