

**Useful AWS-Certified-Machine-Learning-Specialty  
Latest Exam Question | Amazing Pass Rate For AWS-  
Certified-Machine-Learning-Specialty Exam | 100%  
Pass-Rate AWS-Certified-Machine-Learning-Specialty:  
AWS Certified Machine Learning - Specialty**

---

## ► **What Is Amazon SageMaker?**

- Get Started
- SageMaker Studio
- Notebook Instances
- Autopilot: Automated ML
- Label Data
- Prepare and Analyze Datasets
- Process Data
- Create, Store, and Share Features
- Training
- Inference
- Docker containers with SageMaker
- Workflows
- Augmented AI
- Marketplace
- Security
- Monitoring

## Availability Zones

- API and SDK Reference



DOWNLOAD the newest ActualtestPDF AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free:  
[https://drive.google.com/open?id=1WJv5dCPi9SjYVsaHsNildJUAq9CY\\_2A6](https://drive.google.com/open?id=1WJv5dCPi9SjYVsaHsNildJUAq9CY_2A6)

We have three different versions of AWS Certified Machine Learning - Specialty prep torrent for you to choose, including PDF

version, PC version and APP online version. Different versions have their own advantages and user population, and we would like to introduce features of these versions for you. There is no doubt that PDF of AWS-Certified-Machine-Learning-Specialty exam torrent is the most prevalent version among youngsters, mainly due to its convenience for a demo, through which you can have a general understanding and simulation about our AWS-Certified-Machine-Learning-Specialty Test Braindumps to decide whether you are willing to purchase or not, and also convenience for paper printing for you to do some note-taking. As for PC version of our AWS Certified Machine Learning - Specialty prep torrent, it is popular with computer users, and the software is more powerful. Finally when it comes to APP online version of AWS-Certified-Machine-Learning-Specialty test braindumps, as long as you open this study test engine, you are able to study whenever you like and wherever you are.

The AWS Certified Machine Learning - Specialty exam (MLS-C01) is a certification offered by Amazon Web Services (AWS) for individuals who want to validate their expertise in machine learning on the AWS cloud. AWS Certified Machine Learning - Specialty certification is designed to validate a candidate's understanding of the core concepts and best practices of machine learning implementation on AWS, including data preparation and cleaning, feature engineering, model development, and deployment.

Achieving the Amazon MLS-C01 certification is a significant accomplishment for machine learning professionals. It demonstrates a high level of knowledge and expertise in the field of machine learning and validates the skills required to design, build, and deploy machine learning models on AWS. AWS Certified Machine Learning - Specialty certification can help professionals advance their careers and increase their earning potential.

**>> AWS-Certified-Machine-Learning-Specialty Latest Exam Question <<**

## **Trustable AWS-Certified-Machine-Learning-Specialty - AWS Certified Machine Learning - Specialty Latest Exam Question**

The Amazon AWS-Certified-Machine-Learning-Specialty certification exam is one of the top-rated and valuable credentials in the Amazon world. This Amazon AWS-Certified-Machine-Learning-Specialty exam questions is designed to validate the candidate's skills and knowledge. With AWS Certified Machine Learning - Specialty exam dumps everyone can upgrade their expertise and knowledge level. By doing this the successful AWS-Certified-Machine-Learning-Specialty Exam candidates can gain several personal and professional benefits in their career and achieve their professional career objectives in a short time period.

### **Prerequisites**

The potential candidates should have 12-24 months of experience in architecting, developing, or running machine learning or deep learning workloads particularly on AWS Cloud. They must also possess the ability to effectively express the intuition behind basic machine learning algorithms. It is also advisable to have some experience with machine learning and deep learning frameworks, as well as be able to follow operational & deployment best practices.

### **Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q299-Q304):**

#### **NEW QUESTION # 299**

A Machine Learning Specialist has created a deep learning neural network model that performs well on the training data but performs poorly on the test data.

Which of the following methods should the Specialist consider using to correct this? (Select THREE.)

- A. Decrease regularization.
- B. Decrease dropout.
- C. Increase feature combinations.
- D. **Increase regularization.**
- E. **Decrease feature combinations.**
- F. **Increase dropout.**

**Answer: D,E,F**

Explanation:

The problem of poor performance on the test data is a sign of overfitting, which means the model has learned the training data too well and failed to generalize to new and unseen data. To correct this, the Machine Learning Specialist should consider using methods that reduce the complexity of the model and increase its ability to generalize. Some of these methods are:

Increase regularization: Regularization is a technique that adds a penalty term to the loss function of the model, which reduces the

magnitude of the model weights and prevents overfitting. There are different types of regularization, such as L1, L2, and elastic net, that apply different penalties to the weights1.

Increase dropout: Dropout is a technique that randomly drops out some units or connections in the neural network during training, which reduces the co-dependency of the units and prevents overfitting. Dropout can be applied to different layers of the network, and the dropout rate can be tuned to control the amount of dropout2.

Decrease feature combinations: Feature combinations are the interactions between different input features that can be used to create new features for the model. However, too many feature combinations can increase the complexity of the model and cause overfitting. Therefore, the Specialist should decrease the number of feature combinations and select only the most relevant and informative ones for the model3.

References:

1: Regularization for Deep Learning - Amazon SageMaker

2: Dropout - Amazon SageMaker

3: Feature Engineering - Amazon SageMaker

### NEW QUESTION # 300

A Data Scientist is evaluating different binary classification models. A false positive result is 5 times more expensive (from a business perspective) than a false negative result.

The models should be evaluated based on the following criteria:

- 1) Must have a recall rate of at least 80%
- 2) Must have a false positive rate of 10% or less
- 3) Must minimize business costs

After creating each binary classification model, the Data Scientist generates the corresponding confusion matrix.

Which confusion matrix represents the model that satisfies the requirements?

- A. TN = 99, FP = 1  
FN = 21, TP = 79
- B. TN = 91, FP = 9  
FN = 22, TP = 78
- C. TN = 96, FP = 4  
FN = 10, TP = 90
- D. TN = 98, FP = 2  
FN = 18, TP = 82

#### Answer: D

Explanation:

The following calculations are required:

TP = True Positive

FP = False Positive

FN = False Negative

TN = True Negative

FN = False Negative

Recall =  $TP / (TP + FN)$

False Positive Rate (FPR) =  $FP / (FP + TN)$

Cost =  $5 * FP + FN$

□ Options C and D have a recall greater than 80% and an FPR less than 10%, but D is the most cost effective.

### NEW QUESTION # 301

A data scientist stores financial datasets in Amazon S3. The data scientist uses Amazon Athena to query the datasets by using SQL. The data scientist uses Amazon SageMaker to deploy a machine learning (ML) model. The data scientist wants to obtain inferences from the model at the SageMaker endpoint. However, when the data scientist attempts to invoke the SageMaker endpoint, the data scientist receives SQL statement failures. The data scientist's 1AM user is currently unable to invoke the SageMaker endpoint. Which combination of actions will give the data scientist's 1AM user the ability to invoke the SageMaker endpoint? (Select THREE.)

- A. Include a policy statement for the data scientist's 1AM user that allows the 1AM user to perform the `sagemaker:InvokeEndpoint` action,
- B. Include a policy statement for the data scientist's 1AM user that allows the 1AM user to perform the `sagemaker:GetRecord` action.

- C. Include an inline policy for the data scientist's IAM user that allows SageMaker to read S3 objects
- D. Perform a user remapping in SageMaker to map the IAM user to another IAM user that is on the hosted endpoint.
- E. Attach the AmazonAthenaFullAccess AWS managed policy to the user identity.
- F. Include the SQL statement "USING EXTERNAL FUNCTION ml\_function\_name" in the Athena SQL query.

**Answer: A,C,F**

Explanation:

The correct combination of actions to enable the data scientist's IAM user to invoke the SageMaker endpoint is B, C, and E, because they ensure that the IAM user has the necessary permissions, access, and syntax to query the ML model from Athena.

These actions have the following benefits:

- \* B: Including a policy statement for the IAM user that allows the sagemaker:InvokeEndpoint action grants the IAM user the permission to call the SageMaker Runtime InvokeEndpoint API, which is used to get inferences from the model hosted at the endpoint1.
- \* C: Including an inline policy for the IAM user that allows SageMaker to read S3 objects enables the IAM user to access the data stored in S3, which is the source of the Athena queries2.
- \* E: Including the SQL statement "USING EXTERNAL FUNCTION ml\_function\_name" in the Athena SQL query allows the IAM user to invoke the ML model as an external function from Athena, which is a feature that enables querying ML models from SQL statements3.

The other options are not correct or necessary, because they have the following drawbacks:

- \* A: Attaching the AmazonAthenaFullAccess AWS managed policy to the user identity is not sufficient, because it does not grant the IAM user the permission to invoke the SageMaker endpoint, which is required to query the ML model4.
- \* D: Including a policy statement for the IAM user that allows the IAM user to perform the sagemaker: GetRecord action is not relevant, because this action is used to retrieve a single record from a feature group, which is not the case in this scenario5.
- \* F: Performing a user remapping in SageMaker to map the IAM user to another IAM user that is on the hosted endpoint is not applicable, because this feature is only available for multi-model endpoints, which are not used in this scenario.

1: InvokeEndpoint - Amazon SageMaker

2: Querying Data in Amazon S3 from Amazon Athena - Amazon Athena

3: Querying machine learning models from Amazon Athena using Amazon SageMaker | AWS Machine Learning Blog

4: AmazonAthenaFullAccess - AWS Identity and Access Management

5: GetRecord - Amazon SageMaker Feature Store Runtime

[Invoke a Multi-Model Endpoint - Amazon SageMaker]

## NEW QUESTION # 302

An ecommerce company wants to use machine learning (ML) to monitor fraudulent transactions on its website. The company is using Amazon SageMaker to research, train, deploy, and monitor the ML models.

The historical transactions data is in a .csv file that is stored in Amazon S3. The data contains features such as the user's IP address, navigation time, average time on each page, and the number of clicks for ....session.

There is no label in the data to indicate if a transaction is anomalous.

Which models should the company use in combination to detect anomalous transactions? (Select TWO.)

- A. Random Cut Forest (RCF)
- B. K-nearest neighbors (k-NN)
- C. XGBoost
- D. IP Insights
- E. Linear learner with a logistic function

**Answer: A,C**

Explanation:

Explanation

To detect anomalous transactions, the company can use a combination of Random Cut Forest (RCF) and XGBoost models. RCF is an unsupervised algorithm that can detect outliers in the data by measuring the depth of each data point in a collection of random decision trees. XGBoost is a supervised algorithm that can learn from the labeled data points generated by RCF and classify them as normal or anomalous. RCF can also provide anomaly scores that can be used as features for XGBoost to improve the accuracy of the classification. References:

1: Amazon SageMaker Random Cut Forest

2: Amazon SageMaker XGBoost Algorithm

3: Anomaly Detection with Amazon SageMaker Random Cut Forest and Amazon SageMaker XGBoost

### NEW QUESTION # 303

A Machine Learning Specialist has built a model using Amazon SageMaker built-in algorithms and is not getting expected accurate results. The Specialist wants to use hyperparameter optimization to increase the model's accuracy. Which method is the MOST repeatable and requires the LEAST amount of effort to achieve this?

- A. Create an AWS Step Functions workflow that monitors the accuracy in Amazon CloudWatch Logs and relaunches the training job with a defined list of hyperparameters
- B. Launch multiple training jobs in parallel with different hyperparameters
- C. Create a random walk in the parameter space to iterate through a range of values that should be used for each individual hyperparameter
- D. Create a hyperparameter tuning job and set the accuracy as an objective metric.

#### Answer: D

Explanation:

A hyperparameter tuning job is a feature of Amazon SageMaker that allows automatically finding the best combination of hyperparameters for a machine learning model. Hyperparameters are high-level parameters that influence the learning process and the performance of the model, such as the learning rate, the number of layers, the regularization factor, etc. A hyperparameter tuning job works by launching multiple training jobs with different hyperparameters, evaluating the results using an objective metric, and choosing the next set of hyperparameters to try based on a search strategy. The objective metric is a measure of the quality of the model, such as accuracy, precision, recall, etc. The search strategy is a method of exploring the hyperparameter space, such as random search, grid search, or Bayesian optimization.

Among the four options, option C is the most repeatable and requires the least amount of effort to use hyperparameter optimization to increase the model's accuracy. This option involves the following steps:

Create a hyperparameter tuning job: Amazon SageMaker provides an easy-to-use interface for creating a hyperparameter tuning job, either through the AWS Management Console, the AWS CLI, or the AWS SDKs.

To create a hyperparameter tuning job, the Machine Learning Specialist needs to specify the following information:

The name and type of the algorithm to use, either a built-in algorithm or a custom algorithm.

The ranges and types of the hyperparameters to tune, such as categorical, continuous, or integer.

The name and type of the objective metric to optimize, such as accuracy, and whether to maximize or minimize it.

The resource limits for the tuning job, such as the maximum number of training jobs and the maximum parallel training jobs.

The input data channels and the output data location for the training jobs.

The configuration of the training instances, such as the instance type, the instance count, the volume size, etc.

Set the accuracy as an objective metric: To use accuracy as an objective metric, the Machine Learning Specialist needs to ensure that the training algorithm writes the accuracy value to a file called metric\_definitions in JSON format and prints it to stdout or stderr.

For example, the file can contain the following content:

This means that the training algorithm prints a line like this:

Amazon SageMaker reads the accuracy value from the line and uses it to evaluate and compare the training jobs.

The other options are not as repeatable and require more effort than option C for the following reasons:

Option A: This option requires manually launching multiple training jobs in parallel with different hyperparameters, which can be tedious and error-prone. It also requires manually monitoring and comparing the results of the training jobs, which can be time-consuming and subjective.

Option B: This option requires writing code to create an AWS Step Functions workflow that monitors the accuracy in Amazon CloudWatch Logs and relaunches the training job with a defined list of hyperparameters, which can be complex and challenging. It also requires maintaining and updating the list of hyperparameters, which can be inefficient and suboptimal.

Option D: This option requires writing code to create a random walk in the parameter space to iterate through a range of values that should be used for each individual hyperparameter, which can be unreliable and unpredictable. It also requires defining and implementing a stopping criterion, which can be arbitrary and inconsistent.

Automatic Model Tuning - Amazon SageMaker

Define Metrics to Monitor Model Performance

### NEW QUESTION # 304

.....

**AWS-Certified-Machine-Learning-Specialty Certification Cost:** <https://www.actualtestpdf.com/Amazon/AWS-Certified-Machine-Learning-Specialty-practice-exam-dumps.html>

- Vce AWS-Certified-Machine-Learning-Specialty Download  Latest AWS-Certified-Machine-Learning-Specialty Test

Preparation  Reliable AWS-Certified-Machine-Learning-Specialty Test Experience  Download “ AWS-Certified-Machine-Learning-Specialty ” for free by simply entering ✓ [www.prepawayete.com](http://www.prepawayete.com)  ✓  website  AWS-Certified-Machine-Learning-Specialty Test Simulator

DOWNLOAD the newest ActualtestPDF AWS-Certified-Machine-Learning-Specialty PDF dumps from Cloud Storage for free: [https://drive.google.com/open?id=1WJv5dCPi9SjYVsahsNildJUaQ9CY\\_2A6](https://drive.google.com/open?id=1WJv5dCPi9SjYVsahsNildJUaQ9CY_2A6)