

# 100% Pass Quiz 2026 Newest MLS-C01: Clear AWS Certified Machine Learning - Specialty Exam



What's more, part of that PassTestking MLS-C01 dumps now are free: <https://drive.google.com/open?id=1w9wQJzMNBUDaHdKV8mrx-fQLSjZw1NyB>

PassTestking also offers a demo version of the Amazon MLS-C01 exam dumps for free. This way you can easily evaluate the validity of the MLS-C01 prep material before buying it. Downloading a free demo will remove your doubts about purchasing the Amazon MLS-C01 Questions.

Amazon MLS-C01 Exam is a highly respected certification that validates the skills and knowledge of individuals who work with machine learning technologies on the AWS platform. By passing MLS-C01 Exam, candidates can demonstrate their expertise in machine learning, as well as their ability to design, deploy, and maintain machine learning solutions on AWS.

>> Clear MLS-C01 Exam <<

## Amazon MLS-C01 Exam Questions Pdf | Reliable MLS-C01 Exam Testking

In the major environment, people are facing more job pressure. So they want to get Amazon certification rise above the common herd. How to choose valid and efficient MLS-C01 guide torrent should be the key topic most candidates may concern. So now, it is right, you come to us. Our company is famous for its high-quality MLS-C01 Exam Questions in this field especially for Amazon certification exams. It has been accepted by thousands of candidates who practice our MLS-C01 study materials for their exam.

The AWS Certified Machine Learning - Specialty certification is ideal for individuals who want to advance their careers in the field of machine learning and artificial intelligence. It can help them demonstrate their expertise to potential employers and clients, and increase their earning potential. Moreover, it provides them with access to the AWS Certified community, which includes resources and networking opportunities to help them stay up-to-date with the latest trends and technologies in the industry.

## Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q47-Q52):

### NEW QUESTION # 47

An e-commerce company needs a customized training model to classify images of its shirts and pants products. The company needs a proof of concept in 2 to 3 days with good accuracy. Which compute choice should the Machine Learning Specialist select to train and achieve good accuracy on the model quickly?

- A. m5 4xlarge (general purpose)
- B. r5.2xlarge (memory optimized)
- C. p3.2xlarge (GPU accelerated computing)
- D. p3 8xlarge (GPU accelerated computing)

**Answer: C**

#### Explanation:

Image classification is a machine learning task that involves assigning labels to images based on their content.

Image classification can be performed using various algorithms, such as convolutional neural networks (CNNs), which are a type of deep learning model that can learn to extract high-level features from images. To train a customized image classification model, the e-commerce company needs a compute choice that can support the high computational demands of deep learning and provide good accuracy on the model quickly. A GPU accelerated computing instance, such as p3.2xlarge, is a suitable choice for this task, as it can leverage the parallel processing power of GPUs to speed up the training process and reduce the training time. A p3.2xlarge instance has one NVIDIA Tesla V100 GPU, which can provide up to 125 teraflops of mixed-precision performance and 16 GB of GPU memory. A p3.2xlarge instance can also use various deep learning frameworks, such as TensorFlow, PyTorch, MXNet, etc., to build and train the image classification model. A p3.2xlarge instance is also more cost-effective than a p3.8xlarge instance, which has four NVIDIA Tesla V100 GPUs, as the latter may not be necessary for a proof of concept with a small dataset. Therefore, the Machine Learning Specialist should select p3.2xlarge as the compute choice to train and achieve good accuracy on the model quickly.

#### References:

- \* Amazon EC2 P3 Instances - Amazon Web Services
- \* Image Classification - Amazon SageMaker
- \* Convolutional Neural Networks - Amazon SageMaker
- \* Deep Learning AMIs - Amazon Web Services

### NEW QUESTION # 48

A Machine Learning Specialist is building a model that will perform time series forecasting using Amazon SageMaker. The Specialist has finished training the model and is now planning to perform load testing on the endpoint so they can configure Auto Scaling for the model variant. Which approach will allow the Specialist to review the latency, memory utilization, and CPU utilization during the load test?"?

- A. Send Amazon CloudWatch Logs that were generated by Amazon SageMaker to Amazon ES and use Kibana to query and visualize the log data.
- B. Generate an Amazon CloudWatch dashboard to create a single view for the latency, memory utilization, and CPU utilization metrics that are outputted by Amazon SageMaker
- C. Build custom Amazon CloudWatch Logs and then leverage Amazon ES and Kibana to query and visualize the data as it is generated by Amazon SageMaker
- D. Review SageMaker logs that have been written to Amazon S3 by leveraging Amazon Athena and Amazon QuickSight to visualize logs as they are being produced

#### Answer: C

### NEW QUESTION # 49

A Data Scientist wants to gain real-time insights into a data stream of GZIP files.

Which solution would allow the use of SQL to query the stream with the LEAST latency?

- A. Amazon Kinesis Data Firehose to transform the data and put it into an Amazon S3 bucket.
- B. An Amazon Kinesis Client Library to transform the data and save it to an Amazon ES cluster.
- C. Amazon Kinesis Data Analytics with an AWS Lambda function to transform the data.
- D. AWS Glue with a custom ETL script to transform the data.

#### Answer: C

#### Explanation:

<https://aws.amazon.com/big-data/real-time-analytics-featured-partners/>

### NEW QUESTION # 50

When submitting Amazon SageMaker training jobs using one of the built-in algorithms, which common parameters MUST be specified? (Choose three.)

- A. Hyperparameters in a JSON array as documented for the algorithm used.
- B. The validation channel identifying the location of validation data on an Amazon S3 bucket.
- C. The output path specifying where on an Amazon S3 bucket the trained model will persist.

- D. The training channel identifying the location of training data on an Amazon S3 bucket.
- E. The Amazon EC2 instance class specifying whether training will be run using CPU or GPU.
- F. The IAM role that Amazon SageMaker can assume to perform tasks on behalf of the users.

**Answer: C,D,E**

## NEW QUESTION # 51

An insurance company is developing a new device for vehicles that uses a camera to observe drivers' behavior and alert them when they appear distracted. The company created approximately 10,000 training images in a controlled environment that a Machine Learning Specialist will use to train and evaluate machine learning models.

During the model evaluation, the Specialist notices that the training error rate diminishes faster as the number of epochs increases and the model is not accurately inferring on the unseen test images.

Which of the following should be used to resolve this issue? (Choose two.)

- A. Make the neural network architecture complex.
- B. Add vanishing gradient to the model.
- C. Perform data augmentation on the training data.
- D. Use gradient checking in the model.
- E. Add L2 regularization to the model.

**Answer: C,D**

## NEW QUESTION # 52

• • • • •

**MLS-C01 Exam Questions Pdf:** <https://www.passtestking.com/Amazon/MLS-C01-practice-exam-dumps.html>

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, academy.eleven11prod.com, Disposable vapes

2026 Latest PassTestking MLS-C01 PDF Dumps and MLS-C01 Exam Engine Free Share: <https://drive.google.com/open?id=1w9wQJzMNBDUaHdKV8mrxfQLSjZw1NyB>