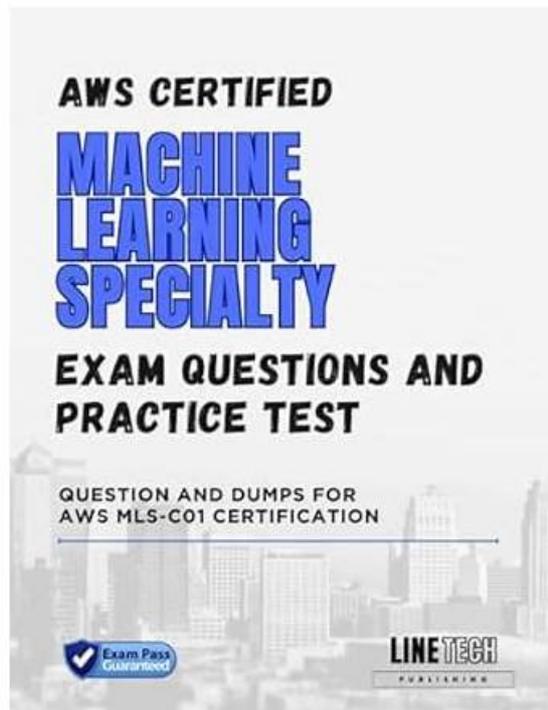


AWS-Certified-Machine-Learning-Specialty Useful Dumps - Reliable AWS-Certified-Machine-Learning-Specialty Test Prep



DOWNLOAD the newest VCEPrep AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1C4Gy5Z1gehko0ps09PIVE4jdxAvOyBNN>

VCEPrep AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam dumps save your study and preparation time. Our experts have added hundreds of AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) questions similar to the real exam. You can prepare for the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam dumps during your job. You don't need to visit the market or any store because VCEPrep AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam questions are easily accessible from the website.

Achieving the AWS Certified Machine Learning - Specialty certification can help individuals advance their careers in the field of machine learning and increase their earning potential. AWS Certified Machine Learning - Specialty certification is recognized by industry leaders and can open up new opportunities for professionals in various industries, including healthcare, finance, and retail, among others.

The AWS Certified Machine Learning - Specialty exam (MLS-C01) is a certification offered by Amazon Web Services (AWS) for individuals who want to validate their expertise in machine learning on the AWS cloud. AWS Certified Machine Learning - Specialty certification is designed to validate a candidate's understanding of the core concepts and best practices of machine learning implementation on AWS, including data preparation and cleaning, feature engineering, model development, and deployment.

>> AWS-Certified-Machine-Learning-Specialty Useful Dumps <<

Download The AWS-Certified-Machine-Learning-Specialty Useful Dumps, Pass The AWS Certified Machine Learning - Specialty

To make sure your situation of passing the certificate efficiently, our AWS-Certified-Machine-Learning-Specialty study materials are compiled by first-rank experts. So the proficiency of our team is unquestionable. They handpicked what the AWS-Certified-Machine-Learning-Specialty training guide usually tested in exam recent years and devoted their knowledge accumulated into these AWS-Certified-Machine-Learning-Specialty Actual Tests. We are on the same team, and it is our common wish to help your realize it. So you can rely on us to success and we won't let you down!

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q44-Q49):

NEW QUESTION # 44

An ecommerce company wants to use machine learning (ML) to monitor fraudulent transactions on its website. The company is using Amazon SageMaker to research, train, deploy, and monitor the ML models.

The historical transactions data is in a .csv file that is stored in Amazon S3. The data contains features such as the user's IP address, navigation time, average time on each page, and the number of clicks forsession. There is no label in the data to indicate if a transaction is anomalous.

Which models should the company use in combination to detect anomalous transactions? (Select TWO.)

- A. K-nearest neighbors (k-NN)
- **B. Random Cut Forest (RCF)**
- C. IP Insights
- D. Linear learner with a logistic function
- **E. XGBoost**

Answer: B,E

Explanation:

To detect anomalous transactions, the company can use a combination of Random Cut Forest (RCF) and XGBoost models. RCF is an unsupervised algorithm that can detect outliers in the data by measuring the depth of each data point in a collection of random decision trees. XGBoost is a supervised algorithm that can learn from the labeled data points generated by RCF and classify them as normal or anomalous. RCF can also provide anomaly scores that can be used as features for XGBoost to improve the accuracy of the classification. References:

1: Amazon SageMaker Random Cut Forest

2: Amazon SageMaker XGBoost Algorithm

3: Anomaly Detection with Amazon SageMaker Random Cut Forest and Amazon SageMaker XGBoost

NEW QUESTION # 45

A Machine Learning Specialist is working with a large cybersecurity company that manages security events in real time for companies around the world. The cybersecurity company wants to design a solution that will allow it to use machine learning to score malicious events as anomalies on the data as it is being ingested. The company also wants to be able to save the results in its data lake for later processing and analysis. What is the MOST efficient way to accomplish these tasks'?

- **A. Ingest the data using Amazon Kinesis Data Firehose, and use Amazon Kinesis Data Analytics Random Cut Forest (RCF) for anomaly detection. Then use Kinesis Data Firehose to stream the results to Amazon S3.**
- B. Ingest the data and store it in Amazon S3. Use AWS Batch along with the AWS Deep Learning AMIs to train a k-means model using TensorFlow on the data in Amazon S3.
- C. Ingest the data into Apache Spark Streaming using Amazon EMR, and use Spark MLlib with k-means to perform anomaly detection. Then store the results in an Apache Hadoop Distributed File System (HDFS) using Amazon EMR with a replication factor of three as the data lake.
- D. Ingest the data and store it in Amazon S3. Have an AWS Glue job that is triggered on demand transform the new data. Then use the built-in Random Cut Forest (RCF) model within Amazon SageMaker to detect anomalies in the data.

Answer: A

Explanation:

Amazon Kinesis Data Firehose is a fully managed service that can capture, transform, and load streaming data into AWS data stores, such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. It can also invoke AWS Lambda

functions to perform custom transformations on the data. Amazon Kinesis Data Analytics is a service that can analyze streaming data in real time using SQL or Apache Flink applications. It can also use machine learning algorithms, such as Random Cut Forest (RCF), to perform anomaly detection on streaming data. RCF is an unsupervised learning algorithm that assigns an anomaly score to each data point based on how different it is from the rest of the data. By using Kinesis Data Firehose and Kinesis Data Analytics, the cybersecurity company can ingest the data in real time, score the malicious events as anomalies, and stream the results to Amazon S3, which can serve as a data lake for later processing and analysis. This is the most efficient way to accomplish these tasks, as it does not require any additional infrastructure, coding, or training.

References:

Amazon Kinesis Data Firehose - Amazon Web Services

Amazon Kinesis Data Analytics - Amazon Web Services

Anomaly Detection with Amazon Kinesis Data Analytics - Amazon Web Services

[AWS Certified Machine Learning - Specialty Sample Questions]

NEW QUESTION # 46

A Machine Learning Specialist is designing a scalable data storage solution for Amazon SageMaker. There is an existing TensorFlow-based model implemented as a `train.py` script that relies on static training data that is currently stored as TFRecords. Which method of providing training data to Amazon SageMaker would meet the business requirements with the LEAST development overhead?

- A. Prepare the data in the format accepted by Amazon SageMaker. Use AWS Glue or AWS Lambda to reformat and store the data in an Amazon S3 bucket.
- **B. Use Amazon SageMaker script mode and use `train.py` unchanged. Put the TFRecord data into an Amazon S3 bucket. Point the Amazon SageMaker training invocation to the S3 bucket without reformatting the training data.**
- C. Rewrite the `train.py` script to add a section that converts TFRecords to protobuf and ingests the protobuf data instead of TFRecords.
- D. Use Amazon SageMaker script mode and use `train.py` unchanged. Point the Amazon SageMaker training invocation to the local path of the data without reformatting the training data.

Answer: B

Explanation:

<https://github.com/aws-samples/amazon-sagemaker-script-mode/blob/master/tf-horovod-inference-pipeline/train.py>

NEW QUESTION # 47

A company wants to forecast the daily price of newly launched products based on 3 years of data for older product prices, sales, and rebates. The time-series data has irregular timestamps and is missing some values.

Data scientist must build a dataset to replace the missing values. The data scientist needs a solution that resamples the data daily and exports the data for further modeling.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Use Amazon SageMaker Studio Notebook with Pandas.
- B. Use AWS Glue DataBrew.
- **C. Use Amazon SageMaker Studio Data Wrangler.**
- D. Use Amazon EMR Serverless with PySpark.

Answer: C

Explanation:

Amazon SageMaker Studio Data Wrangler is a visual data preparation tool that enables users to clean and normalize data without writing any code. Using Data Wrangler, the data scientist can easily import the time-series data from various sources, such as Amazon S3, Amazon Athena, or Amazon Redshift. Data Wrangler can automatically generate data insights and quality reports, which can help identify and fix missing values, outliers, and anomalies in the data. Data Wrangler also provides over 250 built-in transformations, such as resampling, interpolation, aggregation, and filtering, which can be applied to the data with a point-and-click interface. Data Wrangler can also export the prepared data to different destinations, such as Amazon S3, Amazon SageMaker Feature Store, or Amazon SageMaker Pipelines, for further modeling and analysis. Data Wrangler is integrated with Amazon SageMaker Studio, a web-based IDE for machine learning, which makes it easy to access and use the tool. Data Wrangler is a serverless and fully managed service, which means the data scientist does not need to provision, configure, or manage any infrastructure or clusters.

Option A is incorrect because Amazon EMR Serverless is a serverless option for running big data analytics applications using open-

source frameworks, such as Apache Spark. However, using Amazon EMR Serverless would require the data scientist to write PySpark code to perform the data preparation tasks, such as resampling, imputation, and aggregation. This would require more implementation effort than using Data Wrangler, which provides a visual and code-free interface for data preparation.

Option B is incorrect because AWS Glue DataBrew is another visual data preparation tool that can be used to clean and normalize data without writing code. However, DataBrew does not support time-series data as a data type, and does not provide built-in transformations for resampling, interpolation, or aggregation of time-series data. Therefore, using DataBrew would not meet the requirements of the use case.

Option D is incorrect because using Amazon SageMaker Studio Notebook with Pandas would also require the data scientist to write Python code to perform the data preparation tasks. Pandas is a popular Python library for data analysis and manipulation, which supports time-series data and provides various methods for resampling, interpolation, and aggregation. However, using Pandas would require more implementation effort than using Data Wrangler, which provides a visual and code-free interface for data preparation.

1: Amazon SageMaker Data Wrangler documentation

2: Amazon EMR Serverless documentation

3: AWS Glue DataBrew documentation

4: Pandas documentation

NEW QUESTION # 48

A manufacturer is operating a large number of factories with a complex supply chain relationship where unexpected downtime of a machine can cause production to stop at several factories. A data scientist wants to analyze sensor data from the factories to identify equipment in need of preemptive maintenance and then dispatch a service team to prevent unplanned downtime. The sensor readings from a single machine can include up to 200 data points including temperatures, voltages, vibrations, RPMs, and pressure readings. To collect this sensor data, the manufacturer deployed Wi-Fi and LANs across the factories. Even though many factory locations do not have reliable or high-speed internet connectivity, the manufacturer would like to maintain near-real-time inference capabilities. Which deployment architecture for the model will address these business requirements?

- A. Deploy the model in Amazon SageMaker. Run sensor data through this model to predict which machines need maintenance.
- B. Deploy the model in Amazon SageMaker and use an IoT rule to write data to an Amazon DynamoDB table. Consume a DynamoDB stream from the table with an AWS Lambda function to invoke the endpoint.
- C. Deploy the model to an Amazon SageMaker batch transformation job. Generate inferences in a daily batch report to identify machines that need maintenance.
- **D. Deploy the model on AWS IoT Greengrass in each factory. Run sensor data through this model to infer which machines need maintenance.**

Answer: D

Explanation:

<https://aws.amazon.com/blogs/iot/industrial-iot-from-condition-based-monitoring-to-predictive-quality-to-digitize-your-factory-with-aws-iot-services/>

<https://aws.amazon.com/blogs/iot/using-aws-iot-for-predictive-maintenance/>

NEW QUESTION # 49

.....

In order to make your exam easier for every candidate, our AWS-Certified-Machine-Learning-Specialty exam prep is capable of making you test history and review performance, and then you can find your obstacles and overcome them. In addition, once you have used this type of AWS-Certified-Machine-Learning-Specialty exam question online for one time, next time you can practice in an offline environment. The AWS-Certified-Machine-Learning-Specialty test torrent also offer a variety of learning modes for users to choose from, which can be used for multiple clients of computers and mobile phones to study online, as well as to print and print data for offline consolidation. Therefore, for your convenience, more choices are provided for you, we are pleased to suggest you to choose our AWS-Certified-Machine-Learning-Specialty Exam Question for your exam.

Reliable AWS-Certified-Machine-Learning-Specialty Test Prep: <https://www.vceprep.com/AWS-Certified-Machine-Learning-Specialty-latest-vce-prep.html>

- Take Amazon AWS-Certified-Machine-Learning-Specialty Practice Exam Questions (Desktop - Web-Based) Open www.practicevce.com enter **➡** AWS-Certified-Machine-Learning-Specialty and obtain a free download Exam Dumps AWS-Certified-Machine-Learning-Specialty Pdf

