

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.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

Your company manages an application that aggregates news articles from many different online sources and sends them to users. You need to build a recommendation model that will suggest articles to readers that are similar to the articles they are currently reading. Which approach should you use?

- A. Create a collaborative filtering system that recommends articles to a user based on the user's past behavior.
- B. Encode all articles into vectors using word2vec, and build a model that returns articles based on vector similarity.
- C. Build a logistic regression model for each user that predicts whether an article should be recommended to a user.
- D. Manually label a few hundred articles, and then train an SVM classifier based on the manually classified articles that categorizes additional articles into their respective categories.

Correct Answer: B

<https://cloud.google.com/blog/topics/developers-practitioners/meet-ais-multitool-vector-embeddings>

QUESTION 2

You are an ML engineer at a bank that has a mobile application. Management has asked you to build an ML-based biometric authentication for the app that verifies a customer's identity based on their fingerprint. Fingerprints are considered highly sensitive personal information and cannot be downloaded and stored into the bank databases. Which learning strategy should you recommend to train and deploy this ML model?

- A. Data Loss Prevention API
- B. Federated learning
- C. MD5 to encrypt data
- D. Differential privacy

Correct Answer: B

With federated learning, all the data is collected, and the model is trained with algorithms across multiple decentralized edge devices such as cell phones or websites, without exchanging them.

QUESTION 3

You need to design a customized deep neural network in Keras that will predict customer purchases based on their purchase history. You want to explore model performance using multiple model architectures, store training data, and be able to compare the evaluation metrics in the same dashboard. What should you do?

- A. Create multiple models using AutoML Tables.
- B. Automate multiple training runs using Cloud Composer.
- C. Run multiple training jobs on AI Platform with similar job names.

D. Create an experiment in Kubeflow Pipelines to organize multiple runs.

Correct Answer: D

<https://www.kubeflow.org/docs/about/use-cases/>

QUESTION 4

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>

QUESTION 5

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>

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

[E-LEARNING-ENGINEER
VCE Dumps](#)

[E-LEARNING-ENGINEER
Study Guide](#)

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