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 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 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

You are developing an ML model to identify your company\\'s products in images. You have access to over one million images in a Cloud Storage bucket. You plan to experiment with different TensorFlow models by using Vertex AI Training. You need to read images at scale during training while minimizing data I/O bottlenecks. What should you do?

A. Load the images directly into the Vertex AI compute nodes by using Cloud Storage FUSE. Read the images by using the tf.data.Dataset.from_tensor_slices function

B. Create a Vertex AI managed dataset from your image data. Access the AIP_TRAINING_DATA_URI environment variable to read the images by using the tf.data.Dataset.list_files function.

C. Convert the images to TFRecords and store them in a Cloud Storage bucket. Read the TFRecords by using the tf.data.TFRecordDataset function.

D. Store the URLs of the images in a CSV file. Read the file by using the tf.data.experimental.CsvDataset function.

Correct Answer: C

QUESTION 4

You recently deployed a pipeline in Vertex AI Pipelines that trains and pushes a model to a Vertex AI endpoint to serve real-time traffic. You need to continue experimenting and iterating on your pipeline to improve model performance. You plan to use Cloud Build for CI/CD You want to quickly and easily deploy new pipelines into production, and you want to minimize the chance that the new pipeline implementations will break in production. What should you do?

A. Set up a CI/CD pipeline that builds and tests your source code. If the tests are successful, use the Google. Cloud console to upload the built container to Artifact Registry and upload the compiled pipeline to Vertex AI Pipelines.

B. Set up a CI/CD pipeline that builds your source code and then deploys built artifacts into a pre-production environment. Run unit tests in the pre-production environment. If the tests are successful deploy the pipeline to production.

C. Set up a CI/CD pipeline that builds and tests your source code and then deploys built artifacts into a pre-production environment. After a successful pipeline run in the pre-production environment, deploy the pipeline to production.

D. Set up a CI/CD pipeline that builds and tests your source code and then deploys built artifacts into a pre-production environment. After a successful pipeline run in the pre-production environment, rebuild the source code and deploy the artifacts to production.

Correct Answer: C

QUESTION 5

You recently built the first version of an image segmentation model for a self-driving car. After deploying the model, you observe a decrease in the area under the curve (AUC) metric. When analyzing the video recordings, you also discover that the model fails in highly congested traffic but works as expected when there is less traffic. What is the most likely reason for this result?

A. The model is overfitting in areas with less traffic and underfitting in areas with more traffic.

B. AUC is not the correct metric to evaluate this classification model.

C. Too much data representing congested areas was used for model training.

D. Gradients become small and vanish while backpropagating from the output to input nodes.

Correct Answer: A

The most likely reason for this result is the model is overfitting in areas with less traffic and underfitting in areas with more traffic. Probably because the model was trained on a dataset that did not have enough examples of congested traffic. As a result, the model is not able to generalise well. When the model is validated on congested traffic, it makes mistakes because it has not seen this type of data before.

Latest PROFESSIONAL-MACHINE-LEARNING-ENGINEER Dumps

PROFESSIONAL-MACHINE-LEARNING-ENGINEER PDF Dumps

PROFESSIONAL-MACHINE-LEARNING-ENGINEER Braindumps