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

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps

1 / 4

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps

2 / 4

QUESTION 1

You are developing models to classify customer support emails. You created models with TensorFlow Estimators using small datasets on your on-premises system, but you now need to train the models using large datasets to ensure high performance. You will port your models to Google Cloud and want to minimize code refactoring and infrastructure overhead for easier migration from on-prem to cloud. What should you do?

A. Use AI Platform for distributed training.

B. Create a cluster on Dataproc for training.

C. Create a Managed Instance Group with autoscaling.

D. Use Kubeflow Pipelines to train on a Google Kubernetes Engine cluster.

Correct Answer: A

QUESTION 2

You work for a large technology company that wants to modernize their contact center. You have been asked to develop a solution to classify incoming calls by product so that requests can be more quickly routed to the correct support team. You have already transcribed the calls using the Speech-to-Text API. You want to minimize data preprocessing and development time. How should you build the model?

A. Use the AI Platform Training built-in algorithms to create a custom model.

B. Use AutoMIL Natural Language to extract custom entities for classification.

C. Use the Cloud Natural Language API to extract custom entities for classification.

D. Build a custom model to identify the product keywords from the transcribed calls, and then run the keywords through a classification algorithm.

Correct Answer: B

QUESTION 3

Your data science team has requested a system that supports scheduled model retraining, Docker containers, and a service that supports autoscaling and monitoring for online prediction requests. Which platform components should you choose for this system?

A. Vertex AI Pipelines and App Engine

B. Vertex AI Pipelines, Vertex AI Prediction, and Vertex AI Model Monitoring

C. Cloud Composer, BigQuery ML, and Vertex AI Prediction

D. Cloud Composer, Vertex AI Training with custom containers, and App Engine

Correct Answer: B

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps

3 / 4

https://cloud.google.com/vertex-ai/docs/training/containers-overview

QUESTION 4

You are experimenting with a built-in distributed XGBoost model in Vertex AI Workbench user-managed notebooks. You use BigQuery to split your data into training and validation sets using the following queries:

CREATE OR REPLACE TABLE `myproject.mydataset.training` AS (SELECT * FROM `myproject.mydataset.mytable` WHERE RAND()

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Exam Questions |
PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps

4 / 4