



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.passapply.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

PassApply.com

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

2024 Latest passapply 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

While conducting an exploratory analysis of a dataset, you discover that categorical feature A has substantial predictive power, but it is sometimes missing. What should you do?

- A. Drop feature A if more than 15% of values are missing. Otherwise, use feature A as-is.
- B. Compute the mode of feature A and then use it to replace the missing values in feature A.
- C. Replace the missing values with the values of the feature with the highest Pearson correlation with feature A.
- D. Add an additional class to categorical feature A for missing values. Create a new binary feature that indicates whether feature A is missing.

Correct Answer: D

<https://developers.google.com/machine-learning/testing-debugging/common/data-errors>

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

You are an ML engineer responsible for designing and implementing training pipelines for ML models. You need to create an end-to-end training pipeline for a TensorFlow model. The TensorFlow model will be trained on several terabytes of structured data. You need the pipeline to include data quality checks before training and model quality checks after training but prior to deployment. You want to minimize development time and the need for infrastructure maintenance. How should you build and orchestrate your training pipeline?

- A. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components. Orchestrate the pipeline using Vertex AI Pipelines.
- B. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Vertex AI Pipelines.
- C. Create the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud



components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.

D. Create the pipeline using TensorFlow Extended (TFX) and standard TFX components. Orchestrate the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine.

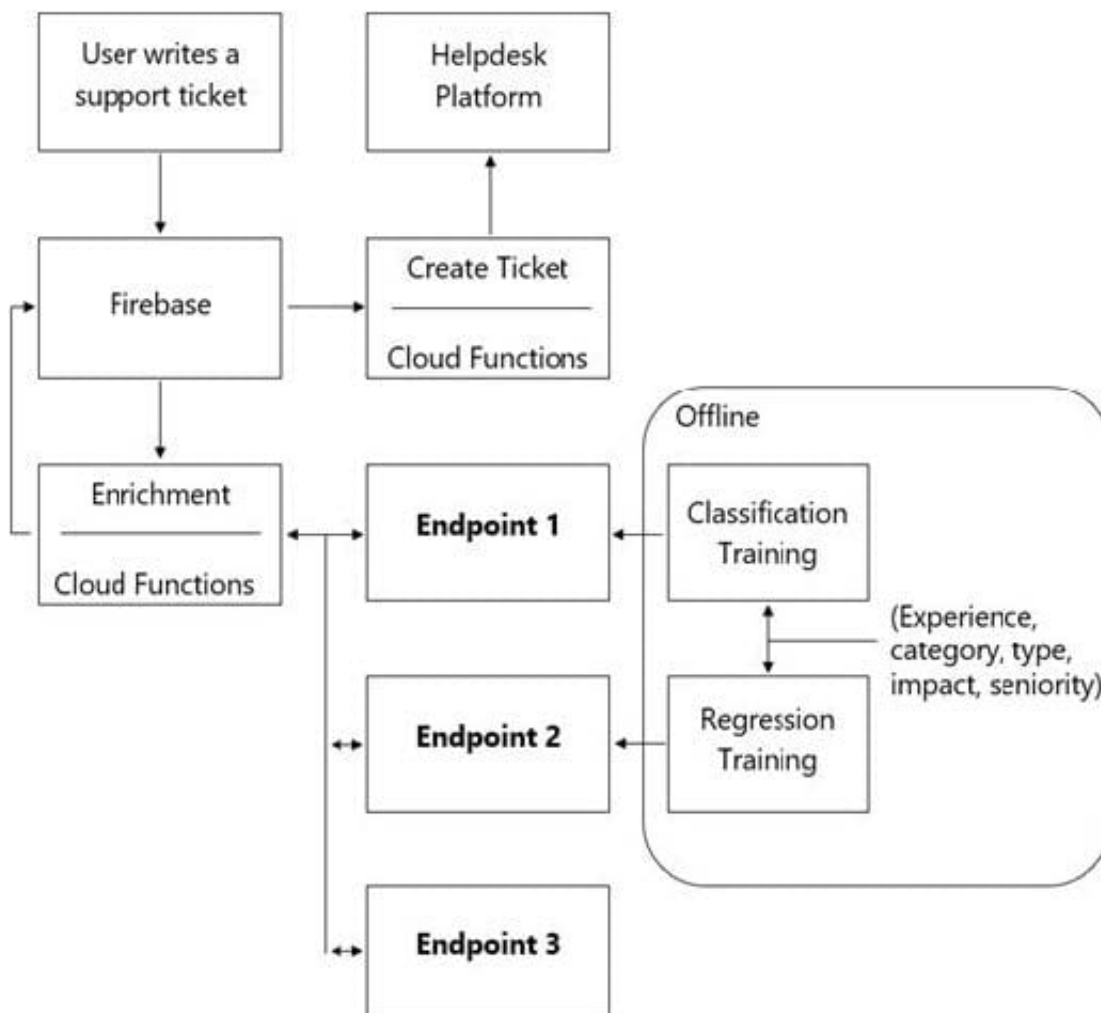
Correct Answer: B

<https://cloud.google.com/vertex-ai/docs/pipelines/build-pipeline#sdk>

QUESTION 4

You are designing an architecture with a serverless ML system to enrich customer support tickets with informative metadata before they are routed to a support agent. You need a set of models to predict ticket priority, predict ticket resolution time, and perform sentiment analysis to help agents make strategic decisions when they process support requests. Tickets are not expected to have any domain-specific terms or jargon.

The proposed architecture has the following flow:



Which endpoints should the Enrichment Cloud Functions call?

A. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Vision



- B. 1 = AI Platform, 2 = AI Platform, 3 = AutoML Natural Language
- C. 1 = AI Platform, 2 = AI Platform, 3 = Cloud Natural Language API
- D. 1 = Cloud Natural Language API, 2 = AI Platform, 3 = Cloud Vision API

Correct Answer: C

<https://cloud.google.com/architecture/architecture-of-a-serverless-ml-model#architecture>

The architecture has the following flow:

A user writes a ticket to Firebase, which triggers a Cloud Function.

-The Cloud Function calls 3 different endpoints to enrich the ticket:

-An AI Platform endpoint, where the function can predict the priority.

-An AI Platform endpoint, where the function can predict the resolution time.

-The Natural Language API to do sentiment analysis and word salience.

-For each reply, the Cloud Function updates the Firebase real-time database.

-The Cloud Function then creates a ticket into the helpdesk platform using the RESTful API.

QUESTION 5

You are training a deep learning model for semantic image segmentation with reduced training time. While using a Deep Learning VM Image, you receive the following error: The resource `\\projects/deeplearning-platform/zones/europe-west4c/acceleratorTypes/nvidia-tesla-k80\\` was not found. What should you do?

- A. Ensure that you have GPU quota in the selected region.
- B. Ensure that the required GPU is available in the selected region.
- C. Ensure that you have preemptible GPU quota in the selected region.
- D. Ensure that the selected GPU has enough GPU memory for the workload.

Correct Answer: B

https://cloud.google.com/deep-learning-vm/docs/troubleshooting#resource_not_found

<https://cloud.google.com/compute/docs/gpus/gpu-regions-zones>

[PROFESSIONAL-MACHINE-LEARNING-ENGINEER Practice Test](#)

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

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