Exam AWS-Certified-Machine-Learning-Specialty Tests - Latest AWS-Certified-Machine-Learning-Specialty Braindumps

What's more, part of that VCE4Plus AWS-Certified-Machine-Learning-Specialty dumps now are free: https://drive.google.com/open?id=1JJgmDniGSD3sDMDqo3kt7G7Yg-gH31ML

In the course of your study, the test engine of AWS-Certified-Machine-Learning-Specialty actual exam will be convenient to strengthen the weaknesses in the learning process. This can be used as an alternative to the process of sorting out the wrong questions of AWS-Certified-Machine-Learning-Specialty learning guide in peacetime learning, which not only help you save time, but also makes you more focused in the follow-up learning process with our AWS-Certified-Machine-Learning-Specialty learning materials.

The AWS Certified Machine Learning - Specialty certification exam covers a wide range of topics, including data preparation, model training, model deployment, and machine learning algorithms. AWS-Certified-Machine-Learning-Specialty exam is designed to test the candidate's knowledge of AWS services and their ability to apply machine learning techniques to real-world problems. AWS-Certified-Machine-Learning-Specialty Exam consists of multiple-choice and multiple-response questions, and the candidate has 170 minutes to complete it.

>> Exam AWS-Certified-Machine-Learning-Specialty Tests <<

Latest Amazon AWS-Certified-Machine-Learning-Specialty Braindumps & Valid Dumps AWS-Certified-Machine-Learning-Specialty Ebook

AWS-Certified-Machine-Learning-Specialty test guide is an examination material written by many industry experts based on the examination outlines of the calendar year and industry development trends. Its main purpose is to help students who want to obtain the certification of AWS-Certified-Machine-Learning-Specialty to successfully pass the exam. Compared with other materials available on the market, the main feature of AWS-Certified-Machine-Learning-Specialty exam materials doesn't like other materials simply list knowledge points. It allows students to find time-saving and efficient learning methods while memorizing knowledge points. With AWS-Certified-Machine-Learning-Specialty study braindumps, learning from day and night will never happen. You can learn more with less time. You will become a master of learning in the eyes of others. With AWS-Certified-Machine-Learning-Specialty study braindumps, successfully passing the exam will no longer be a dream.

The AWS Certified Machine Learning - Specialty certification exam covers a wide range of topics, including exploring data, building and training ML models, deploying models, and managing and optimizing ML solutions. AWS-Certified-Machine-Learning-Specialty Exam also tests the candidate's knowledge of machine learning algorithms, deep learning, neural networks, and other related technologies. Passing this certification exam requires a deep understanding of AWS ML services, including Amazon SageMaker, Amazon Rekognition, and Amazon Comprehend.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q261-Q266):

NEW QUESTION # 261
A company is building a demand forecasting model based on machine learning (ML). In the development stage, an ML specialist uses an Amazon SageMaker notebook to perform feature engineering during work hours that consumes low amounts of CPU and memory resources. A data engineer uses the same notebook to perform data preprocessing once a day on average that requires very high memory and completes in only 2 hours. The data preprocessing is not configured to use GPU. All the processes are running well on an ml.m5.4xlarge notebook instance.
The company receives an AWS Budgets alert that the billing for this month exceeds the allocated budget.
Which solution will result in the MOST cost savings?

- A. Change the notebook instance type to a smaller general-purpose instance. Stop the notebook when it is not in use. Run data preprocessing on an ml. r5 instance with the same memory size as the ml.m5.4xlarge instance by using Amazon SageMaker Processing.
- B. Change the notebook instance type to a smaller general-purpose instance. Stop the notebook when it is not in use. Run data preprocessing on an R5 instance with the same memory size as the ml.m5.4xlarge instance by using the Reserved Instance option.
- C. Keep the notebook instance type and size the same. Stop the notebook when it is not in use. Run data preprocessing on a P3 instance type with the same memory as the ml.m5.4xlarge instance by using Amazon SageMaker Processing.
- D. Change the notebook instance type to a memory optimized instance with the same vCPU number as the ml.m5.4xlarge instance has. Stop the notebook when it is not in use. Run both data preprocessing and feature engineering development on that instance.

Answer: A

Explanation:
Explanation
The best solution to reduce the cost of the notebook instance and the data preprocessing job is to change the notebook instance type to a smaller general-purpose instance, stop the notebook when it is not in use, and run data preprocessing on an ml.r5 instance with the same memory size as the ml.m5.4xlarge instance by using Amazon SageMaker Processing. This solution will result in the most cost savings because:
Changing the notebook instance type to a smaller general-purpose instance will reduce the hourly cost of running the notebook, since the feature engineering development does not require high CPU and memory resources. For example, an ml.t3.medium instance costs $0.0464 per hour, while an ml.m5.4xlarge instance costs $0.888 per hour1.
Stopping the notebook when it is not in use will also reduce the cost, since the notebook will only incur charges when it is running. For example, if the notebook is used for 8 hours per day, 5 days per week, then stopping it when it is not in use will save about 76% of the monthly cost compared to leaving it running all the time2.
Running data preprocessing on an ml.r5 instance with the same memory size as the ml.m5.4xlarge instance by using Amazon SageMaker Processing will reduce the cost of the data preprocessing job, since the ml.r5 instance is optimized for memory-intensive workloads and has a lower cost per GB of memory than the ml.m5 instance. For example, an ml.r5.4xlarge instance has 128 GB of memory and costs $1.008 per hour, while an ml.m5.4xlarge instance has 64 GB of memory and costs $0.888 per hour1. Therefore, the ml.r5.4xlarge instance can process the same amount of data in half the time and at a lower cost than the ml.m5.4xlarge instance. Moreover, using Amazon SageMaker Processing will allow the data preprocessing job to run on a separate, fully managed infrastructure that can be scaled up or down as needed, without affecting the notebook instance.
The other options are not as effective as option C for the following reasons:
Option A is not optimal because changing the notebook instance type to a memory optimized instance with the same vCPU number as the ml.m5.4xlarge instance has will not reduce the cost of the notebook, since the memory optimized instances have a higher cost per vCPU than the general-purpose instances. For example, an ml.r5.4xlarge instance has 16 vCPUs and costs $1.008 per hour, while an ml.m5.4xlarge instance has 16 vCPUs and costs $0.888 per hour1. Moreover, running both data preprocessing and feature engineering development on the same instance will not take advantage of the scalability and flexibility of Amazon SageMaker Processing.
Option B is not suitable because running data preprocessing on a P3 instance type with the same memory as the ml.m5.4xlarge instance by using Amazon SageMaker Processing will not reduce the cost of the data preprocessing job, since the P3 instance type is optimized for GPU-based workloads and has a higher cost per GB of memory than the ml.m5 or ml.r5 instance types. For example, an ml.p3.2xlarge instance has 61 GB of memory and costs $3.06 per hour, while an ml.m5.4xlarge instance has 64 GB of memory and costs $0.888 per hour1. Moreover, the data preprocessing job does not require GPU, so using a P3 instance type will be wasteful and inefficient.
Option D is not feasible because running data preprocessing on an R5 instance with the same memory size as the ml.m5.4xlarge instance by using the Reserved Instance option will not reduce the cost of the data preprocessing job, since the Reserved Instance option requires a commitment to a consistent amount of usage for a period of 1 or 3 years3. However, the data preprocessing job only runs once a day on average and completes in only 2 hours, so it does not have a consistent or predictable usage pattern. Therefore, using the Reserved Instance option will not provide any cost savings and may incur additional charges for unused capacity.

References:
Amazon SageMaker Pricing
Manage Notebook Instances - Amazon SageMaker
Amazon EC2 Pricing - Reserved Instances

NEW QUESTION # 262
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. Use the conda package manager from within the Jupyter notebook console to apply the necessary conda packages to the default kernel of the notebook.
- B. Create an Amazon SageMaker lifecycle configuration with package installation commands and assign the lifecycle configuration to the notebook instance.
- 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: B

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

NEW QUESTION # 263
An online delivery company wants to choose the fastest courier for each delivery at the moment an order is placed. The company wants to implement this feature for existing users and new users of its application. Data scientists have trained separate models with XGBoost for this purpose, and the models are stored in Amazon S3. There is one model fof each city where the company operates. The engineers are hosting these models in Amazon EC2 for responding to the web client requests, with one instance for each model, but the instances have only a 5% utilization in CPU and memory, ...operation engineers want to avoid managing unnecessary resources.
Which solution will enable the company to achieve its goal with the LEAST operational overhead?

- A. Keep only a single EC2 instance for hosting all the models. Install a model server in the instance and load each model by pulling it from Amazon S3. Integrate the instance with the web client using Amazon API Gateway for responding to the requests in real time, specifying the target resource according to the city of each request.
- B. Prepare a Docker container based on the prebuilt images in Amazon SageMaker. Replace the existing instances with separate SageMaker endpoints. one for each city where the company operates. Invoke the endpoints from the web client, specifying the URL and EndpomtName parameter according to the city of each request.
- C. Create an Amazon SageMaker notebook instance for pulling all the models from Amazon S3 using the boto3 library. Remove the existing instances and use the notebook to perform a SageMaker batch transform for performing inferences offline for all the possible users in all the cities. Store the results in different files in Amazon S3. Point the web client to the files.
- D. Prepare an Amazon SageMaker Docker container based on the open-source multi-model server. Remove the existing instances and create a multi-model endpoint in SageMaker instead, pointing to the S3 bucket containing all the models Invoke the endpoint from the web client at runtime, specifying the TargetModel parameter according to the city of each request.

Answer: D

Explanation:
The best solution for this scenario is to use a multi-model endpoint in Amazon SageMaker, which allows hosting multiple models on the same endpoint and invoking them dynamically at runtime. This way, the company can reduce the operational overhead of managing multiple EC2 instances and model servers, and leverage the scalability, security, and performance of SageMaker hosting services. By using a multi-model endpoint, the company can also save on hosting costs by improving endpoint utilization and paying only for the models that are loaded in memory and the API calls that are made. To use a multi-model endpoint, the company needs to prepare a Docker container based on the open-source multi-model server, which is a framework-agnostic library that supports loading and serving multiple models from Amazon S3. The company can then create a multi-model endpoint in SageMaker, pointing to the S3 bucket containing all the models, and invoke the endpoint from the web client at runtime, specifying the TargetModel parameter according to the city of each request. This solution also enables the company to add or remove models from the S3

bucket without redeploying the endpoint, and to use different versions of the same model for different cities if needed. References:
Use Docker containers to build models
Host multiple models in one container behind one endpoint
Multi-model endpoints using Scikit Learn
Multi-model endpoints using XGBoost

NEW QUESTION # 264
A company needs to quickly make sense of a large amount of data and gain insight from it. The data is in different formats, the schemas change frequently, and new data sources are added regularly. The company wants to use AWS services to explore multiple data sources, suggest schemas, and enrich and transform the data. The solution should require the least possible coding effort for the data flows and the least possible infrastructure management.
Which combination of AWS services will meet these requirements?

- A. AWS Glue for data discovery, enrichment, and transformation
  Amazon Athena for querying and analyzing the results in Amazon S3 using standard SQL Amazon QuickSight for reporting and getting insights
- B. AWS Data Pipeline for data transferAWS Step Functions for orchestrating AWS Lambda jobs for data discovery, enrichment, and transformation Amazon Athena for querying and analyzing the results in Amazon S3 using standard SQL Amazon QuickSight for reporting and getting insights
- C. Amazon Kinesis Data Analytics for data ingestion
  Amazon EMR for data discovery, enrichment, and transformation
  Amazon Redshift for querying and analyzing the results in Amazon S3
- D. Amazon EMR for data discovery, enrichment, and transformation
  Amazon Athena for querying and analyzing the results in Amazon S3 using standard SQL Amazon QuickSight for reporting and getting insights

Answer: A

Explanation:
The best combination of AWS services to meet the requirements of data discovery, enrichment, transformation, querying, analysis, and reporting with the least coding and infrastructure management is AWS Glue, Amazon Athena, and Amazon QuickSight. These services are:
* AWS Glue for data discovery, enrichment, and transformation. AWS Glue is a serverless data integration service that automatically crawls, catalogs, and prepares data from various sources and formats. It also provides a visual interface called AWS Glue DataBrew that allows users to apply over
250 transformations to clean, normalize, and enrich data without writing code1
* Amazon Athena for querying and analyzing the results in Amazon S3 using standard SQL. Amazon Athena is a serverless interactive query service that allows users to analyze data in Amazon S3 using standard SQL. It supports a variety of data formats, such as CSV, JSON, ORC, Parquet, and Avro. It also integrates with AWS Glue Data Catalog to provide a unified view of the data sources and schemas2
* Amazon QuickSight for reporting and getting insights. Amazon QuickSight is a serverless business intelligence service that allows users to create and share interactive dashboards and reports. It also provides ML-powered features, such as anomaly detection, forecasting, and natural language queries, to help users discover hidden insights from their data3 The other options are not suitable because they either require more coding effort, more infrastructure management, or do not support the desired use cases. For example:
* Option A uses Amazon EMR for data discovery, enrichment, and transformation. Amazon EMR is a managed cluster platform that runs Apache Spark, Apache Hive, and other open-source frameworks for big data processing. It requires users to write code in languages such as Python, Scala, or SQL to perform data integration tasks. It also requires users to provision, configure, and scale the clusters according to their needs4
* Option B uses Amazon Kinesis Data Analytics for data ingestion. Amazon Kinesis Data Analytics is a service that allows users to process streaming data in real time using SQL or Apache Flink. It is not suitable for data discovery, enrichment, and transformation, which are typically batch-oriented tasks. It also requires users to write code to define the data processing logic and the output destination5
* Option D uses AWS Data Pipeline for data transfer and AWS Step Functions for orchestrating AWS Lambda jobs for data discovery, enrichment, and transformation. AWS Data Pipeline is a service that helps users move data between AWS services and on-premises data sources. AWS Step Functions is a service that helps users coordinate multiple AWS services into workflows. AWS Lambda is a service that lets users run code without provisioning or managing servers. These services require users to write code to define the data sources, destinations, transformations, and workflows. They also require users to manage the scalability, performance, and reliability of the data pipelines.
References:

* 1: AWS Glue - Data Integration Service - Amazon Web Services
* 2: Amazon Athena - Interactive SQL Query Service - AWS
* 3: Amazon QuickSight - Business Intelligence Service - AWS
* 4: Amazon EMR - Amazon Web Services
* 5: Amazon Kinesis Data Analytics - Amazon Web Services
* : AWS Data Pipeline - Amazon Web Services
* : AWS Step Functions - Amazon Web Services
* : AWS Lambda - Amazon Web Services

NEW QUESTION # 265
A Machine Learning Specialist is building a model that will perform time series forecasting using Amazon SageMaker The Specialist has finished training the model and is now planning to perform load testing on the endpoint so they can configure Auto Scaling for the model variant Which approach will allow the Specialist to review the latency, memory utilization, and CPU utilization during the load test"?

- A. Review SageMaker logs that have been written to Amazon S3 by leveraging Amazon Athena and Amazon OuickSight to visualize logs as they are being produced
- B. Send Amazon CloudWatch Logs that were generated by Amazon SageMaker lo Amazon ES and use Kibana to query and visualize the log data.
- C. Generate an Amazon CloudWatch dashboard to create a single view for the latency, memory utilization, and CPU utilization metrics that are outputted by Amazon SageMaker
- D. Build custom Amazon CloudWatch Logs and then leverage Amazon ES and Kibana to query and visualize the data as it is generated by Amazon SageMaker

Answer: D

NEW QUESTION # 266
.....

- Pass-Sure Exam AWS-Certified-Machine-Learning-Specialty Tests Supply you Marvelous Latest Braindumps for AWS-Certified-Machine-Learning-Specialty: AWS Certified Machine Learning - Specialty to Prepare casually ➠ www.pdfvce.com is best website to obtain AWS-Certified-Machine-Learning-Specialty for free download New AWS-Certified-Machine-Learning-Specialty Exam Preparation
- Valid Dumps AWS-Certified-Machine-Learning-Specialty Files ✈ AWS-Certified-Machine-Learning-Specialty Advanced Testing Engine AWS-Certified-Machine-Learning-Specialty Test King Go to website ☀ www.pdfdumps.com ☀ open and search for ➡ AWS-Certified-Machine-Learning-Specialty to download for free AWS-Certified-Machine-Learning-Specialty Free Braindumps
- Valid Dumps AWS-Certified-Machine-Learning-Specialty Files New AWS-Certified-Machine-Learning-Specialty Exam Preparation Exam AWS-Certified-Machine-Learning-Specialty Online Search for ➽ AWS-Certified-Machine-Learning-Specialty and download it for free immediately on ▹ www.pdfvce.com ◃ ↪Exam AWS-Certified-Machine-Learning-Specialty Online
- New AWS-Certified-Machine-Learning-Specialty Exam Price AWS-Certified-Machine-Learning-Specialty Test Questions Fee AWS-Certified-Machine-Learning-Specialty Test Questions Fee Easily obtain ➤ AWS-Certified-Machine-Learning-Specialty for free download through [www.exam4labs.com] New AWS-Certified-Machine-Learning-Specialty Exam Preparation
-

P.S. Free 2026 Amazon AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by VCE4Plus: https://drive.google.com/open?id=1JJgmDniGSD3sDMDqo3kt7G7Yg-gH31ML