

# Valid AWS-Certified-Machine-Learning-Specialty Exam Simulator & Valid AWS-Certified-Machine-Learning-Specialty Exam Questions



2026 Latest Pass4guide AWS-Certified-Machine-Learning-Specialty PDF Dumps and AWS-Certified-Machine-Learning-Specialty Exam Engine Free Share: <https://drive.google.com/open?id=1F24Llstw0JG94AqBZODQjuAHZkLNdWcC>

On the one thing, our company has employed a lot of leading experts in the field to compile the AWS-Certified-Machine-Learning-Specialty exam torrents, so you can definitely feel rest assured about the high quality of our AWS-Certified-Machine-Learning-Specialty question torrents. On the other thing, the pass rate among our customers who prepared the exam under the guidance of our AWS-Certified-Machine-Learning-Specialty Study Materials has reached as high as 98% to 100%. What's more, you will have more opportunities to get promotion as well as a pay raise in the near future after using our AWS-Certified-Machine-Learning-Specialty question torrents since you are sure to get the AWS-Certified-Machine-Learning-Specialty certification.

## Difficulty in preparing for AWS Certified Machine Learning Specialty Exam

In addition to our comprehensive study guide, we also offer exam dumps of certified AWS Certified Machine Learning Specialty, if you want a quick and exam-oriented preparation. Any information in this **AWS Certified Machine Learning Specialty exam dumps** is valuable.

Questions and answers from the AWS Certified Machine Learning Specialty include important topics from the AWS Certified Machine Learning Specialty Certification Program and provides easy-to-learn information for easy access.

The AWS Certified Machine Learning - Specialty (MLS-C01) examination is intended for individuals who perform a development or data science role. This exam validates an examinee's ability to build, train, tune, and deploy machine learning (ML) models using the AWS Cloud.

Candidate must have 1-2 years of hands-on experience developing, architecting, or running ML/deeplearning workloads on the AWS Cloud, along with:

- Experience performing basic hyperparameter optimization
- The ability to express the intuition behind basic ML algorithms
- Experience with ML and deep learning frameworks
- The ability to follow deployment and operational best practices
- The ability to follow model-training best practices

Candidates for the AWS Certified Machine Learning Specialty should have a thorough knowledge and understanding of all the questions and answers of the AWS Certified Machine Learning Specialty in our practice exam and exam dumps.

>> **Valid AWS-Certified-Machine-Learning-Specialty Exam Simulator** <<

**100% Pass Quiz 2026 Pass-Sure AWS-Certified-Machine-Learning-**

## Specialty: Valid AWS Certified Machine Learning - Specialty Exam Simulator

Our AWS-Certified-Machine-Learning-Specialty exam questions are designed from the customer's perspective, and experts that we employed will update our AWS-Certified-Machine-Learning-Specialty learning materials according to changing trends to ensure the high quality of the AWS-Certified-Machine-Learning-Specialty practice materials. What are you still waiting for? Choosing our AWS-Certified-Machine-Learning-Specialty guide questions and work for getting the certificate, you will make your life more colorful and successful.

### Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q304-Q309):

#### NEW QUESTION # 304

A Machine Learning Specialist is building a convolutional neural network (CNN) that will classify 10 types of animals. The Specialist has built a series of layers in a neural network that will take an input image of an animal, pass it through a series of convolutional and pooling layers, and then finally pass it through a dense and fully connected layer with 10 nodes. The Specialist would like to get an output from the neural network that is a probability distribution of how likely it is that the input image belongs to each of the 10 classes. Which function will produce the desired output?

- A. Rectified linear units (ReLU)
- **B. Softmax**
- C. Smooth L1 loss
- D. Dropout

**Answer: B**

Explanation:

Explanation

The softmax function is a function that can transform a vector of arbitrary real values into a vector of real values in the range (0,1) that sum to 1. This means that the softmax function can produce a valid probability distribution over multiple classes. The softmax function is often used as the activation function of the output layer in a neural network, especially for multi-class classification problems. The softmax function can assign higher probabilities to the classes with higher scores, which allows the network to make predictions based on the most likely class. In this case, the Machine Learning Specialist wants to get an output from the neural network that is a probability distribution of how likely it is that the input image belongs to each of the 10 classes of animals. Therefore, the softmax function is the most suitable function to produce the desired output.

References:

Softmax Activation Function for Deep Learning: A Complete Guide

What is Softmax in Machine Learning? - reason.town

machine learning - Why is the softmax function often used as activation ...

Multi-Class Neural Networks: Softmax | Machine Learning | Google for ...

#### NEW QUESTION # 305

A company wants to forecast the daily price of newly launched products based on 3 years of data for older product prices, sales, and rebates. The time-series data has irregular timestamps and is missing some values.

Data scientist must build a dataset to replace the missing values. The data scientist needs a solution that resamples the data daily and exports the data for further modeling.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Use Amazon EMR Serverless with PySpark.
- B. Use AWS Glue DataBrew.
- **C. Use Amazon SageMaker Studio Data Wrangler.**
- D. Use Amazon SageMaker Studio Notebook with Pandas.

**Answer: C**

Explanation:

Amazon SageMaker Studio Data Wrangler is a visual data preparation tool that enables users to clean and normalize data without writing any code. Using Data Wrangler, the data scientist can easily import the time-series data from various sources, such as Amazon S3, Amazon Athena, or Amazon Redshift. Data Wrangler can automatically generate data insights and quality reports,

which can help identify and fix missing values, outliers, and anomalies in the data. Data Wrangler also provides over 250 built-in transformations, such as resampling, interpolation, aggregation, and filtering, which can be applied to the data with a point-and-click interface. Data Wrangler can also export the prepared data to different destinations, such as Amazon S3, Amazon SageMaker Feature Store, or Amazon SageMaker Pipelines, for further modeling and analysis. Data Wrangler is integrated with Amazon SageMaker Studio, a web-based IDE for machine learning, which makes it easy to access and use the tool. Data Wrangler is a serverless and fully managed service, which means the data scientist does not need to provision, configure, or manage any infrastructure or clusters.

Option A is incorrect because Amazon EMR Serverless is a serverless option for running big data analytics applications using open-source frameworks, such as Apache Spark. However, using Amazon EMR Serverless would require the data scientist to write PySpark code to perform the data preparation tasks, such as resampling, imputation, and aggregation. This would require more implementation effort than using Data Wrangler, which provides a visual and code-free interface for data preparation.

Option B is incorrect because AWS Glue DataBrew is another visual data preparation tool that can be used to clean and normalize data without writing code. However, DataBrew does not support time-series data as a data type, and does not provide built-in transformations for resampling, interpolation, or aggregation of time-series data. Therefore, using DataBrew would not meet the requirements of the use case.

Option D is incorrect because using Amazon SageMaker Studio Notebook with Pandas would also require the data scientist to write Python code to perform the data preparation tasks. Pandas is a popular Python library for data analysis and manipulation, which supports time-series data and provides various methods for resampling, interpolation, and aggregation. However, using Pandas would require more implementation effort than using Data Wrangler, which provides a visual and code-free interface for data preparation.

References:

- 1: Amazon SageMaker Data Wrangler documentation
- 2: Amazon EMR Serverless documentation
- 3: AWS Glue DataBrew documentation
- 4: Pandas documentation

### NEW QUESTION # 306

The chief editor for a product catalog wants the research and development team to build a machine learning system that can be used to detect whether or not individuals in a collection of images are wearing the company's retail brand. The team has a set of training data.

Which machine learning algorithm should the researchers use that BEST meets their requirements?

- A. Latent Dirichlet Allocation (LDA)
- B. Recurrent neural network (RNN)
- C. Convolutional neural network (CNN)
- D. K-means

**Answer: C**

Explanation:

The problem of detecting whether or not individuals in a collection of images are wearing the company's retail brand is an example of image recognition, which is a type of machine learning task that identifies and classifies objects in an image. Convolutional neural networks (CNNs) are a type of machine learning algorithm that are well-suited for image recognition, as they can learn to extract features from images and handle variations in size, shape, color, and orientation of the objects. CNNs consist of multiple layers that perform convolution, pooling, and activation operations on the input images, resulting in a high-level representation that can be used for classification or detection. Therefore, option D is the best choice for the machine learning algorithm that meets the requirements of the chief editor.

Option A is incorrect because latent Dirichlet allocation (LDA) is a type of machine learning algorithm that is used for topic modeling, which is a task that discovers the hidden themes or topics in a collection of text documents. LDA is not suitable for image recognition, as it does not preserve the spatial information of the pixels. Option B is incorrect because recurrent neural networks (RNNs) are a type of machine learning algorithm that are used for sequential data, such as text, speech, or time series. RNNs can learn from the temporal dependencies and patterns in the input data, and generate outputs that depend on the previous states. RNNs are not suitable for image recognition, as they do not capture the spatial dependencies and patterns in the input images. Option C is incorrect because k-means is a type of machine learning algorithm that is used for clustering, which is a task that groups similar data points together based on their features. K-means is not suitable for image recognition, as it does not perform classification or detection of the objects in the images.

References:

- \* Image Recognition Software - ML Image & Video Analysis - Amazon ...
- \* Image classification and object detection using Amazon Rekognition ...
- \* AWS Amazon Rekognition - Deep Learning Face and Image Recognition ...

\* GitHub - aws-labs/aws-ai-solution-kit: Machine Learning APIs for common ...

\* Meet iNaturalist, an AWS-powered nature app that helps you identify ...

### NEW QUESTION # 307

An employee found a video clip with audio on a company's social media feed. The language used in the video is Spanish. English is the employee's first language, and they do not understand Spanish. The employee wants to do a sentiment analysis.

What combination of services is the MOST efficient to accomplish the task?

- A. Amazon Transcribe, Amazon Translate, and Amazon SageMaker Neural Topic Model (NTM)
- B. Amazon Transcribe, Amazon Comprehend, and Amazon SageMaker seq2seq
- **C. Amazon Transcribe, Amazon Translate, and Amazon Comprehend**
- D. Amazon Transcribe, Amazon Translate, and Amazon SageMaker BlazingText

**Answer: C**

Explanation:

Amazon Transcribe, Amazon Translate, and Amazon Comprehend are the most efficient combination of services to accomplish the task of sentiment analysis on a video clip with audio in Spanish. Amazon Transcribe is a service that can convert speech to text using deep learning. Amazon Transcribe can transcribe audio from various sources, such as video files, audio files, or streaming audio.

Amazon Transcribe can also recognize multiple speakers, different languages, accents, dialects, and custom vocabularies. In this case, Amazon Transcribe can transcribe the audio from the video clip in Spanish to text in Spanish<sup>1</sup>. Amazon Translate is a service that can translate text from one language to another using neural machine translation.

Amazon Translate can translate text from various sources, such as documents, web pages, chat messages, etc.

Amazon Translate can also support multiple languages, domains, and styles. In this case, Amazon Translate can translate the text from Spanish to English<sup>2</sup>. Amazon Comprehend is a service that can analyze and derive insights from text using natural language processing. Amazon Comprehend can perform various tasks, such as sentiment analysis, entity recognition, key phrase extraction, topic modeling, etc. Amazon Comprehend can also support multiple languages and domains. In this case, Amazon Comprehend can perform sentiment analysis on the text in English and determine whether the feedback is positive, negative, neutral, or mixed<sup>3</sup>. The other options are not valid or efficient for accomplishing the task of sentiment analysis on a video clip with audio in Spanish. Amazon Comprehend, Amazon SageMaker seq2seq, and Amazon SageMaker Neural Topic Model (NTM) are not a good combination, as they do not include a service that can transcribe speech to text, which is a necessary step for processing the audio from the video clip. Amazon Comprehend, Amazon Translate, and Amazon SageMaker BlazingText are not a good combination, as they do not include a service that can perform sentiment analysis, which is the main goal of the task. Amazon SageMaker BlazingText is a service that can train and deploy text classification and word embedding models using deep learning. Amazon SageMaker BlazingText can perform tasks such as text classification, named entity recognition, part-of-speech tagging, etc., but not sentiment analysis<sup>4</sup>.

### NEW QUESTION # 308

A company is building a new version of a recommendation engine. Machine learning (ML) specialists need to keep adding new data from users to improve personalized recommendations. The ML specialists gather data from the users' interactions on the platform and from sources such as external websites and social media.

The pipeline cleans, transforms, enriches, and compresses terabytes of data daily, and this data is stored in Amazon S3. A set of Python scripts was coded to do the job and is stored in a large Amazon EC2 instance.

The whole process takes more than 20 hours to finish, with each script taking at least an hour. The company wants to move the scripts out of Amazon EC2 into a more managed solution that will eliminate the need to maintain servers.

Which approach will address all of these requirements with the LEAST development effort?

- A. Create a set of individual AWS Lambda functions to execute each of the scripts. Build a step function by using the AWS Step Functions Data Science SDK. Store the results in Amazon S3.
- B. Load the data into Amazon DynamoDB. Convert the scripts to an AWS Lambda function. Execute the pipeline by triggering Lambda executions. Store the results in Amazon S3.
- **C. Create an AWS Glue job. Convert the scripts to PySpark. Execute the pipeline. Store the results in Amazon S3.**
- D. Load the data into an Amazon Redshift cluster. Execute the pipeline by using SQL. Store the results in Amazon S3.

**Answer: C**

Explanation:

Explanation

The best approach to address all of the requirements with the least development effort is to create an AWS Glue job, convert the scripts to PySpark, execute the pipeline, and store the results in Amazon S3. This is because:

AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load data for analytics 1. AWS Glue can run Python and Scala scripts to process data from various sources, such as Amazon S3, Amazon DynamoDB, Amazon Redshift, and more 2. AWS Glue also provides a serverless Apache Spark environment to run ETL jobs, eliminating the need to provision and manage servers 3.

PySpark is the Python API for Apache Spark, a unified analytics engine for large-scale data processing 4. PySpark can perform various data transformations and manipulations on structured and unstructured data, such as cleaning, enriching, and compressing 5. PySpark can also leverage the distributed computing power of Spark to handle terabytes of data efficiently and scalably 6.

By creating an AWS Glue job and converting the scripts to PySpark, the company can move the scripts out of Amazon EC2 into a more managed solution that will eliminate the need to maintain servers. The company can also reduce the development effort by using the AWS Glue console, AWS SDK, or AWS CLI to create and run the job 7. Moreover, the company can use the AWS Glue Data Catalog to store and manage the metadata of the data sources and targets 8.

The other options are not as suitable as option C for the following reasons:

Option A is not optimal because loading the data into an Amazon Redshift cluster and executing the pipeline by using SQL will incur additional costs and complexity for the company. Amazon Redshift is a fully managed data warehouse service that enables fast and scalable analysis of structured data .

However, it is not designed for ETL purposes, such as cleaning, transforming, enriching, and compressing data. Moreover, using SQL to perform these tasks may not be as expressive and flexible as using Python scripts. Furthermore, the company will have to provision and configure the Amazon Redshift cluster, and load and unload the data from Amazon S3, which will increase the development effort and time.

Option B is not feasible because loading the data into Amazon DynamoDB and converting the scripts to an AWS Lambda function will not work for the company's use case. Amazon DynamoDB is a fully managed key-value and document database service that provides fast and consistent performance at any scale . However, it is not suitable for storing and processing terabytes of data daily, as it has limits on the size and throughput of each table and item . Moreover, using AWS Lambda to execute the pipeline will not be efficient or cost-effective, as Lambda has limits on the memory, CPU, and execution time of each function . Therefore, using Amazon DynamoDB and AWS Lambda will not meet the company's requirements for processing large amounts of data quickly and reliably.

Option D is not relevant because creating a set of individual AWS Lambda functions to execute each of the scripts and building a step function by using the AWS Step Functions Data Science SDK will not address the main issue of moving the scripts out of Amazon EC2. AWS Step Functions is a fully managed service that lets you coordinate multiple AWS services into serverless workflows . The AWS Step Functions Data Science SDK is an open source library that allows data scientists to easily create workflows that process and publish machine learning models using Amazon SageMaker and AWS Step Functions . However, these services and tools are not designed for ETL purposes, such as cleaning, transforming, enriching, and compressing data. Moreover, as mentioned in option B, using AWS Lambda to execute the scripts will not be efficient or cost-effective for the company's use case.

References:

What Is AWS Glue?

AWS Glue Components

AWS Glue Serverless Spark ETL

PySpark - Overview

PySpark - RDD

PySpark - SparkContext

Adding Jobs in AWS Glue

Populating the AWS Glue Data Catalog

[What Is Amazon Redshift?]

[What Is Amazon DynamoDB?]

[Service, Account, and Table Quotas in DynamoDB]

[AWS Lambda quotas]

[What Is AWS Step Functions?]

[AWS Step Functions Data Science SDK for Python]

## NEW QUESTION # 309

.....

Our Amazon AWS-Certified-Machine-Learning-Specialty practice materials from our company are invulnerable. And we are consigned as the most responsible company in this area. So many competitors concede our superior position in the market. Besides, we offer some promotional benefits for you. The more times you choose our Amazon AWS-Certified-Machine-Learning-Specialty Training Materials, the more benefits you can get, such as free demos of our AWS-Certified-Machine-Learning-Specialty exam dumps, three-version options, rights of updates and so on. So customer orientation is the beliefs we honor.

**Valid AWS-Certified-Machine-Learning-Specialty Exam Questions:** <https://www.pass4guide.com/AWS-Certified-Machine->

