



# MLS-C01<sup>Q&As</sup>

AWS Certified Machine Learning - Specialty (MLS-C01)





## Pass Amazon MLS-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/aws-certified-machine-learning-specialty.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

A machine learning (ML) specialist is administering a production Amazon SageMaker endpoint with model monitoring configured. Amazon SageMaker Model Monitor detects violations on the SageMaker endpoint, so the ML specialist retrains the model with the latest dataset. This dataset is statistically representative of the current production traffic. The ML specialist notices that even after deploying the new SageMaker model and running the first monitoring job, the SageMaker endpoint still has violations.

What should the ML specialist do to resolve the violations?

- A. Manually trigger the monitoring job to re-evaluate the SageMaker endpoint traffic sample.
- B. Run the Model Monitor baseline job again on the new training set. Configure Model Monitor to use the new baseline.
- C. Delete the endpoint and recreate it with the original configuration.
- D. Retrain the model again by using a combination of the original training set and the new training set.

Correct Answer: B

<https://docs.aws.amazon.com/sagemaker/latest/dg/model-monitor-create-baseline.html>

---

### QUESTION 2

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

Correct Answer: AC

Key Words:

1.  
100 continuous numerical features – too many features
- 2.



No feature selection has been done

3.

Easy interpretation - direct relationship between X and Y are preferred

4.

gap between the training and validation set accuracy – overfitting

A: L1 regularization solves overfitting, interpretation is easy, direct relationships between x and y

B: More features, Overfitting will be worse.

C: Recursive feature elimination solves overfitting, interpretation is easy, direct relationships between x and y

D: Perform t-distributed stochastic neighbor embedding (t-SNE)= Amazon's favorite dimensionality reduction technique, frequently show up in the questions. However, same as PCA, less interpretable. You won't be able to see the direct impact of relevant features on the model outcome.

E: If you have more than two classes then Linear Discriminant Analysis is the preferred linear classification technique.

---

### QUESTION 3

A Data Scientist received a set of insurance records, each consisting of a record ID, the final outcome among 200 categories, and the date of the final outcome. Some partial information on claim contents is also provided, but only for a few of the 200 categories. For each outcome category, there are hundreds of records distributed over the past 3 years. The Data Scientist wants to predict how many claims to expect in each category from month to month, a few months in advance.

What type of machine learning model should be used?

A. Classification month-to-month using supervised learning of the 200 categories based on claim contents.

B. Reinforcement learning using claim IDs and timestamps where the agent will identify how many claims in each category to expect from month to month.

C. Forecasting using claim IDs and timestamps to identify how many claims in each category to expect from month to month.

D. Classification with supervised learning of the categories for which partial information on claim contents is provided, and forecasting using claim IDs and timestamps for all other categories.

Correct Answer: C

---

### QUESTION 4

A company stores its documents in Amazon S3 with no predefined product categories. A data scientist needs to build a machine learning model to categorize the documents for all the company's products. Which solution will meet these requirements with the MOST operational efficiency?



- A. Build a custom clustering model. Create a Dockerfile and build a Docker image. Register the Docker image in Amazon Elastic Container Registry (Amazon ECR). Use the custom image in Amazon SageMaker to generate a trained model.
- B. Tokenize the data and transform the data into tabular data. Train an Amazon SageMaker k-means model to generate the product categories.
- C. Train an Amazon SageMaker Neural Topic Model (NTM) model to generate the product categories.
- D. Train an Amazon SageMaker Blazing Text model to generate the product categories.

Correct Answer: C

---

### QUESTION 5

A library is developing an automatic book-borrowing system that uses Amazon Rekognition. Images of library members' faces are stored in an Amazon S3 bucket. When members borrow books, the Amazon Rekognition CompareFaces API operation compares real faces against the stored faces in Amazon S3.

The library needs to improve security by making sure that images are encrypted at rest. Also, when the images are used with Amazon Rekognition, they need to be encrypted in transit. The library also must ensure that the images are not used to improve Amazon Rekognition as a service.

How should a machine learning specialist architect the solution to satisfy these requirements?

- A. Enable server-side encryption on the S3 bucket. Submit an AWS Support ticket to opt out of allowing images to be used for improving the service, and follow the process provided by AWS Support.
- B. Switch to using an Amazon Rekognition collection to store the images. Use the IndexFaces and SearchFacesByImage API operations instead of the CompareFaces API operation.
- C. Switch to using the AWS GovCloud (US) Region for Amazon S3 to store images and for Amazon Rekognition to compare faces. Set up a VPN connection and only call the Amazon Rekognition API operations through the VPN.
- D. Enable client-side encryption on the S3 bucket. Set up a VPN connection and only call the Amazon Rekognition API operations through the VPN.

Correct Answer: A

Rekognition API endpoints only support secure connections over HTTPS and all communication is encrypted in transit with TLS

[Latest MLS-C01 Dumps](#)

[MLS-C01 VCE Dumps](#)

[MLS-C01 Practice Test](#)