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 work for a toy manufacturer that has been experiencing a large increase in demand. You need to build an ML model to reduce the amount of time spent by quality control inspectors checking for product defects. Faster defect detection is a priority. The factory does not have reliable Wi-Fi. Your company wants to implement the new ML model as soon as possible. Which model should you use?

A. AutoML Vision Edge mobile-high-accuracy-1 model

B. AutoML Vision Edge mobile-low-latency-1 model

C. AutoML Vision model

D. AutoML Vision Edge mobile-versatile-1 model

Correct Answer: B

Hence faster defect detection is a priority, AutoML Vision Edge mobile-low-latency-1 model should be the choice. This model is designed to run efficiently on mobile devices and prioritize low latency, which means that it can provide fast defect detection without requiring a connection to the cloud. https://cloud.google.com/vision/automl/docs/train-edge

QUESTION 2

You are deploying a new version of a model to a production Vertex AI endpoint that is serving traffic. You plan to direct all user traffic to the new model. You need to deploy the model with minimal disruption to your application. What should you do?

A. 1. Create a new endpoint

2.

Create a new model. Set it as the default version. Upload the model to Vertex AI Model Registry

3.

Deploy the new model to the new endpoint

4.

Update Cloud DNS to point to the new endpoint

B. 1. Create a new endpoint

2.

Create a new model. Set the parentModel parameter to the model ID of the currently deployed model and set it as the default version. Upload the model to Vertex AI Model Registry

3.

Deploy the new model to the new endpoint, and set the new model to 100% of the traffic.

C. 1. Create a new model. Set the parentModel parameter to the model ID of the currently deployed model. Upload the

model to Vertex AI Model Registry.

2. Deploy the new model to the existing endpoint, and set the new model to 100% of the traffic

D. 1. Create a new model. Set it as the default version. Upload the model to Vertex AI Model Registry

2. Deploy the new model to the existing endpoint

Correct Answer: C

QUESTION 3

Your work for a textile manufacturing company. Your company has hundreds of machines, and each machine has many sensors. Your team used the sensory data to build hundreds of ML models that detect machine anomalies. Models are retrained daily, and you need to deploy these models in a cost-effective way. The models must operate 24/7 without downtime and make sub millisecond predictions. What should you do?

A. Deploy a Dataflow batch pipeline and a Vertex AI Prediction endpoint.

B. Deploy a Dataflow batch pipeline with the RunInference API, and use model refresh.

C. Deploy a Dataflow streaming pipeline and a Vertex AI Prediction endpoint with autoscaling.

D. Deploy a Dataflow streaming pipeline with the RunInference API, and use automatic model refresh.

Correct Answer: D

QUESTION 4

You are training an object detection model using a Cloud TPU v2. Training time is taking longer than expected. Based on this simplified trace obtained with a Cloud TPU profile, what action should you take to decrease training time in a cost-efficient way?

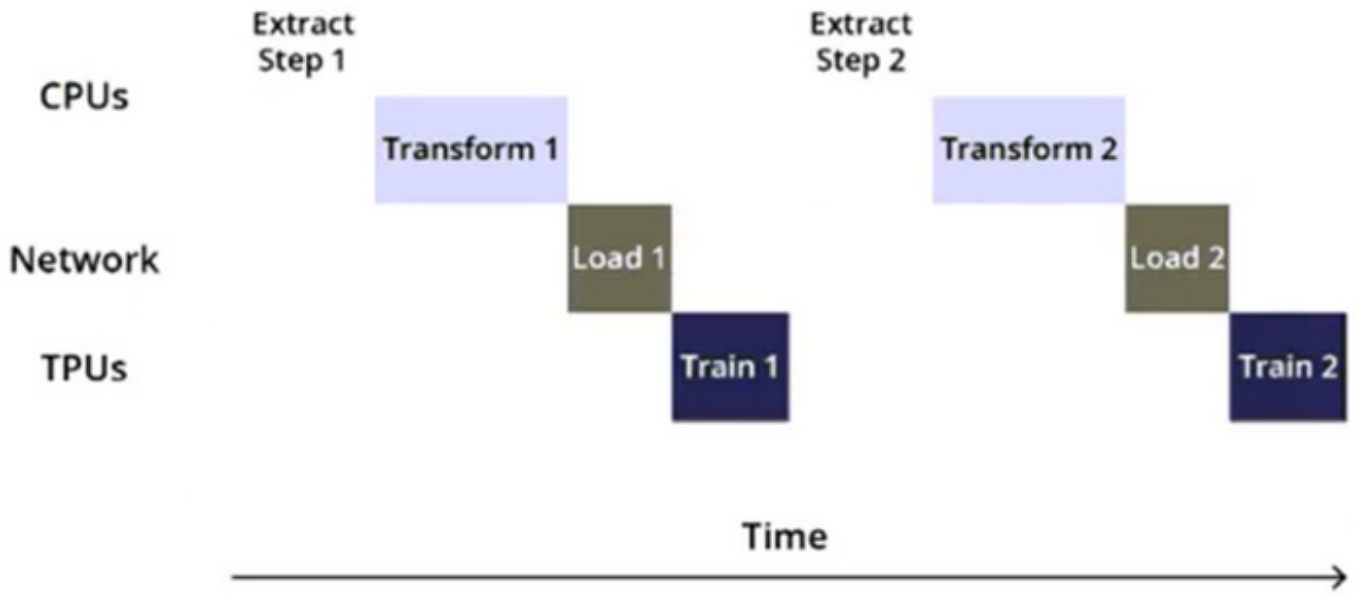

A. Move from Cloud TPU v2 to Cloud TPU v3 and increase batch size.

B. Move from Cloud TPU v2 to 8 NVIDIA V100 GPUs and increase batch size.

C. Rewrite your input function to resize and reshape the input images.

D. Rewrite your input function using parallel reads, parallel processing, and prefetch.

Correct Answer: D

Based on the profile, it appears that the Compute time is relatively low compared to the HostToDevice and DeviceToHost time. This suggests that the data transfer between the host (CPU) and the TPU device is a bottleneck. Therefore, the best action to decrease training time in a cost-efficient way would be to reduce the amount of data transferred between the host and the device.

QUESTION 5

You need to build classification workflows over several structured datasets currently stored in BigQuery. Because you will be performing the classification several times, you want to complete the following steps without writing code: exploratory data analysis, feature selection, model building, training, and hyperparameter tuning and serving. What should you do?

A. Train a TensorFlow model on Vertex AI.

B. Train a classification Vertex AutoML model.

C. Run a logistic regression job on BigQuery ML.

D. Use scikit-learn in Notebooks with pandas library.

Correct Answer: B