



PROFESSIONAL-MACHINE- LEARNING-ENGINEER^{Q&As}

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 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>

QUESTION 2

You are working on a system log anomaly detection model for a cybersecurity organization. You have developed the model using TensorFlow, and you plan to use it for real-time prediction. You need to create a Dataflow pipeline to ingest data via Pub/Sub and write the results to BigQuery. You want to minimize the serving latency as much as possible. What should you do?

- A. Containerize the model prediction logic in Cloud Run, which is invoked by Dataflow.
- B. Load the model directly into the Dataflow job as a dependency, and use it for prediction.
- C. Deploy the model to a Vertex AI endpoint, and invoke this endpoint in the Dataflow job.
- D. Deploy the model in a TF Serving container on Google Kubernetes Engine, and invoke it in the Dataflow job.

Correct Answer: C

<https://cloud.google.com/architecture/minimizing-predictive-serving-latency-in-machine-learning>

QUESTION 3

While conducting an exploratory analysis of a dataset, you discover that categorical feature A has substantial predictive power, but it is sometimes missing. What should you do?

- A. Drop feature A if more than 15% of values are missing. Otherwise, use feature A as-is.
- B. Compute the mode of feature A and then use it to replace the missing values in feature A.



- C. Replace the missing values with the values of the feature with the highest Pearson correlation with feature A.
- D. Add an additional class to categorical feature A for missing values. Create a new binary feature that indicates whether feature A is missing.

Correct Answer: D

<https://developers.google.com/machine-learning/testing-debugging/common/data-errors>

QUESTION 4

Your data science team needs to rapidly experiment with various features, model architectures, and hyperparameters. They need to track the accuracy metrics for various experiments and use an API to query the metrics over time. What should they use to track and report their experiments while minimizing manual effort?

- A. Use Vertex AI Pipelines to execute the experiments. Query the results stored in MetadataStore using the Vertex AI API.
- B. Use Vertex AI Training to execute the experiments. Write the accuracy metrics to BigQuery, and query the results using the BigQuery API.
- C. Use Vertex AI Training to execute the experiments. Write the accuracy metrics to Cloud Monitoring, and query the results using the Monitoring API.
- D. Use Vertex AI Workbench user-managed notebooks to execute the experiments. Collect the results in a shared Google Sheets file, and query the results using the Google Sheets API.

Correct Answer: A

The Vertex AI Pipelines provide a powerful tool for automating machine learning workflows, including data preparation, training, and deployment. MetadataStore can be used to track the performance of different models by logging accuracy metrics and other important information. The Vertex AI API can then be used to query the metadata store and retrieve the results of different experiments.

QUESTION 5

You are developing an image recognition model using PyTorch based on ResNet50 architecture. Your code is working fine on your local laptop on a small subsample. Your full dataset has 200k labeled images. You want to quickly scale your training workload while minimizing cost. You plan to use 4 V100 GPUs. What should you do?

- A. Create a Google Kubernetes Engine cluster with a node pool that has 4 V100 GPUs. Prepare and submit a TFJob operator to this node pool.
- B. Create a Vertex AI Workbench user-managed notebooks instance with 4 V100 GPUs, and use it to train your model.
- C. Package your code with Setuptools, and use a pre-built container. Train your model with Vertex AI using a custom tier that contains the required GPUs.
- D. Configure a Compute Engine VM with all the dependencies that launches the training. Train your model with Vertex AI using a custom tier that contains the required GPUs.

Correct Answer: C



"Vertex AI provides flexible and scalable hardware and secured infrastructure to train PyTorch based deep learning models with pre-built containers and custom containers. (...) use PyTorch ResNet-50 as the example model and train it on ImageNet validation data (50K images) to measure the training performance for different training strategies":

<https://cloud.google.com/blog/products/ai-machine-learning/efficient-pytorch-training-with-vertex-ai>

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Practice Test](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Study Guide](#)

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions](#)