

Latest AWS-Certified-Machine-Learning-Specialty Exam Vce | AWS-Certified-Machine-Learning-Specialty Latest Braindumps Free



P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by UpdateDumps:
<https://drive.google.com/open?id=1pFQexP9PUuZv6T68jlr7hjWTyts7-PRg>

Profit from the opportunity to get these top-notch exam questions for the Amazon AWS-Certified-Machine-Learning-Specialty certification test. We guarantee you that our top-rated Amazon AWS-Certified-Machine-Learning-Specialty practice exam (PDF, desktop practice test software, and web-based practice exam) will enable you to pass the Amazon AWS-Certified-Machine-Learning-Specialty Certification Exam on the very first go.

This updated Amazon AWS-Certified-Machine-Learning-Specialty exam study material of UpdateDumps consists of these 3 formats: AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) PDF, desktop practice test software, and web-based practice exam. Each format of UpdateDumps aids a specific preparation style and offers unique advantages, each of which is beneficial for strong AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam preparation. The features of our three formats are listed below. You can choose any format as per your practice needs.

>> Latest AWS-Certified-Machine-Learning-Specialty Exam Vce <<

AWS-Certified-Machine-Learning-Specialty Latest Braindumps Free & Test AWS-Certified-Machine-Learning-Specialty Dumps Demo

Our AWS-Certified-Machine-Learning-Specialty practice questions are not famous for nothing. As long as you choose our AWS-Certified-Machine-Learning-Specialty study guide, you will find that the exam questions and answers are always the most accurate and up-to-date. It is all due to the hard work of our professionals who always keep a close eye on the updatiog. The AWS-Certified-Machine-Learning-Specialty learning braindumps are regularly updated in line with the changes introduced in the exam contents. You will always find our AWS-Certified-Machine-Learning-Specialty exam simulating highly relevant to your needs.

The Amazon AWS-Certified-Machine-Learning-Specialty Exam consists of 65 multiple-choice and multiple-response questions and

lasts for 180 minutes. AWS-Certified-Machine-Learning-Specialty exam fee is \$300, and it can be taken either at a testing center or online. The passing score for the exam is 750 out of 1000, and the certification is valid for three years. With the increasing demand for machine learning professionals in various industries, the Amazon MLS-C01 certification is a valuable credential to have for those looking to advance their career in this field.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q170-Q175):

NEW QUESTION # 170

A logistics company needs a forecast model to predict next month's inventory requirements for a single item in 10 warehouses. A machine learning specialist uses Amazon Forecast to develop a forecast model from 3 years of monthly data. There is no missing data. The specialist selects the DeepAR+ algorithm to train a predictor. The predictor means absolute percentage error (MAPE) is much larger than the MAPE produced by the current human forecasters.

Which changes to the CreatePredictor API call could improve the MAPE? (Choose two.)

- A. Set FeaturizationMethodName to `filling`.
- B. Set PerformAutoML to `true`.
- C. Set PerformHPO to `true`.
- D. Set ForecastFrequency to `W` for weekly.
- E. Set ForecastHorizon to 4.

Answer: C,D

NEW QUESTION # 171

A data scientist has a dataset of machine part images stored in Amazon Elastic File System (Amazon EFS).

The data scientist needs to use Amazon SageMaker to create and train an image classification machine learning model based on this dataset. Because of budget and time constraints, management wants the data scientist to create and train a model with the least number of steps and integration work required.

How should the data scientist meet these requirements?

- A. Run a SageMaker training job with an EFS file system as the data source.
- B. Mount the EFS file system to an Amazon EC2 instance and use the AWS CLI to copy the data to an Amazon S3 bucket. Run the SageMaker training job with Amazon S3 as the data source.
- C. Launch a transient Amazon EMR cluster. Configure steps to mount the EFS file system and copy the data to an Amazon S3 bucket by using S3DistCp. Run the SageMaker training job with Amazon S3 as the data source.
- D. Mount the EFS file system to a SageMaker notebook and run a script that copies the data to an Amazon FSx for Lustre file system. Run the SageMaker training job with the FSx for Lustre file system as the data source.

Answer: A

Explanation:

Explanation

The simplest and fastest way to use the EFS dataset for SageMaker training is to run a SageMaker training job with an EFS file system as the data source. This option does not require any data copying or additional integration steps. SageMaker supports EFS as a data source for training jobs, and it can mount the EFS file system to the training container using the `FileSystemConfig` parameter. This way, the training script can access the data files as if they were on the local disk of the training instance. References: Access Training Data - Amazon SageMaker

Mount an EFS file system to an Amazon SageMaker notebook (with lifecycle configurations) | AWS Machine Learning Blog

NEW QUESTION # 172

A data scientist uses an Amazon SageMaker notebook instance to conduct data exploration and analysis. This requires certain Python packages that are not natively available on Amazon SageMaker to be installed on the notebook instance.

How can a machine learning specialist ensure that required packages are automatically available on the notebook instance for the data scientist to use?

- A. Create an Amazon SageMaker lifecycle configuration with package installation commands and assign the lifecycle configuration to the notebook instance.
- B. Use the conda package manager from within the Jupyter notebook console to apply the necessary conda packages to the

default kernel of the notebook.

- C. Create a Jupyter notebook file (.ipynb) with cells containing the package installation commands to execute and place the file under the /etc/init directory of each Amazon SageMaker notebook instance.
- D. Install AWS Systems Manager Agent on the underlying Amazon EC2 instance and use Systems Manager Automation to execute the package installation commands.

Answer: A

Explanation:

<https://docs.aws.amazon.com/sagemaker/latest/dg/nbi-add-external.html>

NEW QUESTION # 173

A data scientist at a financial services company used Amazon SageMaker to train and deploy a model that predicts loan defaults. The model analyzes new loan applications and predicts the risk of loan default. To train the model, the data scientist manually extracted loan data from a database. The data scientist performed the model training and deployment steps in a Jupyter notebook that is hosted on SageMaker Studio notebooks.

The model's prediction accuracy is decreasing over time. Which combination of steps in the MOST operationally efficient way for the data scientist to maintain the model's accuracy? (Select TWO.)

- A. Export the training and deployment code from the SageMaker Studio notebooks into a Python script. Package the script into an Amazon Elastic Container Service (Amazon ECS) task that an AWS Lambda function can initiate.
- B. Use SageMaker Pipelines to create an automated workflow that extracts fresh data, trains the model, and deploys a new version of the model.
- C. Rerun the steps in the Jupyter notebook that is hosted on SageMaker Studio notebooks to retrain the model and redeploy a new version of the model.
- D. Store the model predictions in Amazon S3. Create a daily SageMaker Processing job that reads the predictions from Amazon S3, checks for changes in model prediction accuracy, and sends an email notification if a significant change is detected.
- E. Configure SageMaker Model Monitor with an accuracy threshold to check for model drift. Initiate an Amazon CloudWatch alarm when the threshold is exceeded. Connect the workflow in SageMaker Pipelines with the CloudWatch alarm to automatically initiate retraining.

Answer: B,E

Explanation:

* Option A is correct because SageMaker Pipelines is a service that enables you to create and manage automated workflows for your machine learning projects. You can use SageMaker Pipelines to orchestrate the steps of data extraction, model training, and model deployment in a repeatable and scalable way¹.

* Option B is correct because SageMaker Model Monitor is a service that monitors the quality of your models in production and alerts you when there are deviations in the model quality. You can use SageMaker Model Monitor to set an accuracy threshold for your model and configure a CloudWatch alarm that triggers when the threshold is exceeded. You can then connect the alarm to the workflow in SageMaker Pipelines to automatically initiate retraining and deployment of a new version of the model².

* Option C is incorrect because it is not the most operationally efficient way to maintain the model's accuracy. Creating a daily SageMaker Processing job that reads the predictions from Amazon S3 and checks for changes in model prediction accuracy is a manual and time-consuming process. It also requires you to write custom code to perform the data analysis and send the email notification.

Moreover, it does not automatically retrain and deploy the model when the accuracy drops.

* Option D is incorrect because it is not the most operationally efficient way to maintain the model's accuracy. Rerunning the steps in the Jupyter notebook that is hosted on SageMaker Studio notebooks to retrain the model and redeploy a new version of the model is a manual and error-prone process. It also requires you to monitor the model's performance and initiate the retraining and deployment steps yourself. Moreover, it does not leverage the benefits of SageMaker Pipelines and SageMaker Model Monitor to automate and streamline the workflow.

* Option E is incorrect because it is not the most operationally efficient way to maintain the model's accuracy. Exporting the training and deployment code from the SageMaker Studio notebooks into a Python script and packaging the script into an Amazon ECS task that an AWS Lambda function can initiate is a complex and cumbersome process. It also requires you to manage the infrastructure and resources for the Amazon ECS task and the AWS Lambda function. Moreover, it does not leverage the benefits of SageMaker Pipelines and SageMaker Model Monitor to automate and streamline the workflow.

1: SageMaker Pipelines - Amazon SageMaker

2: Monitor data and model quality - Amazon SageMaker

NEW QUESTION # 174

A Marketing Manager at a pet insurance company plans to launch a targeted marketing campaign on social media to acquire new customers. Currently, the company has the following data in Amazon Aurora

- * Profiles for all past and existing customers
- * Profiles for all past and existing insured pets
- * Policy-level information
- * Premiums received
- * Claims paid

What steps should be taken to implement a machine learning model to identify potential new customers on social media?

- A. Use a decision tree classifier engine on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media
- B. Use regression on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media.
- C. Use a recommendation engine on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media
- D. Use clustering on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media.

Answer: D

Explanation:

Clustering is a machine learning technique that can group data points into clusters based on their similarity or proximity. Clustering can help discover the underlying structure and patterns in the data, as well as identify outliers or anomalies. Clustering can also be used for customer segmentation, which is the process of dividing customers into groups based on their characteristics, behaviors, preferences, or needs. Customer segmentation can help understand the key features and needs of different customer segments, as well as design and implement targeted marketing campaigns for each segment. In this case, the Marketing Manager at a pet insurance company plans to launch a targeted marketing campaign on social media to acquire new customers. To do this, the Manager can use clustering on customer profile data to understand the key characteristics of consumer segments, such as their demographics, pet types, policy preferences, premiums paid, claims made, etc. The Manager can then find similar profiles on social media, such as Facebook, Twitter, Instagram, etc., by using the cluster features as filters or keywords. The Manager can then target these potential new customers with personalized and relevant ads or offers that match their segment's needs and interests. This way, the Manager can implement a machine learning model to identify potential new customers on social media.

NEW QUESTION # 175

.....

You can also set the number of Amazon AWS-Certified-Machine-Learning-Specialty dumps questions to attempt in the practice test and time as well. The web-based Amazon AWS-Certified-Machine-Learning-Specialty practice test software needs an active internet connection and can be accessed through all major browsers like Chrome, Edge, Firefox, Opera, and Safari. Our Desktop-based Amazon AWS-Certified-Machine-Learning-Specialty Practice Exam Software is very suitable for those who don't have an internet connection. You can download and install it within a few minutes on Windows-based PCs only and start preparing for the AWS Certified Machine Learning - Specialty exam.

AWS-Certified-Machine-Learning-Specialty Latest Braindumps Free: <https://www.updatedumps.com/Amazon/AWS-Certified-Machine-Learning-Specialty-updated-exam-dumps.html>

- Highly Rated Amazon AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty PDF Dumps Download \Rightarrow AWS-Certified-Machine-Learning-Specialty \Leftarrow for free by simply searching on www.examcollectionpass.com Exam AWS-Certified-Machine-Learning-Specialty Questions
- AWS-Certified-Machine-Learning-Specialty Valid Test Format AWS-Certified-Machine-Learning-Specialty Valid Test Format AWS-Certified-Machine-Learning-Specialty Examcollection Free Dumps Search for \star AWS-Certified-Machine-Learning-Specialty \star and obtain a free download on \blacktriangleright www.pdfvce.com Latest AWS-Certified-Machine-Learning-Specialty Test Blueprint
- AWS-Certified-Machine-Learning-Specialty Test Cram AWS-Certified-Machine-Learning-Specialty Prepaway Dumps Valid AWS-Certified-Machine-Learning-Specialty Real Test Go to website " www.practicevce.com " open and search for " AWS-Certified-Machine-Learning-Specialty " to download for free AWS-Certified-Machine-Learning-Specialty Examinations Actual Questions
- Highly Rated Amazon AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty PDF

Dumps Copy URL { www.pdfvce.com } open and search for **【 AWS-Certified-Machine-Learning-Specialty 】** to download for free **AWS-Certified-Machine-Learning-Specialty Examcollection Free Dumps**

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by UpdateDumps: <https://drive.google.com/open?id=1pFQexP9PUuZv6T68jlr7hjWTyts7-PRg>