Reliable MLS-C01 Test Pattern - MLS-C01 Examcollection Dumps

P.S. Free & New MLS-C01 dumps are available on Google Drive shared by LatestCram: https://drive.google.com/open?id=1OXwaSsoBXWCnrZmtNx92-AXittpMGoMG

Our product boosts many merits and high passing rate. Our products have 3 versions and we provide free update of the Amazon exam torrent to you. If you are the old client you can enjoy the discounts. Most important of all, as long as we have compiled a new version of the MLS-C01 Exam Questions, we will send the latest version of our Amazon exam questions to our customers for free during the whole year after purchasing. Our product can improve your stocks of knowledge and your abilities in some area and help you gain the success in your career.

The Amazon AWS Certified Machine Learning – Specialty certification verifies that you can effectively train, tune, build, and deploy machine learning models using AWS. This certificate validates the professionals' skills, such as:

- Identify the correct AWS services to implement machine learning solutions.
- Selecting and justifying the correct machine learning (ML) approach for a specific business problem;
- Designing and implementing secure, scalable, cost-optimized & reliable machine learning solutions;

To obtain Amazon AWS Certified Machine Learning – Specialty, the applicants must pass the MLS-C01 Exam.

Target Audience

The Amazon MLS-C01 exam is targeted at those individuals who are tasked with performing the data science or development role. It provides that the candidates can design, deploy, implement, and maintain ML or machine learning solutions for given business problems.

>> Reliable MLS-C01 Test Pattern <<

MLS-C01 Examcollection Dumps | Current MLS-C01 Exam Content

There are a lot of experts and professors in or company in the field. In order to meet the demands of all people, these excellent experts and professors from our company have been working day and night. They tried their best to design the best MLS-C01 study materials from our company for all people. By our study materials, all people can prepare for their MLS-C01 exam in the more efficient method. We can guarantee that our study materials will be suitable for all people and meet the demands of all people, including students, workers and housewives and so on. If you decide to buy and use the MLS-C01 Study Materials from our company with dedication on and enthusiasm step and step, it will be very easy for you to pass the exam without doubt. We sincerely hope that you can achieve your dream in the near future by the MLS-C01 study materials of our company.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q96-Q101):

NEW QUESTION # 96
Machine Learning Specialist is training a model to identify the make and model of vehicles in images. The Specialist wants to use transfer learning and an existing model trained on images of general objects. The Specialist collated a large custom dataset of pictures containing different vehicle makes and models.
What should the Specialist do to initialize the model to re-train it with the custom data?

- A. Initialize the model with random weights in all layers including the last fully connected layer.
- B. Initialize the model with pre-trained weights in all layers and replace the last fully connected layer.
- C. Initialize the model with random weights in all layers and replace the last fully connected layer.
- D. Initialize the model with pre-trained weights in all layers including the last fully connected layer.

Answer: B

NEW QUESTION # 97
A company is building a line-counting application for use in a quick-service restaurant. The company wants to use video cameras pointed at the line of customers at a given register to measure how many people are in line and deliver notifications to managers if the line grows too long. The restaurant locations have limited bandwidth for connections to external services and cannot accommodate multiple video streams without impacting other operations.
Which solution should a machine learning specialist implement to meet these requirements?

- A. Install cameras compatible with Amazon Kinesis Video Streams to stream the data to AWS over the restaurant's existing internet connection. Write an AWS Lambda function to take an image and send it to Amazon Rekognition to count the number of faces in the image. Send an Amazon Simple Notification Service (Amazon SNS) notification if the line is too long.
- B. Build a custom model in Amazon SageMaker to recognize the number of people in an image. Deploy AWS DeepLens cameras in the restaurant. Deploy the model to the cameras. Deploy an AWS Lambda function to the cameras to use the model to count people and send an Amazon Simple Notification Service (Amazon SNS) notification if the line is too long.
- C. Build a custom model in Amazon SageMaker to recognize the number of people in an image. Install cameras compatible with Amazon Kinesis Video Streams in the restaurant. Write an AWS Lambda function to take an image. Use the SageMaker endpoint to call the model to count people. Send an Amazon Simple Notification Service (Amazon SNS) notification if the line is too long.
- D. Deploy AWS DeepLens cameras in the restaurant to capture video. Enable Amazon Rekognition on the AWS DeepLens device, and use it to trigger a local AWS Lambda function when a person is recognized. Use the Lambda function to send an Amazon Simple Notification Service (Amazon SNS) notification if the line is too long.

Answer: B

Explanation:
The best solution for building a line-counting application for use in a quick-service restaurant is to use the following steps:
* Build a custom model in Amazon SageMaker to recognize the number of people in an image. Amazon SageMaker is a fully managed service that provides tools and workflows for building, training, and deploying machine learning models. A custom model can be tailored to the specific use case of line- counting and achieve higher accuracy than a generic model1
* Deploy AWS DeepLens cameras in the restaurant to capture video. AWS DeepLens is a wireless video camera that integrates with Amazon SageMaker and AWS Lambda. It can run machine learning inference locally on the device without requiring internet connectivity or streaming video to the cloud. This reduces the bandwidth consumption and latency of the application2
* Deploy the model to the cameras. AWS DeepLens allows users to deploy trained models from Amazon SageMaker to the cameras with a few clicks. The cameras can then use the model to process the video frames and count the number of people in each

frame2
* Deploy an AWS Lambda function to the cameras to use the model to count people and send an Amazon Simple Notification Service (Amazon SNS) notification if the line is too long. AWS Lambda is a serverless computing service that lets users run code without provisioning or managing servers. AWS DeepLens supports running Lambda functions on the device to perform actions based on the inference results. Amazon SNS is a service that enables users to send notifications to subscribers via email, SMS, or mobile push23 The other options are incorrect because they either require internet connectivity or streaming video to the cloud, which may impact the bandwidth and performance of the application. For example:
* Option A uses Amazon Kinesis Video Streams to stream the data to AWS over the restaurant's existing internet connection. Amazon Kinesis Video Streams is a service that enables users to capture, process, and store video streams for analytics and machine learning. However, this option requires streaming multiple video streams to the cloud, which may consume a lot of bandwidth and cause network congestion. It also requires internet connectivity, which may not be reliable or available in some locations4
* Option B uses Amazon Rekognition on the AWS DeepLens device. Amazon Rekognition is a service that provides computer vision capabilities, such as face detection, face recognition, and object detection. However, this option requires calling the Amazon Rekognition API over the internet, which may introduce latency and require bandwidth. It also uses a generic face detection model, which may not be optimized for the line-counting use case.
* Option C uses Amazon SageMaker to build a custom model and an Amazon SageMaker endpoint to call the model. Amazon SageMaker endpoints are hosted web services that allow users to perform inference on their models. However, this option requires sending the images to the endpoint over the internet, which may consume bandwidth and introduce latency. It also requires internet connectivity, which may not be reliable or available in some locations.
1: Amazon SageMaker - Machine Learning Service - AWS
2: AWS DeepLens - Deep learning enabled video camera - AWS
3: Amazon Simple Notification Service (SNS) - AWS
4: Amazon Kinesis Video Streams - Amazon Web Services
Amazon Rekognition - Video and Image - AWS
Deploy a Model - Amazon SageMaker

NEW QUESTION # 98
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. Create an AWS Glue job. Convert the scripts to PySpark. Execute the pipeline. Store the results in Amazon S3.
- B. 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.
- C. Load the data into an Amazon Redshift cluster. Execute the pipeline by using SQL. Store the results in Amazon S3.
- D. 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.

Answer: A

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 # 99
A Mobile Network Operator is building an analytics platform to analyze and optimize a company's operations using Amazon Athena and Amazon S3 The source systems send data in CSV format in real lime The Data Engineering team wants to transform the data to the Apache Parquet format before storing it on Amazon S3 Which solution takes the LEAST effort to implement?

- A. Ingest .CSV data from Amazon Kinesis Data Streams and use Amazon Glue to convert data into Parquet.
- B. Ingest .CSV data using Apache Spark Structured Streaming in an Amazon EMR cluster and use ApacheSpark to convert data into Parquet.
- C. Ingest .CSV data from Amazon Kinesis Data Streams and use Amazon Kinesis Data Firehose to convertdata into Parquet.
- D. Ingest .CSV data using Apache Kafka Streams on Amazon EC2 instances and use Kafka Connect S3 toserialize data as Parquet

Answer: C

Explanation:
Amazon Kinesis Data Streams is a service that can capture, store, and process streaming data in real time.
Amazon Kinesis Data Firehose is a service that can deliver streaming data to various destinations, such as Amazon S3, Amazon Redshift, or Amazon Elasticsearch Service. Amazon Kinesis Data Firehose can also transform the data before delivering it, such as converting the data format, compressing the data, or encrypting the data. One of the supported data formats that Amazon Kinesis Data Firehose can convert to is Apache Parquet, which is a columnar storage format that can improve the performance and cost-

efficiency of analytics queries. By using Amazon Kinesis Data Streams and Amazon Kinesis Data Firehose, the Mobile Network Operator can ingest the .CSV data from the source systems and use Amazon Kinesis Data Firehose to convert the data into Parquet before storing it on Amazon S3. This solution takes the least effort to implement, as it does not require any additional resources, such as Amazon EC2 instances, Amazon EMR clusters, or Amazon Glue jobs. The solution can also leverage the built-in features of Amazon Kinesis Data Firehose, such as data buffering, batching, retry, and error handling.
Amazon Kinesis Data Streams - Amazon Web Services
Amazon Kinesis Data Firehose - Amazon Web Services
Data Transformation - Amazon Kinesis Data Firehose
Apache Parquet - Amazon Athena

NEW QUESTION # 100
A company uses a long short-term memory (LSTM) model to evaluate the risk factors of a particular energy sector. The model reviews multi-page text documents to analyze each sentence of the text and categorize it as either a potential risk or no risk. The model is not performing well, even though the Data Scientist has experimented with many different network structures and tuned the corresponding hyperparameters.
Which approach will provide the MAXIMUM performance boost?

- A. Use gated recurrent units (GRUs) instead of LSTM and run the training process until the validation loss stops decreasing.
- B. Reduce the learning rate and run the training process until the training loss stops decreasing.
- C. Initialize the words by term frequency-inverse document frequency (TF-IDF) vectors pretrained on a large collection of news articles related to the energy sector.
- D. Initialize the words by word2vec embeddings pretrained on a large collection of news articles related to the energy sector.

Answer: D

Explanation:
Explanation
Initializing the words by word2vec embeddings pretrained on a large collection of news articles related to the energy sector will provide the maximum performance boost for the LSTM model. Word2vec is a technique that learns distributed representations of words based on their co-occurrence in a large corpus of text. These representations capture semantic and syntactic similarities between words, which can help the LSTM model better understand the meaning and context of the sentences in the text documents. Using word2vec embeddings that are pretrained on a relevant domain (energy sector) can further improve the performance by reducing the vocabulary mismatch and increasing the coverage of the words in the text documents. References
:
AWS Machine Learning Specialty Exam Guide
AWS Machine Learning Training - Text Classification with TF-IDF, LSTM, BERT: a comparison of performance AWS Machine Learning Training - Machine Learning - Exam Preparation Path

NEW QUESTION # 101
.....

- MLS-C01 New Question □ Reliable MLS-C01 Exam Questions □ MLS-C01 Relevant Answers □ Search for （ MLS-C01 ） and download it for free on { www.pdfvce.com } website □Practice Test MLS-C01 Pdf
- Well-Prepared Reliable MLS-C01 Test Pattern - Leading Offer in Qualification Exams - Updated Amazon AWS Certified Machine Learning - Specialty □ Search for ☀ MLS-C01 □☀□ and download it for free immediately on { www.exam4labs.com } □Positive MLS-C01 Feedback
- Test MLS-C01 Testking □ MLS-C01 Exam Questions Pdf □ Exam MLS-C01 Book □ Search for ▶ MLS-C01 ◀ and easily obtain a free download on [www.pdfvce.com] □New MLS-C01 Test Bootcamp
- MLS-C01 Valid Exam Pattern □ MLS-C01 New Question □ Test MLS-C01 Testking □ Search for ▷ MLS-C01 ◁ and download exam materials for free through □ www.practicevce.com □ □MLS-C01 Real Torrent
- MLS-C01 Pass Test □ MLS-C01 Exam Questions Pdf □ Online MLS-C01 Training □ Simply search for " MLS-C01 " for free download on ➤ www.pdfvce.com □ □Positive MLS-C01 Feedback
- Test MLS-C01 Cram Review □ Reliable MLS-C01 Exam Questions □ Test MLS-C01 Testking □ Immediately open ✔ www.practicevce.com □✔□ and search for ➽ MLS-C01 □ to obtain a free download □MLS-C01 New Question
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, 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, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

BTW, DOWNLOAD part of LatestCram MLS-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1OXwaSsoBXWCnrZmtNx92-AXittpMGoMG