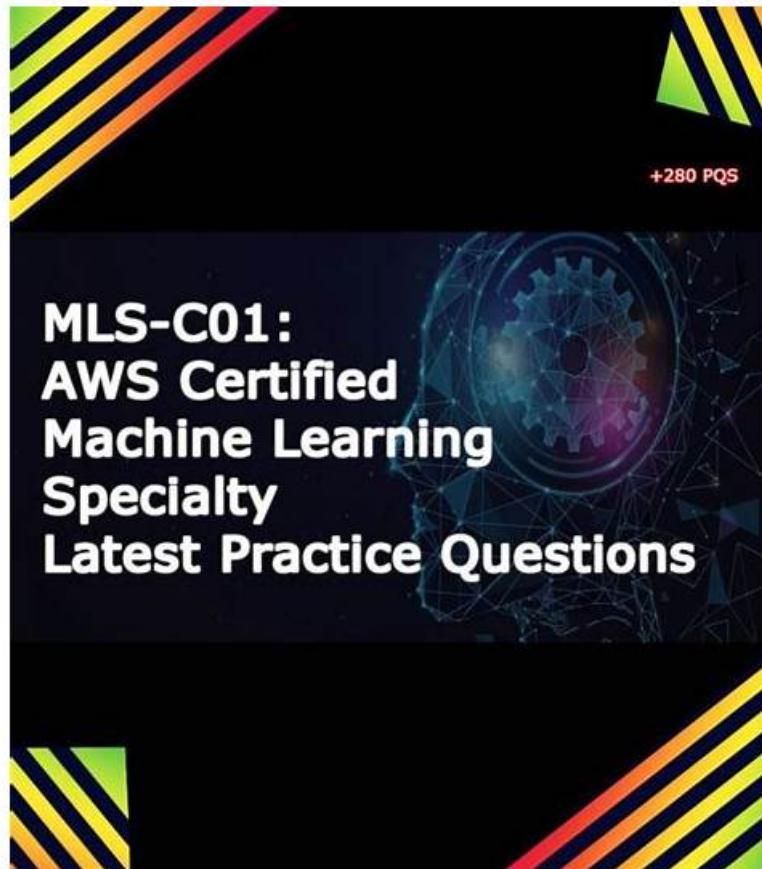


# MLS-C01 Simulation Questions & Training MLS-C01 Pdf



BTW, DOWNLOAD part of TestValid MLS-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1O-7C8B685JmpJEGBtldy6tKKt6HCrB2M>

If you still worry too much about purchasing professional MLS-C01 test guide on the internet, I can tell that it is quite normal. Useful certification MLS-C01 guide materials will help your preparing half work with double results. If you consider about our MLS-C01 exam questoins quality, you can free downlaod the demo of our MLS-C01 Exam Questions. We have thought of your needs and doubts considerably on the MLS-C01 study guide. Our certification MLS-C01 guide materials are collected and compiled by experience experts who have worked in this line more than 10 years.

Amazon MLS-C01 Certification Exam is a rigorous and comprehensive test of an individual's knowledge and skills in ML on AWS. Passing MLS-C01 exam requires a deep understanding of ML concepts, as well as practical experience using AWS services and tools. Candidates who successfully pass the exam will earn the Amazon Certified Machine Learning - Specialty certification, which is recognized globally as a validation of their skills and expertise.

The AWS Certified Machine Learning - Specialty certification exam is intended for individuals with a minimum of one year of experience in designing and implementing machine learning solutions on AWS. AWS Certified Machine Learning - Specialty certification is ideal for data scientists, data engineers, software developers, and IT professionals who want to expand their skill set and demonstrate their expertise in the field of machine learning.

>> [MLS-C01 Simulation Questions](#) <<

## MLS-C01 Simulation Questions Free PDF | Valid Training MLS-C01 Pdf: AWS Certified Machine Learning - Specialty

Contrary to most of the MLS-C01 exam preparatory material available online, TestValid's dumps can be obtained on an affordable price yet their quality and benefits beat all similar products of our competitors. They will prove the best alternative of your time and

money. What's more, our customers' care is available 24/7 for all visitors on our pages. You can put all your queries and get a quick and efficient response as well as advice of our experts on MLS-C01 Certification tests you want to take. Our professional online staff will attend you on priority.

## Amazon AWS Certified Machine Learning - Specialty Sample Questions (Q65-Q70):

### NEW QUESTION # 65

A Marketing Manager at a pet insurance company plans to launch a targeted marketing campaign on social media to acquire new customers. Currently, the company has the following data in Amazon Aurora

- \* Profiles for all past and existing customers
- \* Profiles for all past and existing insured pets
- \* Policy-level information
- \* Premiums received
- \* Claims paid

What steps should be taken to implement a machine learning model to identify potential new customers on social media?

- A. Use a recommendation engine on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media
- B. Use a decision tree classifier engine on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media
- C. Use regression on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media.
- D. Use clustering on customer profile data to understand key characteristics of consumer segments. Find similar profiles on social media.

### Answer: C

### NEW QUESTION # 66

A financial services company is building a robust serverless data lake on Amazon S3. The data lake should be flexible and meet the following requirements:

- \* Support querying old and new data on Amazon S3 through Amazon Athena and Amazon Redshift Spectrum
- \* Support event-driven ETL pipelines.
- \* Provide a quick and easy way to understand metadata.

Which approach meets these requirements?

- A. Use an AWS Glue crawler to crawl S3 data, an AWS Lambda function to trigger an AWS Batch job, and an external Apache Hive metastore to search and discover metadata.
- B. Use an AWS Glue crawler to crawl S3 data, an AWS Lambda function to trigger an AWS Glue ETL job, and an AWS Glue Data Catalog to search and discover metadata.
- C. Use an AWS Glue crawler to crawl S3 data, an Amazon CloudWatch alarm to trigger an AWS Batch job, and an AWS Glue Data Catalog to search and discover metadata.
- D. Use an AWS Glue crawler to crawl S3 data, an Amazon CloudWatch alarm to trigger an AWS Glue ETL job, and an external Apache Hive metastore to search and discover metadata.

### Answer: B

Explanation:

Explanation

To build a robust serverless data lake on Amazon S3 that meets the requirements, the financial services company should use the following AWS services:

AWS Glue crawler: This is a service that connects to a data store, progresses through a prioritized list of classifiers to determine the schema for the data, and then creates metadata tables in the AWS Glue Data Catalog1. The company can use an AWS Glue crawler to crawl the S3 data and infer the schema, format, and partition structure of the data. The crawler can also detect schema changes and update the metadata tables accordingly. This enables the company to support querying old and new data on Amazon S3 through Amazon Athena and Amazon Redshift Spectrum, which are serverless interactive query services that use the AWS Glue Data Catalog as a central location for storing and retrieving table metadata23.

AWS Lambda function: This is a service that lets you run code without provisioning or managing servers. You pay only for the compute time you consume - there is no charge when your code is not running. You can also use AWS Lambda to create event-driven ETL pipelines, by triggering other AWS services based on events such as object creation or deletion in S3 buckets4. The

company can use an AWS Lambda function to trigger an AWS Glue ETL job, which is a serverless way to extract, transform, and load data for analytics. The AWS Glue ETL job can perform various data processing tasks, such as converting data formats, filtering, aggregating, joining, and more.

**AWS Glue Data Catalog:** This is a managed service that acts as a central metadata repository for data assets across AWS and on-premises data sources. The AWS Glue Data Catalog provides a uniform repository where disparate systems can store and find metadata to keep track of data in data silos, and use that metadata to query and transform the data. The company can use the AWS Glue Data Catalog to search and discover metadata, such as table definitions, schemas, and partitions. The AWS Glue Data Catalog also integrates with Amazon Athena, Amazon Redshift Spectrum, Amazon EMR, and AWS Glue ETL jobs, providing a consistent view of the data across different query and analysis services.

References:

- 1: What Is a Crawler? - AWS Glue
- 2: What Is Amazon Athena? - Amazon Athena
- 3: Amazon Redshift Spectrum - Amazon Redshift
- 4: What is AWS Lambda? - AWS Lambda
- 5: AWS Glue ETL Jobs - AWS Glue
- 6: What Is the AWS Glue Data Catalog? - AWS Glue

### NEW QUESTION # 67

A Machine Learning Specialist observes several performance problems with the training portion of a machine learning solution on Amazon SageMaker. The solution uses a large training dataset 2 TB in size and is using the SageMaker k-means algorithm. The observed issues include the unacceptable length of time it takes before the training job launches and poor I/O throughput while training the model. What should the Specialist do to address the performance issues with the current solution?

- A. Ensure that the input mode for the training job is set to Pipe.
- B. Use the SageMaker batch transform feature
- **C. Compress the training data into Apache Parquet format.**
- D. Copy the training dataset to an Amazon EFS volume mounted on the SageMaker instance.

**Answer: C**

### NEW QUESTION # 68

A company wants to segment a large group of customers into subgroups based on shared characteristics. The company's data scientist is planning to use the Amazon SageMaker built-in k-means clustering algorithm for this task. The data scientist needs to determine the optimal number of subgroups (k) to use.

Which data visualization approach will MOST accurately determine the optimal value of k?

- A. Calculate the principal component analysis (PCA) components. Run the k-means clustering algorithm for a range of k by using only the first two PCA components. For each value of k, create a scatter plot with a different color for each cluster. The optimal value of k is the value where the clusters start to look reasonably separated.
- B. Calculate the principal component analysis (PCA) components. Create a line plot of the number of components against the explained variance. The optimal value of k is the number of PCA components after which the curve starts decreasing in a linear fashion.
- **C. Run the k-means clustering algorithm for a range of k. For each value of k, calculate the sum of squared errors (SSE). Plot a line chart of the SSE for each value of k. The optimal value of k is the point after which the curve starts decreasing in a linear fashion.**
- D. Create a t-distributed stochastic neighbor embedding (t-SNE) plot for a range of perplexity values. The optimal value of k is the value of perplexity, where the clusters start to look reasonably separated.

**Answer: C**

Explanation:

Explanation

The solution D is the best data visualization approach to determine the optimal value of k for the k-means clustering algorithm. The solution D involves the following steps:

Run the k-means clustering algorithm for a range of k. For each value of k, calculate the sum of squared errors (SSE). The SSE is a measure of how well the clusters fit the data. It is calculated by summing the squared distances of each data point to its closest cluster center. A lower SSE indicates a better fit, but it will always decrease as the number of clusters increases. Therefore, the goal is to find the smallest value of k that still has a low SSE1.

Plot a line chart of the SSE for each value of k. The line chart will show how the SSE changes as the value of k increases. Typically, the line chart will have a shape of an elbow, where the SSE drops rapidly at first and then levels off. The optimal value of k is the point after which the curve starts decreasing in a linear fashion. This point is also known as the elbow point, and it represents the balance between the number of clusters and the SSE<sup>1</sup>.

The other options are not suitable because:

Option A: Calculating the principal component analysis (PCA) components, running the k-means clustering algorithm for a range of k by using only the first two PCA components, and creating a scatter plot with a different color for each cluster will not accurately determine the optimal value of k. PCA is a technique that reduces the dimensionality of the data by transforming it into a new set of features that capture the most variance in the data. However, PCA may not preserve the original structure and distances of the data, and it may lose some information in the process. Therefore, running the k-means clustering algorithm on the PCA components may not reflect the true clusters in the data. Moreover, using only the first two PCA components may not capture enough variance to represent the data well. Furthermore, creating a scatter plot may not be reliable, as it depends on the subjective judgment of the data scientist to decide when the clusters look reasonably separated<sup>2</sup>.

Option B: Calculating the PCA components and creating a line plot of the number of components against the explained variance will not determine the optimal value of k. This approach is used to determine the optimal number of PCA components to use for dimensionality reduction, not for clustering. The explained variance is the ratio of the variance of each PCA component to the total variance of the data. The optimal number of PCA components is the point where adding more components does not significantly increase the explained variance. However, this number may not correspond to the optimal number of clusters, as PCA and k-means clustering have different objectives and assumptions<sup>2</sup>.

Option C: Creating a t-distributed stochastic neighbor embedding (t-SNE) plot for a range of perplexity values will not determine the optimal value of k. t-SNE is a technique that reduces the dimensionality of the data by embedding it into a lower-dimensional space, such as a two-dimensional plane. t-SNE preserves the local structure and distances of the data, and it can reveal clusters and patterns in the data.

However, t-SNE does not assign labels or centroids to the clusters, and it does not provide a measure of how well the clusters fit the data. Therefore, t-SNE cannot determine the optimal number of clusters, as it only visualizes the data. Moreover, t-SNE depends on the perplexity parameter, which is a measure of how many neighbors each point considers. The perplexity parameter can affect the shape and size of the clusters, and there is no optimal value for it. Therefore, creating a t-SNE plot for a range of perplexity values may not be consistent or reliable<sup>3</sup>.

References:

1: How to Determine the Optimal K for K-Means?

2: Principal Component Analysis

3: t-Distributed Stochastic Neighbor Embedding

## NEW QUESTION # 69

A Machine Learning Specialist is developing a daily ETL workflow containing multiple ETL jobs.

The workflow consists of the following processes:

- Start the workflow as soon as data is uploaded to Amazon S3.
- When all the datasets are available in Amazon S3, start an ETL job to join the uploaded datasets with multiple terabyte-sized datasets already stored in Amazon S3.
- Store the results of joining datasets in Amazon S3.
- If one of the jobs fails, send a notification to the Administrator.

Which configuration will meet these requirements?

- A. Develop the ETL workflow using AWS Batch to trigger the start of ETL jobs when data is uploaded to Amazon S3. Use AWS Glue to join the datasets in Amazon S3. Use an Amazon CloudWatch alarm to send an SNS notification to the Administrator in the case of a failure.
- B. Use AWS Lambda to trigger an AWS Step Functions workflow to wait for dataset uploads to complete in Amazon S3. Use AWS Glue to join the datasets. Use an Amazon CloudWatch alarm to send an SNS notification to the Administrator in the case of a failure.
- C. Develop the ETL workflow using AWS Lambda to start an Amazon SageMaker notebook instance. Use a lifecycle configuration script to join the datasets and persist the results in Amazon S3. Use an Amazon CloudWatch alarm to send an SNS notification to the Administrator in the case of a failure.
- D. Use AWS Lambda to chain other Lambda functions to read and join the datasets in Amazon S3 as soon as the data is uploaded to Amazon S3. Use an Amazon CloudWatch alarm to send an SNS notification to the Administrator in the case of a failure.

## Answer: B

Explanation:

<https://aws.amazon.com/step-functions/use-cases/>

## NEW QUESTION # 70

Without a doubt, there is one thing that can assist them with perceiving this interest and clearing their AWS Certified Machine Learning - Specialty (MLS-C01) exam with flying colors. Amazon MLS-C01 dumps merge all that gigantic and the competitor doesn't require to purchase the aide or different books to review. They have this test material and need nothing else for planning AWS Certified Machine Learning - Specialty exam.

Training MLS-C01 Pdf: <https://www.testvalid.com/MLS-C01-exam-collection.html>

BTW, DOWNLOAD part of TestValid MLS-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1O-7C8B685JmpJEGBtldy6tKKt6HCrB2M>