

Three Easy-to-Use PassLeaderVCE Amazon MLS-C01 Exam Dumps Formats



P.S. Free & New MLS-C01 dumps are available on Google Drive shared by PassLeaderVCE: <https://drive.google.com/open?id=1d4E6LDxKYe3IHs73NpT11TzCG1DHo4LM>

Your AWS Certified Machine Learning - Specialty (MLS-C01) exam anxiety will be reduced by having the chance to practice under the MLS-C01 real exam environment created by this software. The objective of PassLeaderVCE is to offer excellent AWS Certified Machine Learning - Specialty (MLS-C01) test simulation software to its customers. Thus it is offering an exceptional and dedicated 24/7 customer support team to assist its users.

The AWS-Certified-Machine-Learning-Specialty exam is a challenging certification exam aimed at validating the skills and knowledge of individuals who want to design, implement, and deploy machine learning solutions using AWS services. Candidates who Pass MLS-C01 Exam can demonstrate their expertise in machine learning and can enhance their career prospects in the field of data science and machine learning.

>> MLS-C01 New Dumps Files <<

MLS-C01 Valid Test Test, MLS-C01 Test Duration

In the 21 Century, the MLS-C01 certification became more and more recognized in the society because it represented the certain ability of examinees. However, in order to obtain MLS-C01 certification, you have to spend a lot of time preparing for the MLS-C01 Exam. Many people gave up because of all kinds of difficulties before the examination, and finally lost the opportunity to enhance their self-worth. But our MLS-C01 exam questions will help you pass the exam for sure.

For more info visit:

AWS Certified Solutions Architect - Associate

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q174-Q179):

NEW QUESTION # 174

A Machine Learning Specialist needs to create a data repository to hold a large amount of time-based training data for a new model. In the source system, new files are added every hour. Throughout a single 24-hour period, the volume of hourly updates will change significantly. The Specialist always wants to train on the last 24 hours of the data.

Which type of data repository is the MOST cost-effective solution?

- A. An Amazon RDS database with hourly table partitions
- B. An Amazon EMR cluster with hourly hive partitions on Amazon EBS volumes
- C. An Amazon EBS-backed Amazon EC2 instance with hourly directories
- D. An Amazon S3 data lake with hourly object prefixes

Answer: D

Explanation:

Explanation

An Amazon S3 data lake is a cost-effective solution for storing and analyzing large amounts of time-based training data for a new model. Amazon S3 is a highly scalable, durable, and secure object storage service that can store any amount of data in any format. Amazon S3 also offers low-cost storage classes, such as S3 Standard-IA and S3 One Zone-IA, that can reduce the storage costs for infrequently accessed data. By using hourly object prefixes, the Machine Learning Specialist can organize the data into logical partitions based on the time of ingestion. This can enable efficient data access and management, as well as support incremental updates and deletes. The Specialist can also use Amazon S3 lifecycle policies to automatically transition the data to lower-cost storage classes or delete the data after a certain period of time. This way, the Specialist can always train on the last 24 hours of the data and optimize the storage costs.

References:

What is a data lake? - Amazon Web Services

Amazon S3 Storage Classes - Amazon Simple Storage Service

Managing your storage lifecycle - Amazon Simple Storage Service

Best Practices Design Patterns: Optimizing Amazon S3 Performance

NEW QUESTION # 175

A data scientist is building a linear regression model. The scientist inspects the dataset and notices that the mode of the distribution is lower than the median, and the median is lower than the mean.

Which data transformation will give the data scientist the ability to apply a linear regression model?

- **A. Logarithmic transformation**
- B. Polynomial transformation
- C. Exponential transformation
- D. Sinusoidal transformation

Answer: A

Explanation:

A logarithmic transformation is a suitable data transformation for a linear regression model when the data has a skewed distribution, such as when the mode is lower than the median and the median is lower than the mean. A logarithmic transformation can reduce the skewness and make the data more symmetric and normally distributed, which are desirable properties for linear regression. A logarithmic transformation can also reduce the effect of outliers and heteroscedasticity (unequal variance) in the data. An exponential transformation would have the opposite effect of increasing the skewness and making the data more asymmetric. A polynomial transformation may not be able to capture the nonlinearity in the data and may introduce multicollinearity among the transformed variables. A sinusoidal transformation is not appropriate for data that does not have a periodic pattern.

References:

Data Transformation - Scaler Topics

Linear Regression - GeeksforGeeks

Linear Regression - Scribbr

NEW QUESTION # 176

A Machine Learning Specialist is developing a custom video recommendation model for an application. The dataset used to train this model is very large with millions of data points and is hosted in an Amazon S3 bucket. The Specialist wants to avoid loading all of this data onto an Amazon SageMaker notebook instance because it would take hours to move and will exceed the attached 5 GB Amazon EBS volume on the notebook instance.

Which approach allows the Specialist to use all the data to train the model?

- **A. Load a smaller subset of the data into the SageMaker notebook and train locally. Confirm that the training code is executing and the model parameters seem reasonable. Launch an Amazon EC2 instance with an AWS Deep Learning AMI and attach the S3 bucket to train the full dataset.**
- B. Launch an Amazon EC2 instance with an AWS Deep Learning AMI and attach the S3 bucket to the instance. Train on a small amount of the data to verify the training code and hyperparameters. Go back to Amazon SageMaker and train using the full dataset
- C. Load a smaller subset of the data into the SageMaker notebook and train locally. Confirm that the training code is executing and the model parameters seem reasonable. Initiate a SageMaker training job using the full dataset from the S3 bucket using Pipe input mode.

- D. Use AWS Glue to train a model using a small subset of the data to confirm that the data will be compatible with Amazon SageMaker. Initiate a SageMaker training job using the full dataset from the S3 bucket using Pipe input mode.

Answer: A

NEW QUESTION # 177

A retail company is ingesting purchasing records from its network of 20,000 stores to Amazon S3 by using Amazon Kinesis Data Firehose. The company uses a small, server-based application in each store to send the data to AWS over the internet. The company uses this data to train a machine learning model that is retrained each day. The company's data science team has identified existing attributes on these records that could be combined to create an improved model.

Which change will create the required transformed records with the LEAST operational overhead?

- A. Deploy an Amazon EMR cluster that runs Apache Spark and includes the transformation logic. Use Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to launch the cluster each day and transform the records that accumulate in Amazon S3. Deliver the transformed records to Amazon S3.
- **B. Create an AWS Lambda function that can transform the incoming records. Enable data transformation on the ingestion Kinesis Data Firehose delivery stream. Use the Lambda function as the invocation target.**
- C. Launch a fleet of Amazon EC2 instances that include the transformation logic. Configure the EC2 instances with a daily cron job to transform the records that accumulate in Amazon S3. Deliver the transformed records to Amazon S3.
- D. Deploy an Amazon S3 File Gateway in the stores. Update the in-store software to deliver data to the S3 File Gateway. Use a scheduled daily AWS Glue job to transform the data that the S3 File Gateway delivers to Amazon S3.

Answer: B

Explanation:

Explanation

The solution A will create the required transformed records with the least operational overhead because it uses AWS Lambda and Amazon Kinesis Data Firehose, which are fully managed services that can provide the desired functionality. The solution A involves the following steps:

Create an AWS Lambda function that can transform the incoming records. AWS Lambda is a service that can run code without provisioning or managing servers. AWS Lambda can execute the transformation logic on the purchasing records and add the new attributes to the records¹.

Enable data transformation on the ingestion Kinesis Data Firehose delivery stream. Use the Lambda function as the invocation target. Amazon Kinesis Data Firehose is a service that can capture, transform, and load streaming data into AWS data stores. Amazon Kinesis Data Firehose can enable data transformation and invoke the Lambda function to process the incoming records before delivering them to Amazon S3. This can reduce the operational overhead of managing the transformation process and the data storage².

The other options are not suitable because:

Option B: Deploying an Amazon EMR cluster that runs Apache Spark and includes the transformation logic, using Amazon EventBridge (Amazon CloudWatch Events) to schedule an AWS Lambda function to launch the cluster each day and transform the records that accumulate in Amazon S3, and delivering the transformed records to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the Amazon EMR cluster, the Apache Spark application, the AWS Lambda function, and the Amazon EventBridge rule. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day³.

Option C: Deploying an Amazon S3 File Gateway in the stores, updating the in-store software to deliver data to the S3 File Gateway, and using a scheduled daily AWS Glue job to transform the data that the S3 File Gateway delivers to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the S3 File Gateway, the in-store software, and the AWS Glue job. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day⁴.

Option D: Launching a fleet of Amazon EC2 instances that include the transformation logic, configuring the EC2 instances with a daily cron job to transform the records that accumulate in Amazon S3, and delivering the transformed records to Amazon S3 will incur more operational overhead than using AWS Lambda and Amazon Kinesis Data Firehose. The company will have to manage the EC2 instances, the transformation code, and the cron job. Moreover, this solution will introduce a delay in the transformation process, as it will run only once a day⁵.

References:

1: AWS Lambda

2: Amazon Kinesis Data Firehose

3: Amazon EMR

4: Amazon S3 File Gateway

5: Amazon EC2

An e-commerce company needs a customized training model to classify images of its shirts and pants products. The company needs a proof of concept in 2 to 3 days with good accuracy. Which compute choice should the Machine Learning Specialist select to train and achieve good accuracy on the model quickly?

- Answer: B**

• • • • •

[illegible]

BONUS!!! Download part of PassLeaderVCE MLS-C01 dumps for free: <https://drive.google.com/open?id=1d4E6LDxKYe3IHs73NpT11TzCG1DHo4LM>