



# 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.passapply.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



VCE & PDF

PassApply.com

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

2024 Latest passapply PROFESSIONAL-MACHINE-LEARNING-ENGINEER

PDF and VCE dumps Download

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





### QUESTION 1

You are responsible for building a unified analytics environment across a variety of on-premises data marts. Your company is experiencing data quality and security challenges when integrating data across the servers, caused by the use of a wide range of disconnected tools and temporary solutions. You need a fully managed, cloud-native data integration service that will lower the total cost of work and reduce repetitive work. Some members on your team prefer a codeless interface for building Extract, Transform, Load (ETL) process. Which service should you use?

- A. Dataflow
- B. Dataprep
- C. Apache Flink
- D. Cloud Data Fusion

Correct Answer: D

<https://cloud.google.com/data-fusion>

---

### QUESTION 2

You are training a TensorFlow model on a structured dataset with 100 billion records stored in several CSV files. You need to improve the input/output execution performance. What should you do?

- A. Load the data into BigQuery, and read the data from BigQuery.
- B. Load the data into Cloud Bigtable, and read the data from Bigtable.
- C. Convert the CSV files into shards of TFRecords, and store the data in Cloud Storage.
- D. Convert the CSV files into shards of TFRecords, and store the data in the Hadoop Distributed File System (HDFS).

Correct Answer: C

Reference: <https://cloud.google.com/dataflow/docs/guides/templates/provided-batch>

---

### QUESTION 3

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 4

You are building a linear regression model on BigQuery ML to predict a customer's likelihood of purchasing your company's products. Your model uses a city name variable as a key predictive component. In order to train and serve the model, your data must be organized in columns. You want to prepare your data using the least amount of coding while maintaining the predictable variables. What should you do?

- A. Use TensorFlow to create a categorical variable with a vocabulary list. Create the vocabulary file, and upload it as part of your model to BigQuery ML.
- B. Create a new view with BigQuery that does not include a column with city information
- C. Use Cloud Data Fusion to assign each city to a region labeled as 1, 2, 3, 4, or 5, and then use that number to represent the city in the model.
- D. Use Dataprep to transform the state column using a one-hot encoding method, and make each city a column with binary values.

Correct Answer: B

<https://academic.oup.com/rheumatology/article/54/7/1141/1849688> <https://cloud.google.com/bigquery-ml/docs/reference/standard-sql/bigqueryml-auto-preprocessing>

#### QUESTION 5

You work as an ML engineer at a social media company, and you are developing a visual filter for users' profile photos. This requires you to train an ML model to detect bounding boxes around human faces. You want to use this filter in your company's iOS-based mobile phone application. You want to minimize code development and want the model to be optimized for inference on mobile phones. What should you do?

- A. Train a model using AutoML Vision and use the "export for Core ML" option.
- B. Train a model using AutoML Vision and use the "export for Coral" option.
- C. Train a model using AutoML Vision and use the "export for TensorFlow.js" option.
- D. Train a custom TensorFlow model and convert it to TensorFlow Lite (TFLite).

Correct Answer: A

<https://cloud.google.com/vision/automl/docs/export-edge> Core ML -> iOS and macOS Coral -> Edge TPU-based device TensorFlow.js -> web

[Latest PROFESSIONAL-MA](#) [PROFESSIONAL-MACHIN](#) [PROFESSIONAL-MACHIN](#)  
[CHINE-LEARNING-](#) [E-LEARNING-ENGINEER](#) [E-LEARNING-ENGINEER](#)



VCE & PDF

PassApply.com

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

2024 Latest passapply PROFESSIONAL-MACHINE-LEARNING-ENGINEER  
PDF and VCE dumps Download

---

[ENGINEER Dumps](#)

[Practice Test](#)

[Exam Questions](#)