

Quiz 2026 Amazon High Hit-Rate AWS-Certified-Machine-Learning-Specialty: Exam AWS Certified Machine Learning - Specialty Details



BONUS!!! Download part of DumpsMaterials AWS-Certified-Machine-Learning-Specialty dumps for free:
https://drive.google.com/open?id=1LnVV2M704ntuPH6eLEr65IZYZ8_c0all

To help people pass exam easily, we bring you the latest AWS-Certified-Machine-Learning-Specialty exam prep for the actual test which enable you get high passing score easily in test. Our study materials are the up-to-dated and all AWS-Certified-Machine-Learning-Specialty Test Answers you practiced are tested by our professional experts. Once you have well prepared with our AWS-Certified-Machine-Learning-Specialty dumps collection, you will go through the formal test without any difficulty.

Being respected and gaining a high social status maybe what you always long for. But if you want to achieve that you must own good abilities and profound knowledge in some certain area. Passing the AWS-Certified-Machine-Learning-Specialty certification can prove that and help you realize your goal and if you buy our AWS-Certified-Machine-Learning-Specialty Quiz prep you will pass the exam successfully. Our product is compiled by experts and approved by professionals with years of experiences. You can download and try out our latest AWS-Certified-Machine-Specialty quiz torrent freely before your purchase.

[**>> Exam AWS-Certified-Machine-Learning-Specialty Details <<**](#)

Amazon AWS-Certified-Machine-Learning-Specialty Exam Questions Come With Free 12 Months Updates

The hit rate for AWS-Certified-Machine-Learning-Specialty exam guide is as high as 99%. Obviously such positive pass rate will establish you confidence as well as strengthen your will to pass your AWS-Certified-Machine-Learning-Specialty exam. No other vendors can challenge our data in this market. At the same time, by studying with our AWS-Certified-Machine-Learning-Specialty practice materials, you avoid wasting your precious time on randomly looking for the key point information. We provide a smooth road for you to success.

Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q87-Q92):

NEW QUESTION # 87

A Machine Learning Specialist has completed a proof of concept for a company using a small data sample and now the Specialist is ready to implement an end-to-end solution in AWS using Amazon SageMaker. The historical training data is stored in Amazon RDS. Which approach should the Specialist use for training a model using that data?

- A. Move the data to Amazon ElastiCache using AWS DMS and set up a connection within the notebook to pull data in for fast access.
- B. Write a direct connection to the SQL database within the notebook and pull data in
- C. Move the data to Amazon DynamoDB and set up a connection to DynamoDB within the notebook to pull data in
- D. Push the data from Microsoft SQL Server to Amazon S3 using an AWS Data Pipeline and provide the S3 location within

the notebook.

Answer: D

NEW QUESTION # 88

A machine learning (ML) engineer has created a feature repository in Amazon SageMaker Feature Store for the company. The company has AWS accounts for development, integration, and production. The company hosts a feature store in the development account. The company uses Amazon S3 buckets to store feature values offline. The company wants to share features and to allow the integration account and the production account to reuse the features that are in the feature repository.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Share the feature repository that is associated with the S3 buckets from the development account to the integration account and the production account by using AWS Resource Access Manager (AWS RAM).
- B. Create an AWS PrivateLink endpoint in the development account for SageMaker.
- C. Set up S3 replication between the development S3 buckets and the integration and production S3 buckets.
- D. Use AWS Security Token Service (AWS STS) from the integration account and the production account to retrieve credentials for the development account.
- E. Create an IAM role in the development account that the integration account and production account can assume. Attach IAM policies to the role that allow access to the feature repository and the S3 buckets.

Answer: A,E

Explanation:

The combination of steps that will meet the requirements are to create an IAM role in the development account that the integration account and production account can assume, attach IAM policies to the role that allow access to the feature repository and the S3 buckets, and share the feature repository that is associated with the S3 buckets from the development account to the integration account and the production account by using AWS Resource Access Manager (AWS RAM). This approach will enable cross-account access and sharing of the features stored in Amazon SageMaker Feature Store and Amazon S3.

Amazon SageMaker Feature Store is a fully managed, purpose-built repository to store, update, search, and share curated data used in training and prediction workflows. The service provides feature management capabilities such as enabling easy feature reuse, low latency serving, time travel, and ensuring consistency between features used in training and inference workflows. A feature group is a logical grouping of ML features whose organization and structure is defined by a feature group schema. A feature group schema consists of a list of feature definitions, each of which specifies the name, type, and metadata of a feature. Amazon SageMaker Feature Store stores the features in both an online store and an offline store. The online store is a low-latency, high-throughput store that is optimized for real-time inference. The offline store is a historical store that is backed by an Amazon S3 bucket and is optimized for batch processing and model training¹.

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization). An IAM role is an IAM identity that you can create in your account that has specific permissions. You can use an IAM role to delegate access to users, applications, or services that don't normally have access to your AWS resources. For example, you can create an IAM role in your development account that allows the integration account and the production account to assume the role and access the resources in the development account. You can attach IAM policies to the role that specify the permissions for the feature repository and the S3 buckets. You can also use IAM conditions to restrict the access based on the source account, IP address, or other factors².

AWS Resource Access Manager (AWS RAM) is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organization. You can share AWS resources that you own with other accounts using resource shares. A resource share is an entity that defines the resources that you want to share, and the principals that you want to share with. For example, you can share the feature repository that is associated with the S3 buckets from the development account to the integration account and the production account by creating a resource share in AWS RAM. You can specify the feature group ARN and the S3 bucket ARN as the resources, and the integration account ID and the production account ID as the principals. You can also use IAM policies to further control the access to the shared resources³.

The other options are either incorrect or unnecessary. Using AWS Security Token Service (AWS STS) from the integration account and the production account to retrieve credentials for the development account is not required, as the IAM role in the development account can provide temporary security credentials for the cross-account access. Setting up S3 replication between the development S3 buckets and the integration and production S3 buckets would introduce redundancy and inconsistency, as the S3 buckets are already shared through AWS RAM. Creating an AWS PrivateLink endpoint in the development account for SageMaker is not relevant, as it is used to securely connect to SageMaker services from a VPC, not from another account.

References:

1: Amazon SageMaker Feature Store - Amazon Web Services

2: What Is IAM? - AWS Identity and Access Management

NEW QUESTION # 89

A manufacturing company needs to identify returned smartphones that have been damaged by moisture. The company has an automated process that produces 2,000 diagnostic values for each phone. The database contains more than five million phone evaluations. The evaluation process is consistent, and there are no missing values in the data. A machine learning (ML) specialist has trained an Amazon SageMaker linear learner ML model to classify phones as moisture damaged or not moisture damaged by using all available features. The model's F1 score is 0.6.

What changes in model training would MOST likely improve the model's F1 score? (Select TWO.)

- A. Continue to use the SageMaker linear learner algorithm. Reduce the number of features with the scikit-learn multi-dimensional scaling (MDS) algorithm.
- B. Continue to use the SageMaker linear learner algorithm. Set the predictor type to regressor.
- C. Continue to use the SageMaker linear learner algorithm. Reduce the number of features with the SageMaker principal component analysis (PCA) algorithm.
- D. Use the SageMaker k-means algorithm with k of less than 1,000 to train the model
- E. Use the SageMaker k-nearest neighbors (k-NN) algorithm. Set a dimension reduction target of less than 1,000 to train the model.

Answer: C,E

Explanation:

- * Option A is correct because reducing the number of features with the SageMaker PCA algorithm can help remove noise and redundancy from the data, and improve the model's performance. PCA is a dimensionality reduction technique that transforms the original features into a smaller set of linearly uncorrelated features called principal components. The SageMaker linear learner algorithm supports PCA as a built-in feature transformation option.
- * Option E is correct because using the SageMaker k-NN algorithm with a dimension reduction target of less than 1,000 can help the model learn from the similarity of the data points, and improve the model's performance. k-NN is a non-parametric algorithm that classifies an input based on the majority vote of its k nearest neighbors in the feature space. The SageMaker k-NN algorithm supports dimension reduction as a built-in feature transformation option.
- * Option B is incorrect because using the scikit-learn MDS algorithm to reduce the number of features is not a feasible option, as MDS is a computationally expensive technique that does not scale well to large datasets. MDS is a dimensionality reduction technique that tries to preserve the pairwise distances between the original data points in a lower-dimensional space.
- * Option C is incorrect because setting the predictor type to regressor would change the model's objective from classification to regression, which is not suitable for the given problem. A regressor model would output a continuous value instead of a binary label for each phone.
- * Option D is incorrect because using the SageMaker k-means algorithm with k of less than 1,000 would not help the model classify the phones, as k-means is a clustering algorithm that groups the data points into k clusters based on their similarity, without using any labels. A clustering model would not output a binary label for each phone.

Amazon SageMaker Linear Learner Algorithm

Amazon SageMaker K-Nearest Neighbors (k-NN) Algorithm

[Principal Component Analysis - Scikit-learn]

[Multidimensional Scaling - Scikit-learn]

NEW QUESTION # 90

A Machine Learning Specialist built an image classification deep learning model. However, the Specialist ran into an overfitting problem in which the training and testing accuracies were 99% and 75%, respectively.

How should the Specialist address this issue and what is the reason behind it?

- A. The learning rate should be increased because the optimization process was trapped at a local minimum.
- B. The epoch number should be increased because the optimization process was terminated before it reached the global minimum.
- C. The dimensionality of dense layer next to the flatten layer should be increased because the model is not complex enough.
- D. The dropout rate at the flatten layer should be increased because the model is not generalized enough.

Answer: D

Explanation:

<https://kharshit.github.io/blog/2018/05/04/dropout-prevent-overfitting>

NEW QUESTION # 91

A Machine Learning Specialist works for a credit card processing company and needs to predict which transactions may be fraudulent in near-real time. Specifically, the Specialist must train a model that returns the probability that a given transaction may be fraudulent.

How should the Specialist frame this business problem?

- A. Regression classification
- B. Streaming classification
- C. Multi-category classification
- D. **Binary classification**

Answer: D

Explanation:

The business problem of predicting whether a new credit card applicant will default on a credit card payment can be framed as a binary classification problem. Binary classification is the task of predicting a discrete class label output for an example, where the class label can only take one of two possible values. In this case, the class label can be either "default" or "no default", indicating whether the applicant will or will not default on a credit card payment. A binary classification model can return the probability that a given applicant belongs to each class, and then assign the applicant to the class with the highest probability. For example, if the model predicts that an applicant has a 0.8 probability of defaulting and a 0.2 probability of not defaulting, then the model will classify the applicant as "default". Binary classification is suitable for this problem because the outcome of interest is categorical and binary, and the model needs to return the probability of each outcome.

AWS Machine Learning Specialty Exam Guide

AWS Machine Learning Training - Classification vs Regression in Machine Learning

NEW QUESTION # 92

.....

You will need to pass the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam to achieve the AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) certification. Due to extremely high competition, passing the Amazon AWS-Certified-Machine-Learning-Specialty exam is not easy; however, possible. You can use DumpsMaterials products to pass the AWS-Certified-Machine-Learning-Specialty Exam on the first attempt. The AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice exam gives you confidence and helps you understand the criteria of the testing authority and pass the Amazon AWS-Certified-Machine-Learning-Specialty exam on the first attempt.

AWS-Certified-Machine-Learning-Specialty Mock Exam: <https://www.dumpsmaterials.com/AWS-Certified-Machine-Learning-Specialty-real-torrent.html>

Amazon AWS-Certified-Machine-Learning-Specialty Mock Exam recently announced they will change the exam structure from the end of November, so they will focus a bit more on IaaS and the new ARM rather than old PaaS and classic portal, Amazon Exam AWS-Certified-Machine-Learning-Specialty Details Updates with development, We always have one-hand news from official exam center, and then our professional experts pay in the effort on AWS-Certified-Machine-Learning-Specialty actual test dumps day and night so that we can provide the best test VCE engine as soon as possible, AWS-Certified-Machine-Learning-Specialty certification not only improves your personal ability but also helps you be a senior leader in your company successfully.

You can experience the effects of outside products in advance by downloading clue versions of our AWS-Certified-Machine-Learning-Specialty exam torrent, The global area is consistent across the portal and often provides AWS-Certified-Machine-Learning-Specialty Exam Braindumps links to a home page, contact information, accessories, or other frequently used applications.

With Amazon AWS-Certified-Machine-Learning-Specialty Real Questions Nothing Can Stop You from Getting Success

Amazon recently announced they will change the exam structure AWS-Certified-Machine-Learning-Specialty from the end of November, so they will focus a bit more on IaaS and the new ARM rather than old PaaS and classic portal.

Updates with development, We always have one-hand AWS-Certified-Machine-Learning-Specialty Exam Braindumps news from official exam center, and then our professional experts pay in the effort on AWS-Certified-Machine-Learning-Specialty actual test

dumps day and night so that we can provide the best test VCE engine as soon as possible.

AWS-Certified-Machine-Learning-Specialty certification not only improves your personal ability but also helps you be a senior leader in your company successfully. We believe all candidates who purchase our practice AWS-Certified-Machine-Learning-Specialty exam dumps pdf and practice exam dumps online can pass exam.

DOWNLOAD the newest DumpsMaterials AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1LnVV2M704ntuPH6eLEr65IZYZ8_c0aIL