



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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-machine-learning-engineer.html>

2024 Latest pass4itsure 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 rapidly growing social media company. Your team builds TensorFlow recommender models in an on-premises CPU cluster. The data contains billions of historical user events and 100,000 categorical features. You notice that as the data increases, the model training time increases. You plan to move the models to Google Cloud. You want to use the most scalable approach that also minimizes training time. What should you do?

- A. Deploy the training jobs by using TPU VMs with TPUv3 Pod slices, and use the TPUEmbedding API
- B. Deploy the training jobs in an autoscaling Google Kubernetes Engine cluster with CPUs
- C. Deploy a matrix factorization model training job by using BigQuery ML
- D. Deploy the training jobs by using Compute Engine instances with A100 GPUs, and use the tf.nn.embedding_lookup API

Correct Answer: A

QUESTION 2

You have recently used TensorFlow to train a classification model on tabular data. You have created a Dataflow pipeline that can transform several terabytes of data into training or prediction datasets consisting of TFRecords. You now need to productionize the model, and you want the predictions to be automatically uploaded to a BigQuery table on a weekly schedule. What should you do?

- A. Import the model into Vertex AI and deploy it to a Vertex AI endpoint. On Vertex AI Pipelines, create a pipeline that uses the DataflowPythonJobOp and the ModelBatchPredictOp components.
- B. Import the model into Vertex AI and deploy it to a Vertex AI endpoint. Create a Dataflow pipeline that reuses the data processing logic sends requests to the endpoint, and then uploads predictions to a BigQuery table.
- C. Import the model into Vertex AI. On Vertex AI Pipelines, create a pipeline that uses the DataflowPythonJobOp and the ModelBatchPredictOp components.
- D. Import the model into BigQuery. Implement the data processing logic in a SQL query. On Vertex AI Pipelines create a pipeline that uses the BigQueryQueryJobOp and the BigQueryPredictModelJobOp components.

Correct Answer: C

QUESTION 3

You need to train a computer vision model that predicts the type of government ID present in a given image using a GPU-powered virtual machine on Compute Engine. You use the following parameters: Optimizer: SGD Batch size = 64 Epochs = 10 Verbose = 2

During training you encounter the following error: ResourceExhaustedError: Out Of Memory (OOM) when allocating tensor. What should you do?

- A. Change the optimizer.



- B. Reduce the batch size.
- C. Change the learning rate.
- D. Reduce the image shape.

Correct Answer: B

Reference: <https://github.com/tensorflow/tensorflow/issues/136>

QUESTION 4

You work for an online grocery store. You recently developed a custom ML model that recommends a recipe when a user arrives at the website. You chose the machine type on the Vertex AI endpoint to optimize costs by using the queries per second (QPS) that the model can serve, and you deployed it on a single machine with 8 vCPUs and no accelerators.

A holiday season is approaching and you anticipate four times more traffic during this time than the typical daily traffic. You need to ensure that the model can scale efficiently to the increased demand. What should you do?

A. 1. Maintain the same machine type on the endpoint.

2.

Set up a monitoring job and an alert for CPU usage.

3.

If you receive an alert, add a compute node to the endpoint.

B. 1. Change the machine type on the endpoint to have 32 vCPUs.

2.

Set up a monitoring job and an alert for CPU usage.

3.

If you receive an alert, scale the vCPUs further as needed.

C. 1. Maintain the same machine type on the endpoint Configure the endpoint to enable autoscaling based on vCPU usage.

2.

Set up a monitoring job and an alert for CPU usage.

3.

If you receive an alert, investigate the cause.

D. 1. Change the machine type on the endpoint to have a GPU. Configure the endpoint to enable autoscaling based on the GPU usage.

2.



Set up a monitoring job and an alert for GPU usage.

3.

If you receive an alert, investigate the cause.

Correct Answer: C

QUESTION 5

You recently developed a deep learning model using Keras, and now you are experimenting with different training strategies. First, you trained the model using a single GPU, but the training process was too slow. Next, you distributed the training across 4 GPUs using `tf.distribute.MirroredStrategy` (with no other changes), but you did not observe a decrease in training time. What should you do?

- A. Distribute the dataset with `tf.distribute.Strategy.experimental_distribute_dataset`
- B. Create a custom training loop.
- C. Use a TPU with `tf.distribute.TPUStrategy`.
- D. Increase the batch size.

Correct Answer: D

https://www.tensorflow.org/guide/gpu_performance_analysis

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

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

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps](#)