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.geekcert.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

QUESTION 1

You are creating a model training pipeline to predict sentiment scores from text-based product reviews. You want to have control over how the model parameters are tuned, and you will deploy the model to an endpoint after it has been trained. You will use Vertex AI Pipelines to run the pipeline. You need to decide which Google Cloud pipeline components to use. What components should you choose?

A. TabularDatasetCreateOp, CustomTrainingJobOp, and EndpointCreateOp

B. TextDatasetCreateOp, AutoMLTextTrainingOp, and EndpointCreateOp

C. TabularDatasetCreateOp. AutoMLTextTrainingOp, and ModelDeployOp

D. TextDatasetCreateOp, CustomTrainingJobOp, and ModelDeployOp

Correct Answer: D

QUESTION 2

You recently deployed a model to a Vertex AI endpoint and set up online serving in Vertex AI Feature Store. You have configured a daily batch ingestion job to update your featurestore. During the batch ingestion jobs, you discover that CPU utilization is high in your featurestore\\'s online serving nodes and that feature retrieval latency is high. You need to improve online serving performance during the daily batch ingestion. What should you do?

A. Schedule an increase in the number of online serving nodes in your featurestore prior to the batch ingestion jobs

B. Enable autoscaling of the online serving nodes in your featurestore

C. Enable autoscaling for the prediction nodes of your DeployedModel in the Vertex AI endpoint

D. Increase the worker_count in the ImportFeatureValues request of your batch ingestion job

Correct Answer: B

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 training and deploying updated versions of a regression model with tabular data by using Vertex AI Pipelines, Vertex AI Training, Vertex AI Experiments, and Vertex AI Endpoints. The model is deployed in a Vertex AI endpoint, and your users call the model by using the Vertex AI endpoint. You want to receive an email when the feature data distribution changes significantly, so you can retrigger the training pipeline and deploy an updated version of your model. What should you do?

A. Use Vertex Al Model Monitoring. Enable prediction drift monitoring on the endpoint, and specify a notification email.

B. In Cloud Logging, create a logs-based alert using the logs in the Vertex Al endpoint. Configure Cloud Logging to send an email when the alert is triggered.

C. In Cloud Monitoring create a logs-based metric and a threshold alert for the metric. Configure Cloud Monitoring to send an email when the alert is triggered.

D. Export the container logs of the endpoint to BigQuery. Create a Cloud Function to run a SQL query over the exported logs and send an email. Use Cloud Scheduler to trigger the Cloud Function.

Correct Answer: A

QUESTION 5

You work for an online retail company that is creating a visual search engine. You have set up an end-to-end ML pipeline on Google Cloud to classify whether an image contains your company\\'s product. Expecting the release of new products in the near future, you configured a retraining functionality in the pipeline so that new data can be fed into your ML models. You also want to use AI Platform\\'s continuous evaluation service to ensure that the models have high accuracy on your test dataset. What should you do?

A. Keep the original test dataset unchanged even if newer products are incorporated into retraining.

B. Extend your test dataset with images of the newer products when they are introduced to retraining.

C. Replace your test dataset with images of the newer products when they are introduced to retraining.

D. Update your test dataset with images of the newer products when your evaluation metrics drop below a pre-decided threshold.

Correct Answer: B

You need to correctly classify newer products, so you need the new training data ==> A is wrong;

You need to keep doing a good job on older dataset, you can\\'t just ignore it ==> C is wrong;

You know when you are introducing new products, there is no need to wait for a drop in preformaces ==> D is wrong;

B is correct