

# High Professional-Machine-Learning-Engineer Passing Score | Latest Professional-Machine-Learning-Engineer Exam Camp



P.S. Free & New Professional-Machine-Learning-Engineer dumps are available on Google Drive shared by ValidDumps: <https://drive.google.com/open?id=1Asze8uniGKT1XE5A0c9y-FKGC18JqTFL>

We provide Professional-Machine-Learning-Engineer exam torrent which are of high quality and can boost high passing rate and hit rate. Our passing rate of Professional-Machine-Learning-Engineer training guide is 99% and thus you can reassure yourself to buy our product and enjoy the benefits brought by our Professional-Machine-Learning-Engineer exam materials. Our Professional-Machine-Learning-Engineer Learning Engine is efficient and can help you master the Professional-Machine-Learning-Engineer guide torrent in a short time and save your energy. The Professional-Machine-Learning-Engineer exam material we provide is compiled by experts and approved by the professionals who boost profound experiences.

## Exam Topics

The successful performance in the Google Professional Machine Learning Engineer certification test requires a good comprehension of its topics. The exam syllabus consists of six sections that are described below:

- **Framing Problems Related to Machine Learning**

Within this subject area, the candidates should be capable of translating business challenges into the Machine Learning use cases. They should also possess the skills in determining the Machine Learning problems, identifying the business success criteria, as well as defining risks to the feasibility of the Machine Learning solutions.

- **Developing Machine Learning Models**

To answer the questions related to this section, the learners should know how to build, test, and train models. They should also possess the skills in scaling model training as well as serving, including distributed training and scaling prediction service (for instance, containerized serving, AI Platform Prediction, etc.).

- **Automating & Orchestrating Machine Learning Pipelines**

This module encompasses one's competency in designing & implementing training pipelines. This includes your ability to define the components, triggers, parameters, and compute needs; understanding of the orchestration framework; familiarity with the multi-Cloud or hybrid strategies; knowledge of system design involving the TFX components/Kubeflow DSL. The candidates should also possess the skills in implementing serving pipelines, including serving (online, caching, batch), testing for target performance, configuring trigger & pipeline schedules, among other skills. Apart from that, this part requires the students' expertise in tracking & auditing metadata.

To become certified as a Google Professional Machine Learning Engineer, candidates must pass a two-hour long exam that consists of multiple-choice and scenario-based questions. Professional-Machine-Learning-Engineer Exam evaluates the candidate's proficiency in various areas of machine learning, such as data preparation, model development, optimization, and deployment.

Candidates must demonstrate their ability to use various machine learning tools and frameworks, such as TensorFlow, Keras, and Scikit-learn.

>> **High Professional-Machine-Learning-Engineer Passing Score** <<

## **Pass Guaranteed Quiz 2026 The Best Google High Professional-Machine-Learning-Engineer Passing Score**

As we all know, the world does not have two identical leaves. People's tastes also vary a lot. So we have tried our best to develop the three packages of our Professional-Machine-Learning-Engineer exam braindumps for you to choose. Now we have free demo of the Professional-Machine-Learning-Engineer study materials exactly according to the three packages on the website for you to download before you pay for the Professional-Machine-Learning-Engineer Practice Engine, and the free demos are a small part of the questions and answers. You can check the quality and validity by them.

### **Google Professional Machine Learning Engineer Sample Questions (Q283-Q288):**

#### **NEW QUESTION # 283**

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 featurestores 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. Enable autoscaling of the online serving nodes in your featurestore**
- B. Increase the worker counts in the importFeaturevalues request of your batch ingestion job.
- C. Schedule an increase in the number of online serving nodes in your featurestore prior to the batch ingestion jobs.
- D. Enable autoscaling for the prediction nodes of your DeployedModel in the Vertex AI endpoint.

**Answer: A**

#### **NEW QUESTION # 284**

You have written unit tests for a Kubeflow Pipeline that require custom libraries. You want to automate the execution of unit tests with each new push to your development branch in Cloud Source Repositories. What should you do?

- A. Set up a Cloud Logging sink to a Pub/Sub topic that captures interactions with Cloud Source Repositories. Execute the unit tests using a Cloud Function that is triggered when messages are sent to the Pub/Sub topic
- B. Write a script that sequentially performs the push to your development branch and executes the unit tests on Cloud Run
- **C. Using Cloud Build, set an automated trigger to execute the unit tests when changes are pushed to your development branch.**
- D. Set up a Cloud Logging sink to a Pub/Sub topic that captures interactions with Cloud Source Repositories. Configure a Pub/Sub trigger for Cloud Run, and execute the unit tests on Cloud Run.

**Answer: C**

Explanation:

Cloud Build is a service that executes your builds on Google Cloud Platform infrastructure. Cloud Build can import source code from Cloud Source Repositories, Cloud Storage, GitHub, or Bitbucket, execute a build to your specifications, and produce artifacts such as Docker containers or Java archives<sup>1</sup>. Cloud Build allows you to set up automated triggers that start a build when changes are pushed to a source code repository. You can configure triggers to filter the changes based on the branch, tag, or file path<sup>2</sup>. To automate the execution of unit tests for a Kubeflow Pipeline that require custom libraries, you can use Cloud Build to set an automated trigger to execute the unit tests when changes are pushed to your development branch in Cloud Source Repositories. You can specify the steps of the build in a YAML or JSON file, such as installing the custom libraries, running the unit tests, and reporting the results. You can also use Cloud Build to build and deploy the Kubeflow Pipeline components if the unit tests pass<sup>3</sup>. The other options are not recommended or feasible. Writing a script that sequentially performs the push to your development branch and executes the unit tests on Cloud Run is not a good practice, as it does not leverage the benefits of Cloud Build and its integration with Cloud Source Repositories. Setting up a Cloud Logging sink to a Pub/Sub topic that captures interactions with Cloud Source Repositories and using a Pub/Sub trigger for Cloud Run or Cloud Function to execute the unit tests is unnecessarily complex and inefficient, as it adds extra steps and latency to the process. Cloud Run and Cloud Function are also not designed for executing unit

tests, as they have limitations on the memory, CPU, and execution time<sup>45</sup> References: 1: Cloud Build overview 2: Creating and managing build triggers 3: Building and deploying Kubeflow Pipelines using Cloud Build 4: Cloud Run documentation 5: Cloud Functions documentation

### NEW QUESTION # 285

You are building a TensorFlow text-to-image generative model by using a dataset that contains billions of images with their respective captions. You want to create a low maintenance, automated workflow that reads the data from a Cloud Storage bucket collects statistics, splits the dataset into training/validation/test datasets performs data transformations, trains the model using the training/validation datasets. and validates the model by using the test dataset. What should you do?

- A. Use the MLFlow SDK and deploy it on a Google Kubernetes Engine Cluster Create multiple components that use Dataflow and Vertex AI services.
- **B. Use the TensorFlow Extended (TFX) SDK to create multiple components that use Dataflow and Vertex AI services Deploy the workflow on Vertex AI Pipelines.**
- C. Use the Kubeflow Pipelines (KFP) SDK to create multiple components that use Dataflow and Vertex AI services Deploy the workflow on Vertex AI Pipelines.
- D. Use the Apache Airflow SDK to create multiple operators that use Dataflow and Vertex AI services Deploy the workflow on Cloud Composer.

**Answer: B**

Explanation:

According to the web search results, TensorFlow Extended (TFX) is a platform for building end-to-end machine learning pipelines using TensorFlow<sup>1</sup>. TFX provides a set of components that can be orchestrated using either the TFX SDK or Kubeflow Pipelines. TFX components can handle different aspects of the pipeline, such as data ingestion, data validation, data transformation, model training, model evaluation, model serving, and more. TFX components can also leverage other Google Cloud services, such as Dataflow<sup>2</sup> and Vertex AI<sup>3</sup>. Dataflow is a fully managed service for running Apache Beam pipelines on Google Cloud.

Dataflow handles the provisioning and management of the compute resources, as well as the optimization and execution of the pipelines. Vertex AI is a unified platform for machine learning development and deployment.

Vertex AI offers various services and tools for building, managing, and serving machine learning models.

Therefore, option D is the best way to create a low maintenance, automated workflow for the given use case, as it allows you to use the TFX SDK to define and execute your pipeline components, and use Dataflow and Vertex AI services to scale and optimize your pipeline. The other options are not relevant or optimal for this scenario. References:

\* TensorFlow Extended

\* Dataflow

\* Vertex AI

\* Google Professional Machine Learning Certification Exam 2023

\* Latest Google Professional Machine Learning Engineer Actual Free Exam Questions

### NEW QUESTION # 286

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

**Answer: C**

Explanation:

The best option for creating and orchestrating an end-to-end training pipeline for a TensorFlow model is to use TensorFlow

Extended (TFX) and standard TFX components, and deploy the pipeline to Vertex AI Pipelines. TFX is an end-to-end platform for deploying production ML pipelines, which consists of several built-in components that cover the entire ML lifecycle, from data ingestion and validation, to model training and evaluation, to model deployment and monitoring. TFX also supports custom components and integrations with other Google Cloud services, such as BigQuery, Dataflow, and Cloud Storage. Vertex AI Pipelines is a fully managed service that allows you to run TFX pipelines on Google Cloud, without having to worry about infrastructure provisioning, scaling, or maintenance. Vertex AI Pipelines also provides a user-friendly interface to monitor and manage your pipelines, as well as tools to track and compare experiments. The other options are not as suitable for creating and orchestrating an end-to-end training pipeline for a TensorFlow model, because:

Creating the pipeline using Kubeflow Pipelines domain-specific language (DSL) and predefined Google Cloud components would require more development time and effort, as Kubeflow Pipelines DSL is not as expressive or compatible with TensorFlow as TFX. Predefined Google Cloud components might not cover all the stages of the ML lifecycle, and might not be optimized for TensorFlow models.

Orchestrating the pipeline using Kubeflow Pipelines deployed on Google Kubernetes Engine would require more infrastructure maintenance, as Kubeflow Pipelines is not a fully managed service, and you would have to provision and manage your own Kubernetes cluster. This would also incur more costs, as you would have to pay for the cluster resources, regardless of the pipeline usage. Reference:

TFX | ML Production Pipelines | TensorFlow

Vertex AI Pipelines | Google Cloud

Kubeflow Pipelines | Google Cloud

Google Cloud launches machine learning engineer certification

Google Professional Machine Learning Engineer Certification

Professional ML Engineer Exam Guide

#### NEW QUESTION # 287

You are building a MLOps platform to automate your company's ML experiments and model retraining. You need to organize the artifacts for dozens of pipelines. How should you store the pipelines' artifacts?

- A. Store parameters in Cloud SQL, store the models' source code in GitHub, and store the models' binaries in Cloud Storage.
- B. Store parameters in Vertex ML Metadata and store the models' source code and binaries in GitHub.
- **C. Store parameters in Vertex ML Metadata, store the models' source code in GitHub, and store the models' binaries in Cloud Storage.**
- D. Store parameters in Cloud SQL and store the models' source code and binaries in GitHub.

**Answer: C**

Explanation:

To organize the artifacts for dozens of pipelines, you should store the parameters in Vertex ML Metadata, store the models' source code in GitHub, and store the models' binaries in Cloud Storage. This option has the following advantages:

\* Vertex ML Metadata is a service that helps you track and manage the metadata of your ML workflows, such as datasets, models, metrics, and parameters<sup>1</sup>. It can also help you with data lineage, model versioning, and model performance monitoring<sup>2</sup>.

\* GitHub is a popular platform for hosting and collaborating on code repositories. It can help you manage the source code of your models, as well as the configuration files, scripts, and notebooks that are part of your ML pipelines<sup>3</sup>.

\* Cloud Storage is a scalable and durable object storage service that can store any type of data, including model binaries<sup>4</sup>. It can also integrate with other services, such as Vertex AI, Cloud Functions, and Cloud Run, to enable easy deployment and serving of your models<sup>5</sup>.

References:

\* 1: Introduction to Vertex ML Metadata | Vertex AI | Google Cloud

\* 2: Manage metadata for ML workflows | Vertex AI | Google Cloud

\* 3: GitHub - Where the world builds software

\* 4: Cloud Storage | Google Cloud

\* 5: Deploying models | Vertex AI | Google Cloud

#### NEW QUESTION # 288

.....

Since the content of the examination is also updating daily, you will need real and latest Google Professional-Machine-Learning-Engineer Exam Dumps to prepare successfully for the Professional-Machine-Learning-Engineer certification exam in a short time. People who don't study from updated Google Professional-Machine-Learning-Engineer Questions fail the examination and lose time.

