

# Amazon AWS-Certified-Machine-Learning-Specialty Exam | Test AWS-Certified-Machine-Learning-Specialty Result - Test Engine Simulation of Latest AWS-Certified-Machine-Learning-Specialty Exam Duration



BONUS!!! Download part of Prep4sureGuide AWS-Certified-Machine-Learning-Specialty dumps for free:  
[https://drive.google.com/open?id=1273j-y8VdDarPa6KoI16eCmJJdNR\\_tik](https://drive.google.com/open?id=1273j-y8VdDarPa6KoI16eCmJJdNR_tik)

There may be a lot of people feel that the preparation process for exams is hard and boring, and hard work does not necessarily mean good results, which is an important reason why many people are afraid of examinations. Today, our AWS-Certified-Machine-Learning-Specialty study materials will radically change this. High question hit rate makes you no longer aimless when preparing for the exam, so you just should review according to the content of our AWS-Certified-Machine-Learning-Specialty Study Materials prepared for you. Instant answer feedback allows you to identify your vulnerabilities in a timely manner, so as to make up for your weaknesses.

The AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty certification is a unique way to level up your knowledge and skills. With the AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty credential, you become eligible to get high-paying jobs in the constantly advancing tech sector. Success in the Amazon AWS-Certified-Machine-Learning-Specialty examination also boosts your skills to land promotions within your current organization. Are you looking for a simple and quick way to crack the Amazon AWS-Certified-Machine-Learning-Specialty examination? If you are, then rely on AWS-Certified-Machine-Learning-Specialty Exam Dumps.

>> **Test AWS-Certified-Machine-Learning-Specialty Result** <<

## Latest AWS-Certified-Machine-Learning-Specialty Exam Duration - Exam AWS-Certified-Machine-Learning-Specialty Labs

In addition to the advantages of high quality, our AWS-Certified-Machine-Learning-Specialty study materials also provide various versions. In order to meet your personal habits, you can freely choose any version within PDF, APP or PC version. Among them, the PDF version is most suitable for candidates who prefer paper materials, because it supports printing. If you want to use our AWS-Certified-Machine-Learning-Specialty Study Materials on your phone at any time, then APP version is your best choice as long as you have browsers on your phone.

## Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q104-Q109):

### NEW QUESTION # 104

A company's Machine Learning Specialist needs to improve the training speed of a time-series forecasting model using TensorFlow. The training is currently implemented on a single-GPU machine and takes approximately 23 hours to complete. The training needs to be run daily.

The model accuracy is acceptable, but the company anticipates a continuous increase in the size of the training data and a need to

update the model on an hourly, rather than a daily, basis. The company also wants to minimize coding effort and infrastructure changes.

What should the Machine Learning Specialist do to the training solution to allow it to scale for future demand?

- **A. Change the TensorFlow code to implement a Horovod distributed framework supported by Amazon SageMaker. Parallelize the training to as many machines as needed to achieve the business goals.**
- B. Do not change the TensorFlow code. Change the machine to one with a more powerful GPU to speed up the training.
- C. Move the training to Amazon EMR and distribute the workload to as many machines as needed to achieve the business goals.
- D. Switch to using a built-in AWS SageMaker DeepAR model. Parallelize the training to as many machines as needed to achieve the business goals.

**Answer: A**

Explanation:  
Explanation

### NEW QUESTION # 105

A beauty supply store wants to understand some characteristics of visitors to the store. The store has security video recordings from the past several years. The store wants to generate a report of hourly visitors from the recordings. The report should group visitors by hair style and hair color.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Use an object detection algorithm to identify a visitor's hair in video frames. Pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color.
- B. Use a semantic segmentation algorithm to identify a visitor's hair in video frames. Pass the identified hair to an XGBoost algorithm to determine hair style and hair.
- **C. Use a semantic segmentation algorithm to identify a visitor's hair in video frames. Pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color.**
- D. Use an object detection algorithm to identify a visitor's hair in video frames. Pass the identified hair to an XGBoost algorithm to determine hair style and hair color.

**Answer: C**

Explanation:

The solution that will meet the requirements with the least amount of effort is to use a semantic segmentation algorithm to identify a visitor's hair in video frames, and pass the identified hair to an ResNet-50 algorithm to determine hair style and hair color. This solution can leverage the existing Amazon SageMaker algorithms and frameworks to perform the tasks of hair segmentation and classification.

Semantic segmentation is a computer vision technique that assigns a class label to every pixel in an image, such that pixels with the same label share certain characteristics. Semantic segmentation can be used to identify and isolate different objects or regions in an image, such as a visitor's hair in a video frame. Amazon SageMaker provides a built-in semantic segmentation algorithm that can train and deploy models for semantic segmentation tasks. The algorithm supports three state-of-the-art network architectures: Fully Convolutional Network (FCN), Pyramid Scene Parsing Network (PSP), and DeepLab v3. The algorithm can also use pre-trained or randomly initialized ResNet-50 or ResNet-101 as the backbone network. The algorithm can be trained using P2/P3 type Amazon EC2 instances in single machine configurations<sup>1</sup>.

ResNet-50 is a convolutional neural network that is 50 layers deep and can classify images into 1000 object categories. ResNet-50 is trained on more than a million images from the ImageNet database and can achieve high accuracy on various image recognition tasks. ResNet-50 can be used to determine hair style and hair color from the segmented hair regions in the video frames. Amazon SageMaker provides a built-in image classification algorithm that can use ResNet-50 as the network architecture. The algorithm can also perform transfer learning by fine-tuning the pre-trained ResNet-50 model with new data. The algorithm can be trained using P2/P3 type Amazon EC2 instances in single or multiple machine configurations<sup>2</sup>.

The other options are either less effective or more complex to implement. Using an object detection algorithm to identify a visitor's hair in video frames would not segment the hair at the pixel level, but only draw bounding boxes around the hair regions. This could result in inaccurate or incomplete hair segmentation, especially if the hair is occluded or has irregular shapes. Using an XGBoost algorithm to determine hair style and hair color would require transforming the segmented hair images into numerical features, which could lose some information or introduce noise. XGBoost is also not designed for image classification tasks, and may not achieve high accuracy or performance.

1: Semantic Segmentation Algorithm - Amazon SageMaker

2: Image Classification Algorithm - Amazon SageMaker

### NEW QUESTION # 106

A Data Scientist received a set of insurance records, each consisting of a record ID, the final outcome among 200 categories, and the date of the final outcome. Some partial information on claim contents is also provided, but only for a few of the 200 categories. For each outcome category, there are hundreds of records distributed over the past 3 years. The Data Scientist wants to predict how many claims to expect in each category from month to month, a few months in advance. What type of machine learning model should be used?

- A. Reinforcement learning using claim IDs and timestamps where the agent will identify how many claims in each category to expect from month to month.
- B. Classification with supervised learning of the categories for which partial information on claim contents is provided, and forecasting using claim IDs and timestamps for all other categories.
- **C. Forecasting using claim IDs and timestamps to identify how many claims in each category to expect from month to month.**
- D. Classification month-to-month using supervised learning of the 200 categories based on claim contents.

**Answer: C**

### NEW QUESTION # 107

A data scientist is developing a pipeline to ingest streaming web traffic data. The data scientist needs to implement a process to identify unusual web traffic patterns as part of the pipeline. The patterns will be used downstream for alerting and incident response. The data scientist has access to unlabeled historic data to use, if needed.

The solution needs to do the following:

Calculate an anomaly score for each web traffic entry.

Adapt unusual event identification to changing web patterns over time.

Which approach should the data scientist implement to meet these requirements?

- **A. Collect the streaming data using Amazon Kinesis Data Firehose. Map the delivery stream as an input source for Amazon Kinesis Data Analytics. Write a SQL query to run in real time against the streaming data with the Amazon Random Cut Forest (RCF) SQL extension to calculate anomaly scores for each record using a sliding window.**
- B. Use historic web traffic data to train an anomaly detection model using the Amazon SageMaker Random Cut Forest (RCF) built-in model. Use an Amazon Kinesis Data Stream to process the incoming web traffic data. Attach a preprocessing AWS Lambda function to perform data enrichment by calling the RCF model to calculate the anomaly score for each record.
- C. Collect the streaming data using Amazon Kinesis Data Firehose. Map the delivery stream as an input source for Amazon Kinesis Data Analytics. Write a SQL query to run in real time against the streaming data with the k-Nearest Neighbors (kNN) SQL extension to calculate anomaly scores for each record using a tumbling window.
- D. Use historic web traffic data to train an anomaly detection model using the Amazon SageMaker built-in XGBoost model. Use an Amazon Kinesis Data Stream to process the incoming web traffic data. Attach a preprocessing AWS Lambda function to perform data enrichment by calling the XGBoost model to calculate the anomaly score for each record.

**Answer: A**

Explanation:

Explanation

Amazon Kinesis Data Analytics is a service that allows users to analyze streaming data in real time using SQL queries. Amazon Random Cut Forest (RCF) is a SQL extension that enables anomaly detection on streaming data. RCF is an unsupervised machine learning algorithm that assigns an anomaly score to each data point based on how different it is from the rest of the data. A sliding window is a type of window that moves along with the data stream, so that the anomaly detection model can adapt to changing patterns over time. A tumbling window is a type of window that has a fixed size and does not overlap with other windows, so that the anomaly detection model is based on a fixed period of time. Therefore, option D is the best approach to meet the requirements of the question, as it uses RCF to calculate anomaly scores for each web traffic entry and uses a sliding window to adapt to changing web patterns over time.

Option A is incorrect because Amazon SageMaker Random Cut Forest (RCF) is a built-in model that can be used to train and deploy anomaly detection models on batch or streaming data, but it requires more steps and resources than using the RCF SQL extension in Amazon Kinesis Data Analytics. Option B is incorrect because Amazon SageMaker XGBoost is a built-in model that can be used for supervised learning tasks such as classification and regression, but not for unsupervised learning tasks such as anomaly detection. Option C is incorrect because k-Nearest Neighbors (kNN) is a SQL extension that can be used for classification and regression tasks on streaming data, but not for anomaly detection. Moreover, using a tumbling window would not allow the anomaly detection model to adapt to changing web patterns over time.

References:

Using CloudWatch anomaly detection  
Anomaly Detection With CloudWatch  
Performing Real-time Anomaly Detection using AWS  
What Is AWS Anomaly Detection? (And Is There A Better Option?)

### NEW QUESTION # 108

A real-estate company is launching a new product that predicts the prices of new houses. The historical data for the properties and prices is stored in .csv format in an Amazon S3 bucket. The data has a header, some categorical fields, and some missing values. The company's data scientists have used Python with a common open-source library to fill the missing values with zeros. The data scientists have dropped all of the categorical fields and have trained a model by using the open-source linear regression algorithm with the default parameters.

The accuracy of the predictions with the current model is below 50%. The company wants to improve the model performance and launch the new product as soon as possible.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM role for Amazon SageMaker with access to the S3 bucket. Create a SageMaker AutoML job with SageMaker Autopilot pointing to the bucket with the dataset. Specify the price as the target attribute. Wait for the job to complete. Deploy the best model for predictions.
- B. Create an Amazon SageMaker notebook with a new IAM role that is associated with the notebook. Pull the dataset from the S3 bucket. Explore different combinations of feature engineering transformations, regression algorithms, and hyperparameters. Compare all the results in the notebook, and deploy the most accurate configuration in an endpoint for predictions.
- C. Create an IAM role with access to Amazon S3, Amazon SageMaker, and AWS Lambda. Create a training job with the SageMaker built-in XGBoost model pointing to the bucket with the dataset. Specify the price as the target feature. Wait for the job to complete. Load the model artifact to a Lambda function for inference on prices of new houses.
- D. Create a service-linked role for Amazon Elastic Container Service (Amazon ECS) with access to the S3 bucket. Create an ECS cluster that is based on an AWS Deep Learning Containers image. Write the code to perform the feature engineering. Train a logistic regression model for predicting the price, pointing to the bucket with the dataset. Wait for the training job to complete. Perform the inferences.

**Answer: A**

Explanation:  
Explanation

The solution D meets the requirements with the least operational overhead because it uses Amazon SageMaker Autopilot, which is a fully managed service that automates the end-to-end process of building, training, and deploying machine learning models. Amazon SageMaker Autopilot can handle data preprocessing, feature engineering, algorithm selection, hyperparameter tuning, and model deployment. The company only needs to create an IAM role for Amazon SageMaker with access to the S3 bucket, create a SageMaker AutoML job pointing to the bucket with the dataset, specify the price as the target attribute, and wait for the job to complete. Amazon SageMaker Autopilot will generate a list of candidate models with different configurations and performance metrics, and the company can deploy the best model for predictions1.

The other options are not suitable because:

Option A: Creating a service-linked role for Amazon Elastic Container Service (Amazon ECS) with access to the S3 bucket, creating an ECS cluster based on an AWS Deep Learning Containers image, writing the code to perform the feature engineering, training a logistic regression model for predicting the price, and performing the inferences will incur more operational overhead than using Amazon SageMaker Autopilot. The company will have to manage the ECS cluster, the container image, the code, the model, and the inference endpoint. Moreover, logistic regression may not be the best algorithm for predicting the price, as it is more suitable for binary classification tasks2.

Option B: Creating an Amazon SageMaker notebook with a new IAM role that is associated with the notebook, pulling the dataset from the S3 bucket, exploring different combinations of feature engineering transformations, regression algorithms, and hyperparameters, comparing all the results in the notebook, and deploying the most accurate configuration in an endpoint for predictions will incur more operational overhead than using Amazon SageMaker Autopilot. The company will have to write the code for the feature engineering, the model training, the model evaluation, and the model deployment. The company will also have to manually compare the results and select the best configuration3.

Option C: Creating an IAM role with access to Amazon S3, Amazon SageMaker, and AWS Lambda, creating a training job with the SageMaker built-in XGBoost model pointing to the bucket with the dataset, specifying the price as the target feature, loading the model artifact to a Lambda function for inference on prices of new houses will incur more operational overhead than using Amazon SageMaker Autopilot. The company will have to create and manage the Lambda function, the model artifact, and the inference endpoint. Moreover, XGBoost may not be the best algorithm for predicting the price, as it is more suitable for classification and

ranking tasks4.

References:

- 1: Amazon SageMaker Autopilot
- 2: Amazon Elastic Container Service
- 3: Amazon SageMaker Notebook Instances
- 4: Amazon SageMaker XGBoost Algorithm

## NEW QUESTION # 109

.....

The AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) exam questions are the real, valid, and updated AWS-Certified-Machine-Learning-Specialty Exam Questions that are specifically designed for quick and complete AWS-Certified-Machine-Learning-Specialty exam preparation. With Prep4sureGuide AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice test questions you can start Amazon AWS-Certified-Machine-Learning-Specialty exam preparation immediately.

**Latest AWS-Certified-Machine-Learning-Specialty Exam Duration:** <https://www.prep4sureguide.com/AWS-Certified-Machine-Learning-Specialty-prep4sure-exam-guide.html>

Money Back Guarantee GuaranteePrep4sureGuide Latest AWS-Certified-Machine-Learning-Specialty Exam Duration provides hassle-free money back guarantee with our products, If you come across questions about our AWS-Certified-Machine-Learning-Specialty training materials, you can browser the module, AWS-Certified-Machine-Learning-Specialty exam dumps not only have the quality, but also have the quantity, and it's enough for you to practice, If you want to pass your exam and get your certification, we can make sure that our Latest AWS-Certified-Machine-Learning-Specialty Exam Duration guide questions will be your ideal choice.

Windows Sound Recorder, Duke Professor of Biology emeritus at AWS-Certified-Machine-Learning-Specialty Duke University, Money Back Guarantee GuaranteePrep4sureGuide provides hassle-free money back guarantee with our products.

If you come across questions about our AWS-Certified-Machine-Learning-Specialty Training Materials, you can browser the module, AWS-Certified-Machine-Learning-Specialty exam dumps not only have the quality, but also have the quantity, and it's enough for you to practice.

## Test AWS-Certified-Machine-Learning-Specialty Result & Free PDF Amazon Realistic Latest AWS Certified Machine Learning - Specialty Exam Duration

If you want to pass your exam and get your Latest AWS-Certified-Machine-Learning-Specialty Exam Duration certification, we can make sure that our AWS Certified Machine Learning guide questions will be your ideal choice, If you decide to refund, please AWS-Certified-Machine-Learning-Specialty Accurate Answers send the score report to us, we will get your money back after confirmation.

- Latest AWS-Certified-Machine-Learning-Specialty exam pdf, valid Amazon AWS-Certified-Machine-Learning-Specialty questions, AWS-Certified-Machine-Learning-Specialty free demo  Download ➔ AWS-Certified-Machine-Learning-Specialty  for free by simply searching on  [www.pass4test.com](http://www.pass4test.com)   New AWS-Certified-Machine-Learning-Specialty Exam Discount
- New AWS-Certified-Machine-Learning-Specialty Exam Objectives  Updated AWS-Certified-Machine-Learning-Specialty Testkings  Exam AWS-Certified-Machine-Learning-Specialty Study Solutions  Easily obtain free download of  AWS-Certified-Machine-Learning-Specialty  by searching on ▶ [www.pdfvce.com](http://www.pdfvce.com) ◀  AWS-Certified-Machine-Learning-Specialty Examcollection Dumps Torrent
- Pass Guaranteed 2026 Amazon Accurate Test AWS-Certified-Machine-Learning-Specialty Result  Easily obtain free download of ▶ AWS-Certified-Machine-Learning-Specialty ◀ by searching on ☀: [www.testkingpass.com](http://www.testkingpass.com) ☀   Valid Braindumps AWS-Certified-Machine-Learning-Specialty Questions
- Latest AWS-Certified-Machine-Learning-Specialty Quiz Prep Aim at Assisting You to Pass the AWS-Certified-Machine-Learning-Specialty Exam - Pdfvce  Search on 「 [www.pdfvce.com](http://www.pdfvce.com) 」 for  AWS-Certified-Machine-Learning-Specialty  to obtain exam materials for free download  AWS-Certified-Machine-Learning-Specialty Valid Exam Tips
- Amazon AWS-Certified-Machine-Learning-Specialty Exam | Test AWS-Certified-Machine-Learning-Specialty Result - Offer you Valid Latest AWS-Certified-Machine-Learning-Specialty Exam Duration  The page for free download of ➔ AWS-Certified-Machine-Learning-Specialty  on 《 [www.practicevce.com](http://www.practicevce.com) 》 will open immediately  New AWS-

