

# PROFESSIONAL-MACHINE- LEARNING-ENGINEER<sup>Q&As</sup>

Professional Machine Learning Engineer

**Pass Google PROFESSIONAL-MACHINE-LEARNING-  
ENGINEER Exam with 100% Guarantee**

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

<https://www.leads4pass.com/professional-machine-learning-engineer.html>

100% Passing Guarantee  
100% Money Back Assurance

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

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



**QUESTION 1**

You work for a toy manufacturer that has been experiencing a large increase in demand. You need to build an ML model to reduce the amount of time spent by quality control inspectors checking for product defects. Faster defect detection is a priority. The factory does not have reliable Wi-Fi. Your company wants to implement the new ML model as soon as possible. Which model should you use?

- A. AutoML Vision Edge mobile-high-accuracy-1 model
- B. AutoML Vision Edge mobile-low-latency-1 model
- C. AutoML Vision model
- D. AutoML Vision Edge mobile-versatile-1 model

Correct Answer: B

Hence faster defect detection is a priority, AutoML Vision Edge mobile-low-latency-1 model should be the choice. This model is designed to run efficiently on mobile devices and prioritize low latency, which means that it can provide fast defect detection without requiring a connection to the cloud. <https://cloud.google.com/vision/automl/docs/train-edge>

---

**QUESTION 2**

You want to rebuild your ML pipeline for structured data on Google Cloud. You are using PySpark to conduct data transformations at scale, but your pipelines are taking over 12 hours to run. To speed up development and pipeline run time, you want to use a serverless tool and SQL syntax. You have already moved your raw data into Cloud Storage. How should you build the pipeline on Google Cloud while meeting the speed and processing requirements?

- A. Use Data Fusion's GUI to build the transformation pipelines, and then write the data into BigQuery.
- B. Convert your PySpark into SparkSQL queries to transform the data, and then run your pipeline on Dataproc to write the data into BigQuery.
- C. Ingest your data into Cloud SQL, convert your PySpark commands into SQL queries to transform the data, and then use federated queries from BigQuery for machine learning.
- D. Ingest your data into BigQuery using BigQuery Load, convert your PySpark commands into BigQuery SQL queries to transform the data, and then write the transformations to a new table.

Correct Answer: D

---

**QUESTION 3**

You have been given a dataset with sales predictions based on your company's marketing activities. The data is structured and stored in BigQuery, and has been carefully managed by a team of data analysts. You need to prepare a report providing insights into the predictive capabilities of the data. You were asked to run several ML models with different levels of sophistication, including simple models and multilayered neural networks. You only have a few hours to gather the results of your experiments. Which Google Cloud tools should you use to complete this task in the most efficient and self-serviced way?

- A. Use BigQuery ML to run several regression models, and analyze their performance.

- B. Read the data from BigQuery using Dataproc, and run several models using SparkML.
- C. Use Vertex AI Workbench user-managed notebooks with scikit-learn code for a variety of ML algorithms and performance metrics.
- D. Train a custom TensorFlow model with Vertex AI, reading the data from BigQuery featuring a variety of ML algorithms.

Correct Answer: A

---

## QUESTION 4

You built a custom ML model using scikit-learn. Training time is taking longer than expected. You decide to migrate your model to Vertex AI Training, and you want to improve the model's training time. What should you try out first?

- A. Migrate your model to TensorFlow, and train it using Vertex AI Training.
- B. Train your model in a distributed mode using multiple Compute Engine VMs.
- C. Train your model with DLVM images on Vertex AI, and ensure that your code utilizes NumPy and SciPy internal methods whenever possible.
- D. Train your model using Vertex AI Training with GPUs.

Correct Answer: C

---

## QUESTION 5

You work for a large retailer and have been asked to segment your customers by their purchasing habits. The purchase history of all customers has been uploaded to BigQuery. You suspect that there may be several distinct customer segments, however you are unsure of how many, and you don't yet understand the commonalities in their behavior. You want to find the most efficient solution. What should you do?

- A. Create a k-means clustering model using BigQuery ML. Allow BigQuery to automatically optimize the number of clusters.
- B. Create a new dataset in Dataprep that references your BigQuery table. Use Dataprep to identify similarities within each column.
- C. Use the Data Labeling Service to label each customer record in BigQuery. Train a model on your labeled data using AutoML Tables. Review the evaluation metrics to understand whether there is an underlying pattern in the data.
- D. Get a list of the customer segments from your company's Marketing team. Use the Data Labeling Service to label each customer record in BigQuery according to the list. Analyze the distribution of labels in your dataset using Data Studio.

Correct Answer: A

<https://cloud.google.com/bigquery/docs/kmeans-tutorial>

[Latest PROFESSIONAL-MA](#) [PROFESSIONAL-MACHIN](#) [PROFESSIONAL-MACHIN](#)

[CHINE-LEARNING-ENGINEER Dumps](#)

[E-LEARNING-ENGINEER Exam Questions](#)

[E-LEARNING-ENGINEER Braindumps](#)