



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 are an ML engineer responsible for designing and implementing training pipelines for ML models. You need to create an end-to-end training pipeline for a TensorFlow model. The TensorFlow model will be trained on several terabytes of structured data. You need the pipeline to include data quality checks before training and model quality checks after training but prior to deployment. You want to minimize development time and the need for infrastructure maintenance. How should you build and orchestrate your training pipeline?

- A. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components. Orchestrate the pipeline using Vertex AI Pipelines.
- B. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Vertex AI Pipelines.
- C. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.
- D. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.

Correct Answer: B

<https://cloud.google.com/vertex-ai/docs/pipelines/build-pipeline#sdk>

QUESTION 2

You recently joined an enterprise-scale company that has thousands of datasets. You know that there are accurate descriptions for each table in BigQuery, and you are searching for the proper BigQuery table to use for a model you are building on AI Platform. How should you find the data that you need?

- A. Use Data Catalog to search the BigQuery datasets by using keywords in the table description.
- B. Tag each of your model and version resources on AI Platform with the name of the BigQuery table that was used for training.
- C. Maintain a lookup table in BigQuery that maps the table descriptions to the table ID. Query the lookup table to find the correct table ID for the data that you need.
- D. Execute a query in BigQuery to retrieve all the existing table names in your project using the INFORMATION_SCHEMA metadata tables that are native to BigQuery. Use the result to find the table that you need.

Correct Answer: A

<https://cloud.google.com/data-catalog/docs/concepts/overview>

QUESTION 3

You are using transfer learning to train an image classifier based on a pre-trained EfficientNet model. Your training dataset has 20,000 images. You plan to retrain the model once per day. You need to minimize the cost of infrastructure. What platform components and configuration environment should you use?



- A. A Deep Learning VM with 4 V100 GPUs and local storage.
- B. A Deep Learning VM with 4 V100 GPUs and Cloud Storage.
- C. A Google Kubernetes Engine cluster with a V100 GPU Node Pool and an NFS Server
- D. An AI Platform Training job using a custom scale tier with 4 V100 GPUs and Cloud Storage

Correct Answer: D

QUESTION 4

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 5

You work for a gaming company that develops massively multiplayer online (MMO) games. You built a TensorFlow model that predicts whether players will make in-app purchases of more than \$10 in the next two weeks. The model's predictions will be used to adapt each user's game experience. User data is stored in BigQuery. How should you serve your model while optimizing cost, user experience, and ease of management?

- A. Import the model into BigQuery ML. Make predictions using batch reading data from BigQuery, and push the data to Cloud SQL
- B. Deploy the model to Vertex AI Prediction. Make predictions using batch reading data from Cloud Bigtable, and push the data to Cloud SQL.
- C. Embed the model in the mobile application. Make predictions after every in-app purchase event is published in Pub/Sub, and push the data to Cloud SQL.
- D. Embed the model in the streaming Dataflow pipeline. Make predictions after every in-app purchase event is published in Pub/Sub, and push the data to Cloud SQL.

Correct Answer: A

<https://cloud.google.com/bigquery-ml/docs/reference/standard-sql/bigqueryml-syntax-create-tensorflow>



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

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps](#)

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

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