

AWS-Certified-Machine-Learning-Specialty Exam Dumps Provider & Test AWS-Certified-Machine-Learning-Specialty Cram



What's more, part of that PassReview AWS-Certified-Machine-Learning-Specialty dumps now are free:
<https://drive.google.com/open?id=1fNUussiO5ICIFLsGEZII8p5sGzDusm-I>

Many companies' executives have a job content that purchasing AWS-Certified-Machine-Learning-Specialty valid exam collection PDF help their engineers to pass exam and obtain a useful certificate. It is not only improving the qualification of engineers personal but also showing the qualification of companies. If they choose right AWS-Certified-Machine-Learning-Specialty valid exam collection PDF they will save a lot of exam cost and dumps fee for companies. Our products will be excellent choice with high passing rate.

If you are a positive and optimistic person and want to improve your personal skills, especially for the IT technology, congratulate you, you have found the right place. Amazon exam certification as an important IT certification has attracted many IT candidates. While PassReview AWS-Certified-Machine-Learning-Specialty real test dumps can help you get your goals. The aim of the PassReview is to help all of you pass your test and get your certification. When you visit our website, you will find that we have three different versions for the dumps. Then focusing on the AWS-Certified-Machine-Learning-Specialty free demo, you can free download it for a try. The questions of the free demo are part of the AWS-Certified-Machine-Learning-Specialty complete exam dumps, so if you want the complete one, you will pay for it. What's more, the AWS-Certified-Machine-Learning-Specialty questions are selected and compiled by our professional team with accurate answers which can ensure you 100% pass.

>> AWS-Certified-Machine-Learning-Specialty Exam Dumps Provider <<

Test AWS-Certified-Machine-Learning-Specialty Cram, Valid AWS-Certified-Machine-Learning-Specialty Test Cost

Our goal is to help you save both time and money by providing you with the AWS-Certified-Machine-Learning-Specialty updated exam questions. Keep up the good work on preparing for the Amazon AWS-Certified-Machine-Learning-Specialty test with our actual Amazon AWS-Certified-Machine-Learning-Specialty Dumps. We are so confident that you will succeed on the first try that we will return your money according to the terms and conditions if you do not.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q119-Q124):

NEW QUESTION # 119

A Machine Learning Specialist needs to move and transform data in preparation for training. Some of the data needs to be processed in near-real time and other data can be moved hourly. There are existing Amazon EMR MapReduce jobs to clean and feature engineering to perform on the data. Which of the following services can feed data to the MapReduce jobs? (Select TWO)

- A. Amazon Kinesis
- B. Amazon ES

- C. AWS Data Pipeline
- D. AWS DMS
- E. Amazon Athena

Answer: A,C

NEW QUESTION # 120

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. Amazon EMR for data discovery, enrichment, and transformationAmazon Athena for querying and analyzing the results in Amazon S3 using standard SQLAmazon 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 transformationAmazon Athena for querying and analyzing the results in Amazon S3 using standard SQLAmazon QuickSight for reporting and getting insights
- C. AWS Glue for data discovery, enrichment, and transformationAmazon Athena for querying and analyzing the results in Amazon S3 using standard SQLAmazon QuickSight for reporting and getting insights
- D. Amazon Kinesis Data Analytics for data ingestionAmazon EMR for data discovery, enrichment, and transformationAmazon Redshift for querying and analyzing the results in Amazon S3

Answer: C

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.

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 # 121

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

Answer: B

NEW QUESTION # 122

A machine learning (ML) specialist is using the Amazon SageMaker DeepAR forecasting algorithm to train a model on CPU-based Amazon EC2 On-Demand instances. The model currently takes multiple hours to train. The ML specialist wants to decrease the training time of the model.

Which approaches will meet this requirement⁷ (SELECT TWO)

- A. Use a pre-trained version of the model. Run incremental training.
- B. Use multiple training instances.
- C. Replace CPU-based EC2 instances with GPU-based EC2 instances.
- D. Configure model auto scaling dynamically to adjust the number of instances automatically.
- E. Replace On-Demand Instances with Spot Instances

Answer: B,C

Explanation:

The best approaches to decrease the training time of the model are C and D, because they can improve the computational efficiency and parallelization of the training process. These approaches have the following benefits:

C: Replacing CPU-based EC2 instances with GPU-based EC2 instances can speed up the training of the DeepAR algorithm, as it can leverage the parallel processing power of GPUs to perform matrix operations and gradient computations faster than CPUs¹². The DeepAR algorithm supports GPU-based EC2 instances such as ml.p2 and ml.p3¹³.

D: Using multiple training instances can also reduce the training time of the DeepAR algorithm, as it can distribute the workload across multiple nodes and perform data parallelism⁴. The DeepAR algorithm supports distributed training with multiple CPU-based or GPU-based EC2 instances³.

The other options are not effective or relevant, because they have the following drawbacks:

A: Replacing On-Demand Instances with Spot Instances can reduce the cost of the training, but not necessarily the time, as Spot Instances are subject to interruption and availability⁵. Moreover, the DeepAR algorithm does not support checkpointing, which means that the training cannot resume from the last saved state if the Spot Instance is terminated³.

B: Configuring model auto scaling dynamically to adjust the number of instances automatically is not applicable, as this feature is only available for inference endpoints, not for training jobs⁶.

E: Using a pre-trained version of the model and running incremental training is not possible, as the DeepAR algorithm does not support incremental training or transfer learning³. The DeepAR algorithm requires a full retraining of the model whenever new data is added or the hyperparameters are changed⁷.

References:

- 1: GPU vs CPU: What Matters Most for Machine Learning? | by Louis (What's AI) Bouchard | Towards Data Science
- 2: How GPUs Accelerate Machine Learning Training | NVIDIA Developer Blog
- 3: DeepAR Forecasting Algorithm - Amazon SageMaker
- 4: Distributed Training - Amazon SageMaker
- 5: Managed Spot Training - Amazon SageMaker
- 6: Automatic Scaling - Amazon SageMaker
- 7: How the DeepAR Algorithm Works - Amazon SageMaker

NEW QUESTION # 123

A data engineer at a bank is evaluating a new tabular dataset that includes customer data. The data engineer will use the customer data to create a new model to predict customer behavior. After creating a correlation matrix for the variables, the data engineer notices that many of the 100 features are highly correlated with each other.

Which steps should the data engineer take to address this issue? (Choose two.)

- A. Apply one-hot encoding category-based variables.
- B. Use a linear-based algorithm to train the model.
- C. Remove a portion of highly correlated features from the dataset.
- D. Apply min-max feature scaling to the dataset.
- E. Apply principal component analysis (PCA).

Answer: C,E

Explanation:

* B. Apply principal component analysis (PCA): PCA is a technique that reduces the dimensionality of a dataset by transforming the original features into a smaller set of new features that capture most of the variance in the data. PCA can help address the issue of multicollinearity, which occurs when some features are highly correlated with each other and can cause problems for some machine learning algorithms. By applying PCA, the data engineer can reduce the number of features and remove the redundancy in the data.

* C. Remove a portion of highly correlated features from the dataset: Another way to deal with multicollinearity is to manually remove some of the features that are highly correlated with each other.

This can help simplify the model and avoid overfitting. The data engineer can use the correlation matrix to identify the features that have a high correlation coefficient (e.g., above 0.8 or below -0.8) and remove one of them from the dataset. References: =

* Principal Component Analysis: This is a document from AWS that explains what PCA is, how it works, and how to use it with Amazon SageMaker.

* Multicollinearity: This is a document from AWS that describes what multicollinearity is, how to detect it, and how to deal with it.

NEW QUESTION # 124

.....

Our experts are well-aware of the problems of exam candidates particularly of those who can't manage to spare time to study the AWS-Certified-Machine-Learning-Specialty exam questions due to their heavy work pressure. Hence, our AWS-Certified-Machine-Learning-Specialty study materials have been developed into a simple content and language for our worthy customers all over the world. What is more, you will find there are only the keypoints in our AWS-Certified-Machine-Learning-Specialty learning guide.

Test AWS-Certified-Machine-Learning-Specialty Cram: https://www.passreview.com/AWS-Certified-Machine-Learning-Specialty_exam-braindumps.html

You can completely trust our Amazon AWS-Certified-Machine-Learning-Specialty learning materials, To do this you just need to pass the one Test AWS-Certified-Machine-Learning-Specialty Cram - AWS Certified Machine Learning - Specialty exam, It is unnecessary for you to fail and try more time, and even pay high exam cost once you purchase our AWS-Certified-Machine-Learning-Specialty practice exam materials, This website has got the tools like updated Test AWS-Certified-Machine-Learning-Specialty Cram - AWS Certified Machine Learning - Specialty Amazon Test AWS-Certified-Machine-Learning-Specialty Cram Test AWS-Certified-Machine-Learning-Specialty Cram - AWS Certified Machine Learning - Specialty, PassReview is committed to offering the easiest and simplest way for Amazon AWS-Certified-Machine-Learning-Specialty exam preparation.

Online chat apps including ChatON and Google Hangouts, Drawing Lines and Shapes, You can completely trust our Amazon AWS-Certified-Machine-Learning-Specialty Learning Materials, To do this you just need to pass the one AWS Certified Machine Learning - Specialty exam.

AWS-Certified-Machine-Learning-Specialty Exam Dumps Provider | High-quality Test AWS-Certified-Machine-Learning-Specialty Cram: AWS Certified Machine Learning - Specialty 100% Pass

It is unnecessary for you to fail and try more time, and even pay high exam cost once you purchase our AWS-Certified-Machine-Learning-Specialty practice exam materials, This website has got the tools like updated AWS Certified Machine Learning - Specialty Amazon AWS Certified Machine Learning - Specialty.

PassReview is committed to offering the easiest and simplest way for Amazon AWS-Certified-Machine-Learning-Specialty exam preparation.

myportal.utt.edu.tt, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free & New AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by PassReview:
<https://drive.google.com/open?id=1fNUussiO5ICIFLsGEZII8p5sGzDusm-I>