

New AWS-Certified-Machine-Learning-Specialty Exam Name & AWS-Certified-Machine-Learning-Specialty Latest Braindumps



DOWNLOAD the newest DumpsActual AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free:
<https://drive.google.com/open?id=1m8D4FXGRLZH1JmGrNPQY9AGMQrO73B3V>

Our clients come from all around the world and our company sends the products to them quickly. The clients only need to choose the version of the product, fill in the correct mails and pay for our AWS-Certified-Machine-Learning-Specialty study materials. Then they will receive our mails in 5-10 minutes. Once the clients click on the links they can use our AWS-Certified-Machine-Learning-Specialty Study Materials immediately. If the clients can't receive the mails they can contact our online customer service and they will help them solve the problem. Finally the clients will receive the mails successfully. The purchase procedures are simple and the delivery of our AWS-Certified-Machine-Learning-Specialty study materials is fast.

Amazon MLS-C01 (AWS Certified Machine Learning - Specialty) Exam is a certification exam offered by Amazon Web Services (AWS) for professionals who want to demonstrate their expertise in machine learning on the AWS platform. As artificial intelligence (AI) and machine learning become increasingly important in today's business world, the demand for professionals skilled in these areas is growing rapidly. AWS Certified Machine Learning - Specialty certification demonstrates that an individual has the knowledge and skills necessary to design, implement, deploy, and maintain machine learning solutions on AWS.

>> New AWS-Certified-Machine-Learning-Specialty Exam Name <<

AWS-Certified-Machine-Learning-Specialty Latest Braindumps & AWS-Certified-Machine-Learning-Specialty Reliable Test Testking

AWS-Certified-Machine-Learning-Specialty latest torrents simulate the real exam environment and does not limit the number of computer installations, which can help you better understand the details of the exam. The online version of AWS-Certified-Machine-Learning-Specialty test questions also support multiple devices and can be used offline permanently after being opened for the first time using the network. On buses or subways, you can use fractional time to test your learning outcomes with AWS-Certified-Machine-Learning-Specialty Test Torrent, which will greatly increase your pro forma efficiency.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q149-Q154):

NEW QUESTION # 149

A Data Scientist is developing a binary classifier to predict whether a patient has a particular disease on a series of test results. The Data Scientist has data on 400 patients randomly selected from the population. The disease is seen in 3% of the population. Which cross-validation strategy should the Data Scientist adopt?

- A. A k-fold cross-validation strategy with $k=5$
- B. An 80/20 stratified split between training and validation
- C. A stratified k-fold cross-validation strategy with $k=5$

- D. A k-fold cross-validation strategy with k=5 and 3 repeats

Answer: C

Explanation:

A stratified k-fold cross-validation strategy is a technique that preserves the class distribution in each fold. This is important for imbalanced datasets, such as the one in the question, where the disease is seen in only 3% of the population. If a random k-fold cross-validation strategy is used, some folds may have no positive cases or very few, which would lead to poor estimates of the model performance. A stratified k-fold cross-validation strategy ensures that each fold has the same proportion of positive and negative cases as the whole dataset, which makes the evaluation more reliable and robust. A k-fold cross-validation strategy with k=5 and 3 repeats is also a possible option, but it is more computationally expensive and may not be necessary if the stratification is done properly. An 80/20 stratified split between training and validation is another option, but it uses less data for training and validation than k-fold cross-validation, which may result in higher variance and lower accuracy of the estimates. References:

AWS Machine Learning Specialty Certification Exam Guide

AWS Machine Learning Training: Model Evaluation

How to Fix k-Fold Cross-Validation for Imbalanced Classification

NEW QUESTION # 150

A Machine Learning Specialist must build out a process to query a dataset on Amazon S3 using Amazon Athena. The dataset contains more than 800,000 records stored as plaintext CSV files. Each record contains 200 columns and is approximately 1.5 MB in size. Most queries will span 5 to 10 columns only. How should the Machine Learning Specialist transform the dataset to minimize query runtime?

- A. Convert the records to JSON format
- **B. Convert the records to Apache Parquet format**
- C. Convert the records to XML format
- D. Convert the records to GZIP CSV format

Answer: B

Explanation:

Amazon Athena is an interactive query service that allows you to analyze data stored in Amazon S3 using standard SQL. Athena is serverless, so you only pay for the queries that you run and there is no infrastructure to manage.

To optimize the query performance of Athena, one of the best practices is to convert the data into a columnar format, such as Apache Parquet or Apache ORC. Columnar formats store data by columns rather than by rows, which allows Athena to scan only the columns that are relevant to the query, reducing the amount of data read and improving the query speed. Columnar formats also support compression and encoding schemes that can reduce the storage space and the data scanned per query, further enhancing the performance and reducing the cost.

In contrast, plaintext CSV files store data by rows, which means that Athena has to scan the entire row even if only a few columns are needed for the query. This increases the amount of data read and the query latency.

Moreover, plaintext CSV files do not support compression or encoding, which means that they take up more storage space and incur higher query costs.

Therefore, the Machine Learning Specialist should transform the dataset to Apache Parquet format to minimize query runtime.

Top 10 Performance Tuning Tips for Amazon Athena

Columnar Storage Formats

Using compressions will reduce the amount of data scanned by Amazon Athena, and also reduce your S3 bucket storage. It's a Win-Win for your AWS bill. Supported formats: GZIP, LZO, SNAPPY (Parquet) and ZLIB.

Reference: <https://www.cloudforecast.io/blog/using-parquet-on-athena-to-save-money-on-aws/>

NEW QUESTION # 151

A Machine Learning Specialist is packaging a custom ResNet model into a Docker container so the company can leverage Amazon SageMaker for training. The Specialist is using Amazon EC2 P3 instances to train the model and needs to properly configure the Docker container to leverage the NVIDIA GPUs.

What does the Specialist need to do?

- A. Set the GPU flag in the Amazon SageMaker CreateTrainingJob request body
- B. Organize the Docker container's file structure to execute on GPU instances.
- C. Build the Docker container to be NVIDIA-Docker compatible.
- **D. Bundle the NVIDIA drivers with the Docker image.**

Answer: D

NEW QUESTION # 152

A Machine Learning Specialist is building a prediction model for a large number of features using linear models, such as linear regression and logistic regression. During exploratory data analysis, the Specialist observes that many features are highly correlated with each other. This may make the model unstable. What should be done to reduce the impact of having such a large number of features?

- **A. Create a new feature space using principal component analysis (PCA)**
- B. Perform one-hot encoding on highly correlated features
- C. Apply the Pearson correlation coefficient
- D. Use matrix multiplication on highly correlated features.

Answer: A

Explanation:

Principal component analysis (PCA) is an unsupervised machine learning algorithm that attempts to reduce the dimensionality (number of features) within a dataset while still retaining as much information as possible. This is done by finding a new set of features called components, which are composites of the original features that are uncorrelated with one another. They are also constrained so that the first component accounts for the largest possible variability in the data, the second component the second most variability, and so on. By using PCA, the impact of having a large number of features that are highly correlated with each other can be reduced, as the new feature space will have fewer dimensions and less redundancy. This can make the linear models more stable and less prone to overfitting. References:

Principal Component Analysis (PCA) Algorithm - Amazon SageMaker

Perform a large-scale principal component analysis faster using Amazon SageMaker | AWS Machine Learning Blog
Machine Learning- Principal Component Analysis | i2tutorials

NEW QUESTION # 153

A data scientist uses Amazon SageMaker Data Wrangler to define and perform transformations and feature engineering on historical data. The data scientist saves the transformations to SageMaker Feature Store.

The historical data is periodically uploaded to an Amazon S3 bucket. The data scientist needs to transform the new historic data and add it to the online feature store. The data scientist needs to prepare thehistoric data for training and inference by using native integrations.

Which solution will meet these requirements with the LEAST development effort?

- A. Use AWS Lambda to run a predefined SageMaker pipeline to perform the transformations on each new dataset that arrives in the S3 bucket.
- **B. Configure Amazon EventBridge to run a predefined SageMaker pipeline to perform the transformations when a new data is detected in the S3 bucket.**
- C. Run an AWS Step Functions step and a predefined SageMaker pipeline to perform the transformations on each new dataset that arrives in the S3 bucket.
- D. Use Apache Airflow to orchestrate a set of predefined transformations on each new dataset that arrives in the S3 bucket.

Answer: B

Explanation:

The best solution is to configure Amazon EventBridge to run a predefined SageMaker pipeline to perform the transformations when a new data is detected in the S3 bucket. This solution requires the least development effort because it leverages the native integration between EventBridge and SageMaker Pipelines, which allows you to trigger a pipeline execution based on an event rule.

EventBridge can monitor the S3 bucket for new data uploads and invoke the pipeline that contains the same transformations and feature engineering steps that were defined in SageMaker Data Wrangler. The pipeline can then ingest the transformed data into the online feature store for training and inference.

The other solutions are less optimal because they require more development effort and additional services.

Using AWS Lambda or AWS Step Functions would require writing custom code to invoke the SageMaker pipeline and handle any errors or retries. Using Apache Airflow would require setting up and maintaining an Airflow server and DAGs, as well as integrating with the SageMaker API.

References:

* Amazon EventBridge and Amazon SageMaker Pipelines integration

* Create a pipeline using a JSON specification

* Ingest data into a feature group

NEW QUESTION # 154

• • • • •

As we always want to do better in this career, our research center has formed a group of professional experts responsible for researching new technology of the AWS-Certified-Machine-Learning-Specialty study materials. The technology of the AWS-Certified-Machine-Learning-Specialty practice prep will be innovated every once in a while. As you can see, we never stop innovating new version of the AWS-Certified-Machine-Learning-Specialty Exam Questions. We really need your strong support. We always adopt the kind and useful advices of our loyal customers who wrote to us and gave us their opinions on their study.

AWS-Certified-Machine-Learning-Specialty Latest Braindumps: <https://www.dumpsactual.com/AWS-Certified-Machine-Learning-Specialty-actualtests-dumps.html>

- [illegible]

Disposable vapes

BONUS!!! Download part of DumpsActual AWS-Certified-Machine-Learning-Specialty dumps for free:
<https://drive.google.com/open?id=1m8D4FXGRLZH1JmGrNPQY9AGMQrO73B3V>