vmexam.com
#1 Cloud and VM Certification Guide

Google GCP-PMLE

GOOGLE PROFESSIONAL MACHINE LEARNING ENGINEER CERTIFICATION QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

GCP-PMLE

Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE)

60 Questions Exam – 70% Cut Score – Duration of 120 minutes

www.VMExam.com

Table of Contents

Know Your GCP-PMLE Certification Well:

The GCP-PMLE is best suitable for candidates who want to gain knowledge in the Google Cloud. Before you start your GCP-PMLE preparation you may struggle to get all the crucial Professional Machine Learning Engineer materials like GCP-PMLE syllabus, sample questions, study guide.

But don't worry the GCP-PMLE PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the GCP-PMLE syllabus?
- How many questions are there in the GCP-PMLE exam?
- Which Practice test would help me to pass the GCP-PMLE exam at the first attempt?

Passing the GCP-PMLE exam makes you Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE). Having the Professional Machine Learning Engineer certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

Google GCP-PMLE Professional Machine Learning Engineer Certification Details:

Exam Name	Google Professional Machine Learning Engineer
Exam Code	GCP-PMLE
Exam Price	$200 USD
Duration	120 minutes
Number of Questions	60
Passing Score	Pass / Fail (Approx 70%)
Recommended Training / Books	**Google Cloud training** **Google Cloud documentation** **Google Cloud solutions**
Schedule Exam	**PEARSON VUE**
Sample Questions	**Google GCP-PMLE Sample Questions**
Recommended Practice	**Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE) Practice Test**

GCP-PMLE Syllabus:

Section	Objectives
Framing ML problems	
Translating business challenges into ML use cases. Considerations include:	- Choosing the best solution (ML vs. non-ML, custom vs. pre-packaged [e.g., AutoML, Vision API]) based on the business requirements - Defining how the model output should be used to solve the business problem - Deciding how incorrect results should be handled - Identifying data sources (available vs. ideal)
Defining ML problems. Considerations include:	- Problem type (e.g., classification, regression, clustering) - Outcome of model predictions - Input (features) and predicted output format
Defining business success criteria. Considerations include:	- Alignment of ML success metrics to the business problem - Key results - Determining when a model is deemed unsuccessful
Identifying risks to feasibility of ML solutions. Considerations include:	- Assessing and communicating business impact - Assessing ML solution readiness - Assessing data readiness and potential limitations - Aligning with Google's Responsible AI practices (e.g., different biases)
Architecting ML solutions	
Designing reliable, scalable, and highly available ML solutions. Considerations include:	- Choosing appropriate ML services for the use case (e.g., Cloud Build, Kubeflow) - Component types (e.g., data collection, data management) - Exploration/analysis - Feature engineering - Logging/management - Automation - Orchestration - Monitoring - Serving
Choosing appropriate Google Cloud hardware components. Considerations include:	- Evaluation of compute and accelerator options (e.g., CPU, GPU, TPU, edge devices)

Section	Objectives
Designing architecture that complies with security concerns across sectors/industries. Considerations include:	- Building secure ML systems (e.g., protecting against unintentional exploitation of data/model, hacking) - Privacy implications of data usage and/or collection (e.g., handling sensitive data such as Personally Identifiable Information [PII] and Protected Health Information [PHI])
Designing data preparation and processing systems	
Exploring data (EDA). Considerations include:	- Visualization - Statistical fundamentals at scale - Evaluation of data quality and feasibility - Establishing data constraints (e.g., TFDV)
Building data pipelines. Considerations include:	- Organizing and optimizing training datasets - Data validation - Handling missing data - Handling outliers - Data leakage
Creating input features (feature engineering). Considerations include:	- Ensuring consistent data pre-processing between training and serving - Encoding structured data types - Feature selection - Class imbalance - Feature crosses - Transformations (TensorFlow Transform)
Developing ML models	
Building models. Considerations include:	- Choice of framework and model - Modeling techniques given interpretability requirements - Transfer learning - Data augmentation - Semi-supervised learning - Model generalization and strategies to handle overfitting and underfitting
Training models. Considerations include:	- Ingestion of various file types into training (e.g., CSV, JSON, IMG, parquet or databases, Hadoop/Spark) - Training a model as a job in different environments - Hyperparameter tuning - Tracking metrics during training - Retraining/redeployment evaluation
Testing models. Considerations include:	- Unit tests for model training and serving - Model performance against baselines, simpler models, and

Section	Objectives
	across the time dimension - Model explainability on AI Platform
Scaling model training and serving. Considerations include:	- Distributed training - Scaling prediction service (e.g., AI Platform Prediction, containerized serving)
Automating and orchestrating ML pipelines	
Designing and implementing training pipelines. Considerations include:	- Identification of components, parameters, triggers, and compute needs (e.g., Cloud Build, Cloud Run) - Orchestration framework (e.g., Kubeflow Pipelines/AI Platform Pipelines, Cloud Composer/Apache Airflow) - Hybrid or multicloud strategies - System design with TFX components/Kubeflow DSL
Implementing serving pipelines. Considerations include:	- Serving (online, batch, caching) - Google Cloud serving options - Testing for target performance - Configuring trigger and pipeline schedules
Tracking and auditing metadata. Considerations include:	- Organizing and tracking experiments and pipeline runs - Hooking into model and dataset versioning - Model/dataset lineage
Monitoring, optimizing, and maintaining ML solutions	
Monitoring and troubleshooting ML solutions. Considerations include:	- Performance and business quality of ML model predictions - Logging strategies - Establishing continuous evaluation metrics (e.g., evaluation of drift or bias) - Understanding Google Cloud permissions model - Identification of appropriate retraining policy - Common training and serving errors (TensorFlow) - ML model failure and resulting biases
Tuning performance of ML solutions for training and serving in production. Considerations include:	- Optimization and simplification of input pipeline for training - Simplification techniques

Google GCP-PMLE Sample Questions:

Question: 1

You work on a team where the process for deploying a model into production starts with data scientists training different versions of models in a Kubeflow pipeline.

The workflow then stores the new model artifact into the corresponding Cloud Storage bucket. You need to build the next steps of the pipeline after the submitted model is ready to be tested and deployed in production on AI Platform.

How should you configure the architecture before deploying the model to production?

a) Deploy model in test environment -> Evaluate and test model -> Create a new AI Platform model version
b) Validate model -> Deploy model in test environment -> Create a new AI Platform model version
c) Create a new AI Platform model version -> Evaluate and test model -> Deploy model in test environment
d) Create a new AI Platform model version - > Deploy model in test environment -> Validate model

Answer: a

Question: 2

You need to write a generic test to verify whether Dense Neural Network (DNN) models automatically released by your team have a sufficient number of parameters to learn the task for which they were built.

What should you do?

a) Train the model for a few iterations, and check for NaN values.
b) Train the model for a few iterations, and verify that the loss is constant.
c) Train a simple linear model, and determine if the DNN model outperforms it.
d) Train the model with no regularization, and verify that the loss function is close to zero.

Answer: d

Question: 3

You are an ML engineer at a media company. You want to use machine learning to analyze video content, identify objects, and alert users if there is inappropriate content.

Which Google Cloud products should you use to build this project?

a) Pub/Sub, Cloud Function, Cloud Vision API
b) Pub/Sub, Cloud IoT, Dataflow, Cloud Vision API, Cloud Logging
c) Pub/Sub, Cloud Function, Video Intelligence API, Cloud Logging
d) Pub/Sub, Cloud Function, AutoML Video Intelligence, Cloud Logging

Answer: c

Question: 4

You work for a large retailer. You want to use ML to forecast future sales leveraging 10 years of historical sales data.

The historical data is stored in Cloud Storage in Avro format. You want to rapidly experiment with all the available data.

How should you build and train your model for the sales forecast?

a) Load data into BigQuery and use the ARIMA model type on BigQuery ML.
b) Convert the data into CSV format and create a regression model on AutoML Tables.
c) Convert the data into TFRecords and create an RNN model on TensorFlow on AI Platform Notebooks.
d) Convert and refactor the data into CSV format and use the built-in XGBoost algorithm on AI Platform Training.

Answer: a

Question: 5

You need to build an object detection model for a small startup company to identify if and where the company's logo appears in an image. You were given a large repository of images, some with logos and some without.

These images are not yet labelled. You need to label these pictures, and then train and deploy the model. What should you do?

a) Use Google Cloud's Data Labelling Service to label your data. Use AutoML Object Detection to train and deploy the model.
b) Use Vision API to detect and identify logos in pictures and use it as a label. Use AI Platform to build and train a convolutional neural network.
c) Create two folders: one where the logo appears and one where it doesn't. Manually place images in each folder. Use AI Platform to build and train a convolutional neural network.
d) Create two folders: one where the logo appears and one where it doesn't. Manually place images in each folder. Use AI Platform to build and train a real time object detection model.

Answer: a

Question: 6

Your team is using a TensorFlow Inception-v3 CNN model pretrained on ImageNet for an image classification prediction challenge on 10,000 images. You will use AI Platform to perform the model training.

What TensorFlow distribution strategy and AI Platform training job configuration should you use to train the model and optimize for wall-clock time?

a) Default Strategy; Custom tier with a single master node and four v100 GPUs.
b) One Device Strategy; Custom tier with a single master node and four v100 GPUs.
c) One Device Strategy; Custom tier with a single master node and eight v100 GPUs.
d) MirroredStrategy; Custom tier with a single master node and four v100 GPUs.

Answer: d

Question: 7

You work for a gaming company that develops and manages a popular massively multiplayer online (MMO) game.

The game's environment is open-ended, and a large number of positions and moves can be taken by a player. Your team has developed an ML model with TensorFlow that predicts the next move of each player.

Edge deployment is not possible, but low-latency serving is required. How should you configure the deployment?

a) Use a Cloud TPU to optimize model training speed.
b) Use AI Platform Prediction with a NVIDIA GPU to make real-time predictions.
c) Use AI Platform Prediction with a high-CPU machine type to get a batch prediction for the players.
d) Use AI Platform Prediction with a high-memory machine type to get a batch prediction for the players.

Answer: b

Question: 8

You work for a manufacturing company that owns a high-value machine which has several machine settings and multiple sensors.

A history of the machine's hourly sensor readings and known failure event data are stored in BigQuery. You need to predict if the machine will fail within the next 3 days in order to schedule maintenance before the machine fails.

Which data preparation and model training steps should you take?

a) Data preparation: Daily max value feature engineering; Model training: AutoML classification with BQML
b) Data preparation: Daily min value feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to True
c) Data preparation: Rolling average feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to False
d) Data preparation: Rolling average feature engineering; Model training: Logistic regression with BQML and AUTO_CLASS_WEIGHTS set to True

Answer: d

Question: 9

You work for a textile manufacturer and have been asked to build a model to detect and classify fabric defects.

You trained a machine learning model with high recall based on high resolution images taken at the end of the production line. You want quality control inspectors to gain trust in your model.

Which technique should you use to understand the rationale of your classifier?

a) Use K-fold cross validation to understand how the model performs on different test datasets.
b) Use the Integrated Gradients method to efficiently compute feature attributions for each predicted image.
c) Use PCA (Principal Component Analysis) to reduce the original feature set to a smaller set of easily understood features.
d) Use k-means clustering to group similar images together, and calculate the Davies-Bouldin index to evaluate the separation between clusters.

Answer: b

Question: 10

You work for a large financial institution that is planning to use Dialogflow to create a chatbot for the company's mobile app.

You have reviewed old chat logs and tagged each conversation for intent based on each customer's stated intention for contacting customer service.

About 70% of customer inquiries are simple requests that are solved within 10 intents. The remaining 30% of inquiries require much longer and more complicated requests.

Which intents should you automate first?

a) Automate a blend of the shortest and longest intents to be representative of all intents.
b) Automate the more complicated requests first because those require more of the agents' time.
c) Automate the 10 intents that cover 70% of the requests so that live agents can handle the more complicated requests.
d) Automate intents in places where common words such as "payment" only appear once to avoid confusing the software.

Answer: c

Study Guide to Crack Google Professional Machine Learning Engineer GCP-PMLE Exam:

- Getting details of the GCP-PMLE syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the GCP-PMLE exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the Google provided training for GCP-PMLE exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the GCP-PMLE sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on GCP-PMLE practice tests is must. Continuous practice will make you an expert in all syllabus areas.