

# **Precise AWS-Certified-Machine-Learning-Specialty Exam Questions Answers bring you First-Grade AWS-Certified-Machine-Learning-Specialty Exam Fees for Amazon AWS Certified Machine Learning - Specialty**



What's more, part of that PassLeader AWS-Certified-Machine-Learning-Specialty dumps now are free:  
<https://drive.google.com/open?id=1jU7cmZvulXiMXnx4BFYRLqKqb0u-6Yio>

You can customize the time and Amazon AWS-Certified-Machine-Learning-Specialty questions of our AWS Certified Machine Learning - Specialty (AWS-Certified-Machine-Learning-Specialty) practice exams according to your needs. Real Amazon AWS-Certified-Machine-Learning-Specialty exam environment which our web-based and desktop AWS-Certified-Machine-Learning-Specialty Practice Exams create is beneficial to get accustomed to the real AWS-Certified-Machine-Learning-Specialty exam pattern.

To take the Amazon MLS-C01 exam, candidates must first pass the AWS Certified Cloud Practitioner or AWS Certified Solutions Architect Associate exam. AWS-Certified-Machine-Learning-Specialty Exam consists of 65 multiple-choice and multiple-response questions and must be completed in 180 minutes. AWS-Certified-Machine-Learning-Specialty exam is available in several languages, including English, Japanese, Korean, and Simplified Chinese. Upon passing the exam, candidates will receive an AWS Certified Machine Learning – Specialty certification, which is valid for three years. AWS Certified Machine Learning - Specialty certification demonstrates to employers and clients that the individual has the skills and knowledge needed to design, implement, and maintain machine learning solutions on the AWS platform.

## **Career Path**

In case you want to specialize in more specific AWS services, you can opt for other Amazon specialty certifications like the AWS Certified Advanced Networking Specialty, the AWS Certified Alexa Skill Builder Specialty, or the AWS Certified Database Specialty if to name a few.

## **Read the registration procedure of the AWS Certified Machine Learning Specialty Exam**

In order to apply for the AWS Certified Machine Learning Specialty, You have to follow these steps

- Step 1: Sign in to AWS Training
- Step 2: Click Certification in the top navigation
- Step 3: Click AWS Certification Account Button
- Step 4: Followed by Schedule New Exam
- Step 5: Search the AWS Certified Machine Learning Specialty exam
- Step 6: Click either the Schedule at PSI or Schedule at Pearson VUE button
- Step 7: Select Date, time and Schedule your test

>> AWS-Certified-Machine-Learning-Specialty Exam Questions Answers <<

## AWS-Certified-Machine-Learning-Specialty Exam Fees - AWS-Certified-Machine-Learning-Specialty Latest Exam

Many learners feel that they have choice phobia disorder while they are choosing reliable AWS-Certified-Machine-Learning-Specialty test guide on the internet. If so you can choose our AWS-Certified-Machine-Learning-Specialty certification materials. We are the leading position in this field and our company is growing faster and faster because of our professional and high pass-rate AWS-Certified-Machine-Learning-Specialty Exam Torrent materials. Every year more than thousands of candidates choose our reliable AWS-Certified-Machine-Learning-Specialty test guide materials we help more than 98% of candidates clear exams, we are proud of our AWS-Certified-Machine-Learning-Specialty exam questions.

## Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q29-Q34):

### NEW QUESTION # 29

A Data Scientist is building a model to predict customer churn using a dataset of 100 continuous numerical features. The Marketing team has not provided any insight about which features are relevant for churn prediction. The Marketing team wants to interpret the model and see the direct impact of relevant features on the model outcome. While training a logistic regression model, the Data Scientist observes that there is a wide gap between the training and validation set accuracy. Which methods can the Data Scientist use to improve the model performance and satisfy the Marketing team's needs? (Choose two.)

- A. Add L1 regularization to the classifier
- B. Perform recursive feature elimination
- C. Perform t-distributed stochastic neighbor embedding (t-SNE)
- D. Perform linear discriminant analysis
- E. Add features to the dataset

**Answer: A,B**

Explanation:

\* The Data Scientist is building a model to predict customer churn using a dataset of 100 continuous numerical features. The Marketing team wants to interpret the model and see the direct impact of relevant features on the model outcome. However, the Data Scientist observes that there is a wide gap between the training and validation set accuracy, which indicates that the model is overfitting the data and generalizing poorly to new data.

\* To improve the model performance and satisfy the Marketing team's needs, the Data Scientist can use the following methods:

\* Add L1 regularization to the classifier: L1 regularization is a technique that adds a penalty term to the loss function of the logistic regression model, proportional to the sum of the absolute values of the coefficients. L1 regularization can help reduce overfitting by shrinking the coefficients of the less important features to zero, effectively performing feature selection. This can simplify the model and make it more interpretable, as well as improve the validation accuracy.

\* Perform recursive feature elimination: Recursive feature elimination (RFE) is a feature selection technique that involves training a model on a subset of the features, and then iteratively removing the least important features one by one until the desired number of features is reached. The idea behind RFE is to determine the contribution of each feature to the model by measuring how well the model performs when that feature is removed. The features that are most important to the model will have the greatest impact on performance when they are removed. RFE can help improve the model performance by eliminating the irrelevant or redundant features that may cause noise or multicollinearity in the data. RFE can also help the Marketing team understand the direct impact of the relevant features on the model outcome, as the remaining features will have the highest weights in the model.

Regularization for Logistic Regression

### NEW QUESTION # 30

A Machine Learning Specialist is working with a large cybersecurity company that manages security events in real time for companies around the world. The cybersecurity company wants to design a solution that will allow it to use machine learning to score malicious events as anomalies on the data as it is being ingested. The company also wants to be able to save the results in its data lake for later processing and analysis. What is the MOST efficient way to accomplish these tasks'?

- A. Ingest the data into Apache Spark Streaming using Amazon EMR, and use Spark MLlib with k-means to perform anomaly detection. Then store the results in an Apache Hadoop Distributed File System (HDFS) using Amazon EMR with a replication factor of three as the data lake.
- B. Ingest the data and store it in Amazon S3. Use AWS Batch along with the AWS Deep Learning AMIs to train a k-means model using TensorFlow on the data in Amazon S3.
- C. Ingest the data and store it in Amazon S3. Have an AWS Glue job that is triggered on demand transform the new data. Then use the built-in Random Cut Forest (RCF) model within Amazon SageMaker to detect anomalies in the data.
- **D. Ingest the data using Amazon Kinesis Data Firehose, and use Amazon Kinesis Data Analytics Random Cut Forest (RCF) for anomaly detection. Then use Kinesis Data Firehose to stream the results to Amazon S3.**

**Answer: D**

Explanation:

Amazon Kinesis Data Firehose is a fully managed service that can capture, transform, and load streaming data into AWS data stores, such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service, and Splunk. It can also invoke AWS Lambda functions to perform custom transformations on the data. Amazon Kinesis Data Analytics is a service that can analyze streaming data in real time using SQL or Apache Flink applications. It can also use machine learning algorithms, such as Random Cut Forest (RCF), to perform anomaly detection on streaming data. RCF is an unsupervised learning algorithm that assigns an anomaly score to each data point based on how different it is from the rest of the data. By using Kinesis Data Firehose and Kinesis Data Analytics, the cybersecurity company can ingest the data in real time, score the malicious events as anomalies, and stream the results to Amazon S3, which can serve as a data lake for later processing and analysis. This is the most efficient way to accomplish these tasks, as it does not require any additional infrastructure, coding, or training.

References:

Amazon Kinesis Data Firehose - Amazon Web Services

Amazon Kinesis Data Analytics - Amazon Web Services

Anomaly Detection with Amazon Kinesis Data Analytics - Amazon Web Services

[AWS Certified Machine Learning - Specialty Sample Questions]

### NEW QUESTION # 31

A company is building a predictive maintenance model for its warehouse equipment. The model must predict the probability of failure of all machines in the warehouse. The company has collected 10,000 event samples within 3 months. The event samples include 100 failure cases that are evenly distributed across 50 different machine types.

How should the company prepare the data for the model to improve the model's accuracy?

- A. Undersample the non-failure events by using the Synthetic Minority Oversampling Technique (SMOTE).
- B. Undersample the non-failure events. Stratify the non-failure events by machine type.
- C. Adjust the class weight to account for each machine type.
- **D. Oversample the failure cases by using the Synthetic Minority Oversampling Technique (SMOTE).**

**Answer: D**

Explanation:

In predictive maintenance, when a dataset is imbalanced (with far fewer failure cases than non-failure cases), oversampling the minority class helps the model learn from the minority class effectively. The Synthetic Minority Oversampling Technique (SMOTE) generates synthetic samples for the minority class by creating data points between existing minority class instances. This can enhance the model's ability to recognize failure patterns, particularly in imbalanced datasets.

SMOTE increases the effective presence of failure cases in the dataset, providing a balanced learning environment for the model. This is more effective than undersampling, which would risk losing important non-failure data.

### NEW QUESTION # 32

A company is launching a new product and needs to build a mechanism to monitor comments about the company and its new product on social media. The company needs to be able to evaluate the sentiment expressed in social media posts, and visualize trends and configure alarms based on various thresholds.

The company needs to implement this solution quickly, and wants to minimize the infrastructure and data science resources needed to evaluate the messages. The company already has a solution in place to collect posts and store them within an Amazon S3 bucket. What services should the data science team use to deliver this solution?

- A. Trigger an AWS Lambda function when social media posts are added to the S3 bucket. Call Amazon Comprehend for each post to capture the sentiment in the message and record the sentiment in a custom Amazon CloudWatch metric and in S3. Use CloudWatch alarms to notify analysts of trends.
- B. Train a model in Amazon SageMaker by using the semantic segmentation algorithm to model the semantic content in the corpus of social media posts. Expose an endpoint that can be called by AWS Lambda. Trigger a Lambda function when objects are added to the S3 bucket to invoke the endpoint and record the sentiment in an Amazon DynamoDB table. Schedule a second Lambda function to query recently added records and send an Amazon Simple Notification Service (Amazon SNS) notification to notify analysts of trends.
- C. Trigger an AWS Lambda function when social media posts are added to the S3 bucket. Call Amazon Comprehend for each post to capture the sentiment in the message and record the sentiment in an Amazon DynamoDB table. Schedule a second Lambda function to query recently added records and send an Amazon Simple Notification Service (Amazon SNS) notification to notify analysts of trends.
- D. Train a model in Amazon SageMaker by using the BlazingText algorithm to detect sentiment in the corpus of social media posts. Expose an endpoint that can be called by AWS Lambda. Trigger a Lambda function when posts are added to the S3 bucket to invoke the endpoint and record the sentiment in an Amazon DynamoDB table and in a custom Amazon CloudWatch metric. Use CloudWatch alarms to notify analysts of trends.

**Answer: A**

Explanation:

The solution that uses Amazon Comprehend and Amazon CloudWatch is the most suitable for the given scenario. Amazon Comprehend is a natural language processing (NLP) service that can analyze text and extract insights such as sentiment, entities, topics, and syntax. Amazon CloudWatch is a monitoring and observability service that can collect and track metrics, create dashboards, and set alarms based on various thresholds. By using these services, the data science team can quickly and easily implement a solution to monitor the sentiment of social media posts without requiring much infrastructure or data science resources. The solution also meets the requirements of storing the sentiment in both S3 and CloudWatch, and using CloudWatch alarms to notify analysts of trends.

Amazon Comprehend

Amazon CloudWatch

### NEW QUESTION # 33

A financial services company wants to adopt Amazon SageMaker as its default data science environment. The company's data scientists run machine learning (ML) models on confidential financial data. The company is worried about data egress and wants an ML engineer to secure the environment.

Which mechanisms can the ML engineer use to control data egress from SageMaker? (Choose three.)

- A. Disable root access on the SageMaker notebook instances.
- B. Restrict notebook presigned URLs to specific IPs used by the company.
- C. Connect to SageMaker by using a VPC interface endpoint powered by AWS PrivateLink.
- D. Protect data with encryption at rest and in transit. Use AWS Key Management Service (AWS KMS) to manage encryption keys.
- E. Enable network isolation for training jobs and models.
- F. Use SCPs to restrict access to SageMaker.

**Answer: B,E,F**

Explanation:

<https://aws.amazon.com/blogs/machine-learning/millennium-management-secure-machine-learning-using-amazon-sagemaker/>

### NEW QUESTION # 34

.....

- P.S. Free 2025 Amazon AWS-Certified-Machine-Learning-Specialty dumps are available on Google Drive shared by PassLeader:  
<https://drive.google.com/open?id=1jU7cmZvuIXiMXnx4BfYRLqKqb0u-6Yio>