AWS-Certified-Machine-Learning-Specialty Valid Braindumps Questions & AWS-Certified-Machine-Learning-Specialty Valid Test Cram

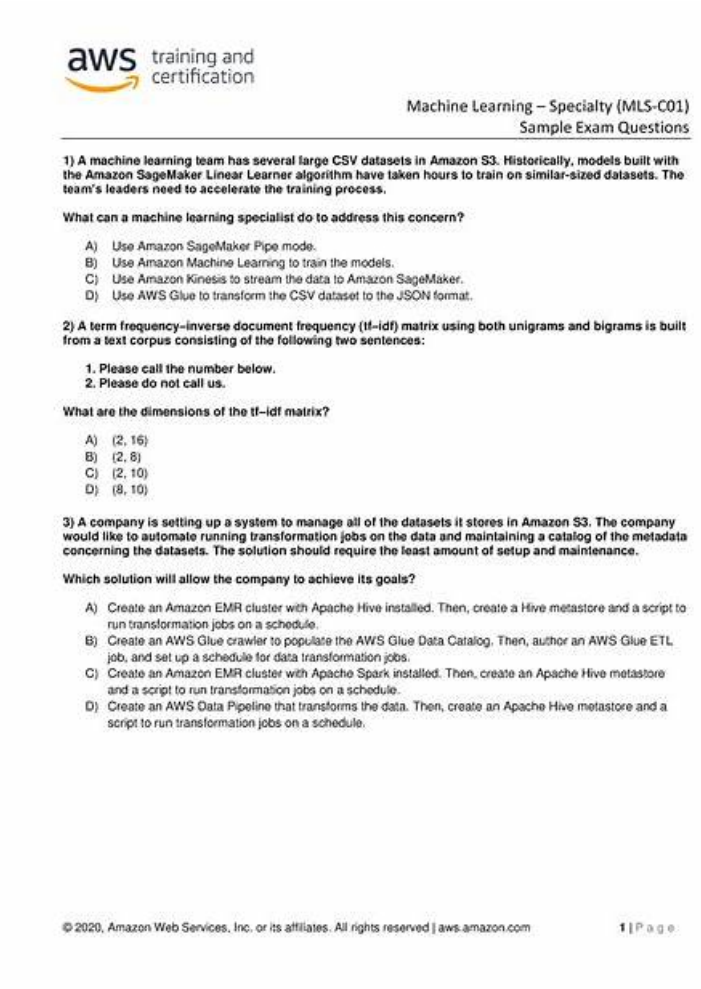

aws training and certification

Machine Learning – Specialty (MLS-C01)
Sample Exam Questions

1) A machine learning team has several large CSV datasets in Amazon S3. Historically, models built with the Amazon SageMaker Linear Learner algorithm have taken hours to train on similar-sized datasets. The team's leaders need to accelerate the training process.

What can a machine learning specialist do to address this concern?

A) Use Amazon SageMaker Pipe mode.
B) Use Amazon Machine Learning to train the models.
C) Use Amazon Kinesis to stream the data to Amazon SageMaker.
D) Use AWS Glue to transform the CSV dataset to the JSON format.

2) A term frequency–inverse document frequency (tf–idf) matrix using both unigrams and bigrams is built from a text corpus consisting of the following two sentences:

1. Please call the number below.
2. Please do not call us.

What are the dimensions of the tf–idf matrix?

A) (2, 16)
B) (2, 8)
C) (2, 10)
D) (8, 10)

3) A company is setting up a system to manage all of the datasets it stores in Amazon S3. The company would like to automate running transformation jobs on the data and maintaining a catalog of the metadata concerning the datasets. The solution should require the least amount of setup and maintenance.

Which solution will allow the company to achieve its goals?

A) Create an Amazon EMR cluster with Apache Hive installed. Then, create a Hive metastore and a script to run transformation jobs on a schedule.
B) Create an AWS Glue crawler to populate the AWS Glue Data Catalog. Then, author an AWS Glue ETL job, and set up a schedule for data transformation jobs.
C) Create an Amazon EMR cluster with Apache Spark installed. Then, create an Apache Hive metastore and a script to run transformation jobs on a schedule.
D) Create an AWS Data Pipeline that transforms the data. Then, create an Apache Hive metastore and a script to run transformation jobs on a schedule.

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by Actual4Cert: https://drive.google.com/open?id=1--DC0RviyQKXBhE9wc5cgNdoQA3qsLEE

The AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice questions (desktop and web-based) are customizable, meaning users can set the questions and time according to their needs to improve their discipline and feel the real-based exam scenario to pass the Amazon AWS-Certified-Machine-Learning-Specialty Certification. Customizable mock tests comprehensively and accurately represent the actual Amazon AWS-Certified-Machine-Learning-Specialty certification exam scenario.

The Amazon AWS-Certified-Machine-Learning-Specialty Exam is intended for professionals who are already working in the field of machine learning or those who are planning to start a career in this field. To take the exam, candidates should have a good understanding of programming languages such as Python or R, as well as experience with AWS services like S3, EC2, and SageMaker.

>> AWS-Certified-Machine-Learning-Specialty Valid Braindumps Questions <<

Amazon AWS-Certified-Machine-Learning-Specialty Valid Test Cram | AWS-Certified-Machine-Learning-Specialty Valid Exam Bootcamp

PDF format is pretty much easy to use for the ones who always have their smart devices and love to prepare for AWS Certified

Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam from them. Applicants can also make notes of printed AWS-Certified-Machine-Learning-Specialty Exam Material so they can use it anywhere in order to pass AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) with a good score.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q33-Q38):

NEW QUESTION # 33
A network security vendor needs to ingest telemetry data from thousands of endpoints that run all over the world. The data is transmitted every 30 seconds in the form of records that contain 50 fields. Each record is up to 1 KB in size. The security vendor uses Amazon Kinesis Data Streams to ingest the dat a. The vendor requires hourly summaries of the records that Kinesis Data Streams ingests. The vendor will use Amazon Athena to query the records and to generate the summaries. The Athena queries will target 7 to 12 of the available data fields.
Which solution will meet these requirements with the LEAST amount of customization to transform and store the ingested data?

- A. Use Amazon Kinesis Data Firehose to read and aggregate the data hourly. Transform the data and store it in Amazon S3 by using AWS Lambda.
- B. Use AWS Lambda to read and aggregate the data hourly. Transform the data and store it in Amazon S3 by using Amazon Kinesis Data Firehose.
- C. Use Amazon Kinesis Data Firehose to read and aggregate the data hourly. Transform the data and store it in Amazon S3 by using a short-lived Amazon EMR cluster.
- D. Use Amazon Kinesis Data Analytics to read and aggregate the data hourly. Transform the data and store it in Amazon S3 by using Amazon Kinesis Data Firehose.

Answer: D

Explanation:
The solution that will meet the requirements with the least amount of customization to transform and store the ingested data is to use Amazon Kinesis Data Analytics to read and aggregate the data hourly, transform the data and store it in Amazon S3 by using Amazon Kinesis Data Firehose. This solution leverages the built-in features of Kinesis Data Analytics to perform SQL queries on streaming data and generate hourly summaries. Kinesis Data Analytics can also output the transformed data to Kinesis Data Firehose, which can then deliver the data to S3 in a specified format and partitioning scheme. This solution does not require any custom code or additional infrastructure to process the data. The other solutions either require more customization (such as using Lambda or EMR) or do not meet the requirement of aggregating the data hourly (such as using Lambda to read the data from Kinesis Data Streams). References:
1: Boosting Resiliency with an ML-based Telemetry Analytics Architecture | AWS Architecture Blog
2: AWS Cloud Data Ingestion Patterns and Practices
3: IoT ingestion and Machine Learning analytics pipeline with AWS IoT ...
4: AWS IoT Data Ingestion Simplified 101: The Complete Guide - Hevo Data

NEW QUESTION # 34
An office security agency conducted a successful pilot using 100 cameras installed at key locations within the main office. Images from the cameras were uploaded to Amazon S3 and tagged using Amazon Rekognition, and the results were stored in Amazon ES. The agency is now looking to expand the pilot into a full production system using thousands of video cameras in its office locations globally. The goal is to identify activities performed by non-employees in real time Which solution should the agency consider?

- A. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Image to detect faces from a collection of known employees and alert when non-employees are detected.
- B. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection on each stream, and alert when non-employees are detected.
- C. Use a proxy server at each local office and for each camera, and stream the RTSP feed to a unique Amazon Kinesis Video Streams video stream. On each stream, use Amazon Rekognition Video and create a stream processor to detect faces from a collection of known employees, and alert when non-employees are detected.
- D. Install AWS DeepLens cameras and use the DeepLens_Kinesis_Video module to stream video to Amazon Kinesis Video Streams for each camera. On each stream, run an AWS Lambda function to capture image fragments and then call Amazon Rekognition Image to detect faces from a collection of known employees, and alert when non-employees are detected.

Answer: B

Explanation:
https://aws.amazon.com/blogs/machine-learning/video-analytics-in-the-cloud-and-at-the-edge- with-aws-deeplens-and-kinesis-video-streams/

NEW QUESTION # 35
A financial company is trying to detect credit card fraud. The company observed that, on average, 2% of credit card transactions were fraudulent. A data scientist trained a classifier on a year's worth of credit card transactions data. The model needs to identify the fraudulent transactions (positives) from the regular ones (negatives). The company's goal is to accurately capture as many positives as possible.
Which metrics should the data scientist use to optimize the model? (Choose two.)

- A. Accuracy
- B. True positive rate
- C. False positive rate
- D. Area under the precision-recall curve
- E. Specificity

Answer: B,D

Explanation:
The data scientist should use the area under the precision-recall curve and the true positive rate to optimize the model. These metrics are suitable for imbalanced classification problems, such as credit card fraud detection, where the positive class (fraudulent transactions) is much rarer than the negative class (non-fraudulent transactions).
The area under the precision-recall curve (AUPRC) is a measure of how well the model can identify the positive class among all the predicted positives. Precision is the fraction of predicted positives that are actually positive, and recall is the fraction of actual positives that are correctly predicted. A higher AUPRC means that the model can achieve a higher precision with a higher recall, which is desirable for fraud detection.
The true positive rate (TPR) is another name for recall. It is also known as sensitivity or hit rate. It measures the proportion of actual positives that are correctly identified by the model. A higher TPR means that the model can capture more positives, which is the company's goal.
References:
* Metrics for Imbalanced Classification in Python - Machine Learning Mastery
* Precision-Recall - scikit-learn

NEW QUESTION # 36
A Machine Learning Specialist needs to be able to ingest streaming data and store it in Apache Parquet files for exploration and analysis. Which of the following services would both ingest and store this data in the correct format?

- A. Amazon Kinesis Data Firehose
- B. AWSDMS
- C. Amazon Kinesis Data Streams
- D. Amazon Kinesis Data Analytics

Answer: A

Explanation:
Explanation
Amazon Kinesis Data Firehose is a service that can ingest streaming data and store it in various destinations, including Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. Amazon Kinesis Data Firehose can also convert the incoming data to Apache Parquet or Apache ORC format before storing it in Amazon S3. This can reduce the storage cost and improve the performance of analytical queries on the data.
Amazon Kinesis Data Firehose supports various data sources, such as Amazon Kinesis Data Streams, Amazon Managed Streaming for Apache Kafka, AWS IoT, and custom applications. Amazon Kinesis Data Firehose can also apply data transformation and compression using AWS Lambda functions.
AWSDMS is not a valid service name. AWS Database Migration Service (AWS DMS) is a service that can migrate data from various sources to various targets, but it does not support streaming data or Parquet format.
Amazon Kinesis Data Streams is a service that can ingest and process streaming data in real time, but it does not store the data in any destination. Amazon Kinesis Data Streams can be integrated with Amazon Kinesis Data Firehose to store the data in Parquet

format.
Amazon Kinesis Data Analytics is a service that can analyze streaming data using SQL or Apache Flink, but it does not store the data in any destination. Amazon Kinesis Data Analytics can be integrated with Amazon Kinesis Data Firehose to store the data in Parquet format. References:
Amazon Kinesis Data Firehose - Amazon Web Services
What Is Amazon Kinesis Data Firehose? - Amazon Kinesis Data Firehose
Amazon Kinesis Data Firehose FAQs - Amazon Web Services

NEW QUESTION # 37
A data scientist stores financial datasets in Amazon S3. The data scientist uses Amazon Athena to query the datasets by using SQL. The data scientist uses Amazon SageMaker to deploy a machine learning (ML) model. The data scientist wants to obtain inferences from the model at the SageMaker endpoint However, when the data .... ntist attempts to invoke the SageMaker endpoint, the data scientist receives SOL statement failures The data scientist's 1AM user is currently unable to invoke the SageMaker endpoint Which combination of actions will give the data scientist's 1AM user the ability to invoke the SageMaker endpoint? (Select THREE.)

- A. Include an inline policy for the data scientist's 1AM user that allows SageMaker to read S3 objects
- B. Attach the AmazonAthenaFullAccess AWS managed policy to the user identity.
- C. Include the SQL statement "USING EXTERNAL FUNCTION ml_function_name" in the Athena SQL query.
- D. Include a policy statement for the data scientist's 1AM user that allows the 1AM user to perform the sagemakerGetRecord action.
- E. Include a policy statement for the data scientist's 1AM user that allows the 1AM user to perform the sagemaker: lnvokeEndpoint action,
- F. Perform a user remapping in SageMaker to map the 1AM user to another 1AM user that is on the hosted endpoint.

Answer: A,C,E

Explanation:
The correct combination of actions to enable the data scientist's IAM user to invoke the SageMaker endpoint is B, C, and E, because they ensure that the IAM user has the necessary permissions, access, and syntax to query the ML model from Athena. These actions have the following benefits:
* B: Including a policy statement for the IAM user that allows the sagemaker:InvokeEndpoint action grants the IAM user the permission to call the SageMaker Runtime InvokeEndpoint API, which is used to get inferences from the model hosted at the endpoint1.
* C: Including an inline policy for the IAM user that allows SageMaker to read S3 objects enables the IAM user to access the data stored in S3, which is the source of the Athena queries2.
* E: Including the SQL statement "USING EXTERNAL FUNCTION ml_function_name" in the Athena SQL query allows the IAM user to invoke the ML model as an external function from Athena, which is a feature that enables querying ML models from SQL statements3.
The other options are not correct or necessary, because they have the following drawbacks:
* A: Attaching the AmazonAthenaFullAccess AWS managed policy to the user identity is not sufficient, because it does not grant the IAM user the permission to invoke the SageMaker endpoint, which is required to query the ML model4.
* D: Including a policy statement for the IAM user that allows the IAM user to perform the sagemaker:
GetRecord action is not relevant, because this action is used to retrieve a single record from a feature group, which is not the case in this scenario5.
* F: Performing a user remapping in SageMaker to map the IAM user to another IAM user that is on the hosted endpoint is not applicable, because this feature is only available for multi-model endpoints, which are not used in this scenario.
References:
* 1: InvokeEndpoint - Amazon SageMaker
* 2: Querying Data in Amazon S3 from Amazon Athena - Amazon Athena
* 3: Querying machine learning models from Amazon Athena using Amazon SageMaker | AWS Machine Learning Blog
* 4: AmazonAthenaFullAccess - AWS Identity and Access Management
* 5: GetRecord - Amazon SageMaker Feature Store Runtime
* : [Invoke a Multi-Model Endpoint - Amazon SageMaker]

NEW QUESTION # 38
.....

Specialty Questions. You should practice with actual AWS-Certified-Machine-Learning-Specialty exam questions that are aligned with the latest content of the AWS-Certified-Machine-Learning-Specialty test. These Amazon AWS-Certified-Machine-Learning-Specialty exam questions remove the need for you to spend time on unnecessary or irrelevant material, allowing you to complete your AWS-Certified-Machine-Learning-Specialty Certification Exam preparation swiftly. You can save time and clear the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) test in one sitting if you skip unnecessary material and focus on our AWS-Certified-Machine-Learning-Specialty actual questions.

AWS-Certified-Machine-Learning-Specialty Valid Test Cram: https://www.actual4cert.com/AWS-Certified-Machine-Learning-Specialty-real-questions.html

- AWS-Certified-Machine-Learning-Specialty Reliable Source 🈵 Reliable AWS-Certified-Machine-Learning-Specialty Exam Sample □ AWS-Certified-Machine-Learning-Specialty Pdf Braindumps □ Open website 《 www.vceengine.com 》 and search for 「 AWS-Certified-Machine-Learning-Specialty 」 for free download □AWS-Certified-Machine-Learning-Specialty Interactive Questions
- Amazon AWS-Certified-Machine-Learning-Specialty Questions - Reduce your Chances of Failure in Exam □ Search for ➥ AWS-Certified-Machine-Learning-Specialty □ and easily obtain a free download on ➽ www.pdfvce.com □ □ □AWS-Certified-Machine-Learning-Specialty Dump
- Examcollection AWS-Certified-Machine-Learning-Specialty Dumps Torrent □ New Braindumps AWS-Certified-Machine-Learning-Specialty Book □ Exam AWS-Certified-Machine-Learning-Specialty Tutorial □ Search for ➠ AWS-Certified-Machine-Learning-Specialty □ and obtain a free download on ▹ www.pdfdumps.com ◃ □Reliable AWS-Certified-Machine-Learning-Specialty Exam Sample
- AWS-Certified-Machine-Learning-Specialty Exam Dump □ AWS-Certified-Machine-Learning-Specialty New Dumps Book □ New Braindumps AWS-Certified-Machine-Learning-Specialty Book □ Download ☀ AWS-Certified-Machine-Learning-Specialty □☀□ for free by simply entering □ www.pdfvce.com □ website □Latest AWS-Certified-Machine-Learning-Specialty Test Vce
- AWS-Certified-Machine-Learning-Specialty New Dumps Book □ AWS-Certified-Machine-Learning-Specialty Pdf Braindumps □ Latest AWS-Certified-Machine-Learning-Specialty Test Vce □ Go to website □ www.examcollectionpass.com □ open and search for （ AWS-Certified-Machine-Learning-Specialty ） to download for free □AWS-Certified-Machine-Learning-Specialty Reliable Test Blueprint
- AWS-Certified-Machine-Learning-Specialty Pass4sure Dumps Pdf □ Sample AWS-Certified-Machine-Learning-Specialty Exam □ Exam AWS-Certified-Machine-Learning-Specialty Tutorial □ ➡ www.pdfvce.com □□□ is best website to obtain ➤ AWS-Certified-Machine-Learning-Specialty □ for free download □Examcollection AWS-Certified-Machine-Learning-Specialty Dumps Torrent
- Superb AWS-Certified-Machine-Learning-Specialty Exam Questions Supply You Marvelous Learning Dumps - www.exam4labs.com □ Search for [AWS-Certified-Machine-Learning-Specialty] on □ www.exam4labs.com □ immediately to obtain a free download □AWS-Certified-Machine-Learning-Specialty Reliable Exam Dumps
- Latest AWS-Certified-Machine-Learning-Specialty Exam Pass4sure ↘ Updated AWS-Certified-Machine-Learning-Specialty CBT □ Exam AWS-Certified-Machine-Learning-Specialty Tutorial □ Search for ➠ AWS-Certified-Machine-Learning-Specialty □ and obtain a free download on ➡ www.pdfvce.com □□□ □AWS-Certified-Machine-Learning-Specialty Pdf Format
- Sample AWS-Certified-Machine-Learning-Specialty Exam □ AWS-Certified-Machine-Learning-Specialty Reliable Exam Dumps □ Sample AWS-Certified-Machine-Learning-Specialty Exam □ Search for ➤ AWS-Certified-Machine-Learning-Specialty □ and download it for free immediately on □ www.troytecdumps.com □ □New Braindumps AWS-Certified-Machine-Learning-Specialty Book
- Latest AWS-Certified-Machine-Learning-Specialty Exam Pass4sure □ AWS-Certified-Machine-Learning-Specialty Reliable Exam Dumps □ AWS-Certified-Machine-Learning-Specialty Pdf Braindumps □ Immediately open ➡ www.pdfvce.com □□□ and search for ➽ AWS-Certified-Machine-Learning-Specialty □ to obtain a free download □ □AWS-Certified-Machine-Learning-Specialty Exam Dump
- AWS-Certified-Machine-Learning-Specialty Reliable Exam Dumps □ Reliable AWS-Certified-Machine-Learning-Specialty Exam Sample □ Sample AWS-Certified-Machine-Learning-Specialty Exam □ Easily obtain （ AWS-Certified-Machine-Learning-Specialty ） for free download through ▸ www.vce4dumps.com ◂ □New Braindumps AWS-Certified-Machine-Learning-Specialty Book
- myportal.utt.edu.tt, www.mixcloud.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, onlyfans.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, github.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

BTW, DOWNLOAD part of Actual4Cert AWS-Certified-Machine-Learning-Specialty dumps from Cloud Storage: https://drive.google.com/open?id=1--DC0RviyQKXBhE9wc5cgNdoQA3qsLEE