

# Pass GCP-PMLE Professional Machine Learning Engineer Exam: Study Tips & Resources!

---

**GOOGLE PROFESSIONAL MACHINE LEARNING  
ENGINEER CERTIFICATION QUESTIONS & ANSWERS**

**Get Instant Access to Vital Exam Acing  
Materials | Study Guide | Sample Questions |  
Practice Test**

**GCP-PMLE**

**[Google Cloud Platform - Professional Machine Learning Engineer \(GCP-PMLE\)](#)**

**50-60 Questions Exam – 70% Cut Score – Duration of 120 minutes**

---

## Table of Contents

- Get Ready for the GCP-PMLE Exam:.....2**
- Know More About the Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE) Certification: .....2**
- Learn More About the GCP-PMLE Syllabus: .....2**
  - Architecting low-code ML solutions (12% of the exam) .....2
  - Collaborating within and across teams to manage data and models (16% of the exam) .....3
  - Scaling prototypes into ML models (18% of the exam).....4
  - Serving and scaling models (19% of the exam) .....5
  - Automating and orchestrating ML pipelines (21% of the exam) .....5
  - Monitoring ML solutions (14% of the exam) .....6
- Prepare with GCP-PMLE Sample Questions: .....7**
- Tips for Success in the Google Professional Machine Learning Engineer Exam:..... 11**
  - Familiarize Yourself with the GCP-PMLE Exam Format: ..... 11
  - Create A Study Timetable for the GCP-PMLE Exam: ..... 11
  - Diversify Your Study Sources: ..... 11
  - Regular Practice for the GCP-PMLE Exam:..... 11
  - Allow for Rest and Breaks:..... 11
  - Maintain Organization Throughout Your GCP-PMLE Exam Preparation:..... 11
  - Seek Guidance from Mentors: ..... 12
  - Regular Review is Crucial for the GCP-PMLE Exam:..... 12
  - Master Time Management for the GCP-PMLE Exam:..... 12
  - Have A Positive Mindset: ..... 12
- Benefits of Passing the GCP-PMLE Exam: ..... 12**
- Explore the Trusted Practice Exam for the GCP-PMLE Certification: ..... 13**
- Final Remarks: ..... 13**

## Get Ready for the GCP-PMLE Exam:

Prepare effectively for the GCP-PMLE exam using reliable [study strategies and methods](#). Enhance your preparedness, deepen your understanding of the Professional, and enhance your likelihood of achieving success in the Google Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE) with our comprehensive guide. Embark on your path to exam excellence today.

## Know More About the Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE) Certification:

<b>Exam Name</b>	Google Professional Machine Learning Engineer
<b>Exam Code</b>	GCP-PMLE
<b>Exam Price</b>	\$200 USD
<b>Duration</b>	120 minutes
<b>Number of Questions</b>	50-60
<b>Passing Score</b>	Pass / Fail (Approx 70%)
<b>Recommended Training / Books</b>	<a href="#">Google Cloud training</a> <a href="#">Google Cloud documentation</a> <a href="#">Google Cloud solutions</a>
<b>Schedule Exam</b>	<a href="#">PEARSON VUE</a>
<b>Sample Questions</b>	<a href="#">Google GCP-PMLE Sample Questions</a>
<b>Recommended Practice</b>	<a href="#">Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE) Practice Test</a>

## Learn More About the GCP-PMLE Syllabus:

Section	Objectives
<b><a href="#">Architecting low-code ML solutions (12% of the exam)</a></b>	
<b>Developing ML models by using</b>	- Building the appropriate BigQuery ML model (e.g., linear and binary classification, regression, time-series, matrix

Section	Objectives
<b>BigQuery ML. Considerations include:</b>	factorization, boosted trees, autoencoders) based on the business problem - Feature engineering or selection by using BigQuery ML - Generating predictions by using BigQuery ML
<b>Building AI solutions by using ML APIs. Considerations include:</b>	- Building applications by using ML APIs (e.g., Cloud Vision API, Natural Language API, Cloud Speech API, Translation) - Building applications by using industry-specific APIs (e.g., Document AI API, Retail API)
<b>Training models by using AutoML. Considerations include:</b>	- Preparing data for AutoML (e.g., feature selection, data labeling, Tabular Workflows on AutoML) - Using available data (e.g., tabular, text, speech, images, videos) to train custom models - Using AutoML for tabular data - Creating forecasting models using AutoML - Configuring and debugging trained models
<b>Collaborating within and across teams to manage data and models (16% of the exam)</b>	
<b>Exploring and preprocessing organization-wide data (e.g., Cloud Storage, BigQuery, Cloud Spanner, Cloud SQL, Apache Spark, Apache Hadoop). Considerations include:</b>	- Organizing different types of data (e.g., tabular, text, speech, images, videos) for efficient training - Managing datasets in Vertex AI - Data preprocessing (e.g., Dataflow, TensorFlow Extended [TFX], BigQuery) - Creating and consolidating features in Vertex AI Feature Store - Privacy implications of data usage and/or collection (e.g., handling sensitive data such as personally identifiable information [PII] and protected health information [PHI])
<b>Model prototyping</b>	- Choosing the appropriate Jupyter backend on Google Cloud

Section	Objectives
<b>using Jupyter notebooks.</b> <b>Considerations include:</b>	(e.g., Vertex AI Workbench, notebooks on Dataproc) - Applying security best practices in Vertex AI Workbench - Using Spark kernels - Integration with code source repositories - Developing models in Vertex AI Workbench by using common frameworks (e.g., TensorFlow, PyTorch, sklearn, Spark, JAX)
<b>Tracking and running ML experiments.</b> <b>Considerations include:</b>	- Choosing the appropriate Google Cloud environment for development and experimentation (e.g., Vertex AI Experiments, Kubeflow Pipelines, Vertex AI TensorBoard with TensorFlow and PyTorch) given the framework
<b>Scaling prototypes into ML models (18% of the exam)</b>	
<b>Building models.</b> <b>Considerations include:</b>	- Choosing ML framework and model architecture - Modeling techniques given interpretability requirements
<b>Training models.</b> <b>Considerations include:</b>	- Organizing training data (e.g., tabular, text, speech, images, videos) on Google Cloud (e.g., Cloud Storage, BigQuery) - Ingestion of various file types (e.g., CSV, JSON, images, Hadoop, databases) into training - Training using different SDKs (e.g., Vertex AI custom training, Kubeflow on Google Kubernetes Engine, AutoML, tabular workflows) - Using distributed training to organize reliable pipelines - Hyperparameter tuning - Troubleshooting ML model training failures
<b>Choosing appropriate hardware for training.</b> <b>Considerations</b>	- Evaluation of compute and accelerator options (e.g., CPU, GPU, TPU, edge devices) - Distributed training with TPUs and GPUs (e.g., Reduction Server on Vertex AI, Horovod)

Section	Objectives
<b>include:</b>	
<b>Serving and scaling models (19% of the exam)</b>	
<b>Serving models. Considerations include:</b>	<ul style="list-style-type: none"> <li>- Batch and online inference (e.g., Vertex AI, Dataflow, BigQuery ML, Dataproc)</li> <li>- Using different frameworks (e.g., PyTorch, XGBoost) to serve models</li> <li>- Organizing a model registry</li> <li>- A/B testing different versions of a model</li> </ul>
<b>Scaling online model serving. Considerations include:</b>	<ul style="list-style-type: none"> <li>- Vertex AI Feature Store</li> <li>- Vertex AI public and private endpoints</li> <li>- Choosing appropriate hardware (e.g., CPU, GPU, TPU, edge)</li> <li>- Scaling the serving backend based on the throughput (e.g., Vertex AI Prediction, containerized serving)</li> <li>- Tuning ML models for training and serving in production (e.g., simplification techniques, optimizing the ML solution for increased performance, latency, memory, throughput)</li> </ul>
<b>Automating and orchestrating ML pipelines (21% of the exam)</b>	
<b>Developing end-to-end ML pipelines. Considerations include:</b>	<ul style="list-style-type: none"> <li>- Data and model validation</li> <li>- Ensuring consistent data pre-processing between training and serving</li> <li>- Hosting third-party pipelines on Google Cloud (e.g., MLFlow)</li> <li>- Identifying components, parameters, triggers, and compute needs (e.g., Cloud Build, Cloud Run)</li> <li>- Orchestration framework (e.g., Kubeflow Pipelines, Vertex AI Pipelines, Cloud Composer)</li> <li>- Hybrid or multicloud strategies</li> <li>- System design with TFX components or Kubeflow DSL</li> </ul>

Section	Objectives
	(e.g., Dataflow)
<b>Automating model retraining.</b> <b>Considerations include:</b>	<ul style="list-style-type: none"> <li>- Determining an appropriate retraining policy</li> <li>- Continuous integration and continuous delivery (CI/CD) model deployment (e.g., Cloud Build, Jenkins)</li> </ul>
<b>Tracking and auditing metadata.</b> <b>Considerations include:</b>	<ul style="list-style-type: none"> <li>- Tracking and comparing model artifacts and versions (e.g., Vertex AI Experiments, Vertex ML Metadata)</li> <li>- Hooking into model and dataset versioning</li> <li>- Model and data lineage</li> </ul>
<b>Monitoring ML solutions (14% of the exam)</b>	
<b>Identifying risks to ML solutions.</b> <b>Considerations include:</b>	<ul style="list-style-type: none"> <li>- Building secure ML systems (e.g., protecting against unintentional exploitation of data or models, hacking)</li> <li>- Aligning with Google’s Responsible AI practices (e.g., biases)</li> <li>- Assessing ML solution readiness (e.g., data bias, fairness)</li> <li>- Model explainability on Vertex AI (e.g., Vertex AI Prediction)</li> </ul>
<b>Monitoring, testing, and troubleshooting ML solutions.</b> <b>Considerations include:</b>	<ul style="list-style-type: none"> <li>- Establishing continuous evaluation metrics (e.g., Vertex AI Model Monitoring, Explainable AI)</li> <li>- Monitoring for training-serving skew</li> <li>- Monitoring for feature attribution drift</li> <li>- Monitoring model performance against baselines, simpler models, and across the time dimension</li> <li>- Common training and serving errors</li> </ul>

## Prepare with GCP-PMLE Sample Questions:

### Question: 1

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**

### Question: 2

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) 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.
- 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) Use Google Cloud's Data Labelling Service to label your data. Use AutoML Object Detection to train and deploy the model.

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

**Answer: b**

**Question: 5**

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: 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 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: 8**

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: 9**

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: 10**

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 the Integrated Gradients method to efficiently compute feature attributions for each predicted image.
- b) Use K-fold cross validation to understand how the model performs on different test datasets.
- 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: a**

# Tips for Success in the Google Professional Machine Learning Engineer Exam:

## Familiarize Yourself with the GCP-PMLE Exam Format:

Before starting your study regimen, it's crucial to acquaint yourself with the structure of the GCP-PMLE exam. Take a moment to review [the exam syllabus](#), grasp the test format, and pinpoint the main areas of concentration. Having prior knowledge of the exam's layout will assist you in customizing your study strategy effectively.

## Create A Study Timetable for the GCP-PMLE Exam:

To prepare efficiently for the GCP-PMLE exam, devise a study schedule that aligns with your lifestyle and preferred learning approach. Allocate dedicated time slots for studying each day, prioritizing topics according to their significance and your level of proficiency. Maintaining consistency by adhering to your schedule and steering clear of procrastination is imperative.

## Diversify Your Study Sources:

Ensure you broaden your study material beyond just one source. Use various resources like textbooks, online courses, practice exams, and study guides to understand the GCP-PMLE exam subjects thoroughly. Each resource provides distinct perspectives and explanations that can enrich your learning journey.

## Regular Practice for the GCP-PMLE Exam:

Consistent practice is essential for effective preparation for the GCP-PMLE exam. Engaging in regular practice enables you to strengthen your grasp of essential concepts, improve your problem-solving abilities, and become accustomed to the exam format. Allocate dedicated time to solving practice questions and sample tests to assess your progress accurately.

## Allow for Rest and Breaks:

While studying is crucial, taking breaks and rest is equally vital. Pushing yourself too hard without sufficient rest can result in burnout and reduced effectiveness. Incorporate short breaks into your study sessions to recharge and stay focused.

## Maintain Organization Throughout Your GCP-PMLE Exam Preparation:

Keep yourself organized as you prepare for the GCP-PMLE exam by monitoring your progress and managing your materials effectively. Ensure your study area remains

neat, utilize folders or digital aids to arrange your notes and resources, and develop a checklist of topics to review. Employing an organized approach will assist you in staying focused and reducing stress levels.

## Seek Guidance from Mentors:

Feel free to ask for clarification when you come across confusing or difficult concepts during your study sessions. Seek support from peers, instructors, or online forums to address any uncertainties. Addressing doubts will prevent misunderstandings and ensure you develop a strong [understanding of the material](#).

## Regular Review is Crucial for the GCP-PMLE Exam:

Frequent revisiting of material is paramount for retaining information over the long term. Revisit topics you've already covered to strengthen your comprehension and pinpoint areas that need further focus. Regular review sessions will [solidify your understanding](#) and enhance your confidence.

## Master Time Management for the GCP-PMLE Exam:

Skillful time management is essential on the exam day to ensure you finish all sections within the designated time limits. During your practice sessions, replicate the conditions of the GCP-PMLE exam and practice managing your time accordingly. Formulate strategies for efficiently addressing each section to optimize your score.

## Have A Positive Mindset:

Finally, maintain a positive attitude and have faith in your capabilities. Stay confident in your preparation and trust that you are well-prepared to handle the GCP-PMLE exam. Envision success, remain focused, and approach the exam calmly and objectively.

## Benefits of Passing the GCP-PMLE Exam:

- Completing the GCP-PMLE exam unlocks pathways to fresh career prospects and progression within your industry.
- The extensive preparation needed for the GCP-PMLE certification equips you with comprehensive knowledge and practical expertise applicable to your field.
- Possessing the GCP-PMLE certification showcases your mastery and dedication to excellence, garnering acknowledgment from both peers and employers.
- Certified professionals often command higher salaries and have greater potential for earning than those without certification.

- Acquiring the GCP-PMLE certification validates your competence and trustworthiness, fostering confidence among clients, employers, and peers.

## Explore the Trusted Practice Exam for the GCP-PMLE Certification:

At VMExam.com, you'll find comprehensive resources for the GCP-PMLE exam. Our platform offers authentic practice exams tailored specifically for the GCP-PMLE certification. What advantages do these practice exams provide? You'll encounter genuine exam-style questions expertly crafted by industry professionals, allowing you to improve your performance in the exam. Rely on VMExam.com for rigorous, unlimited access to [GCP-PMLE practice exams](#) for two months, allowing you to boost your confidence steadily. Through focused practice, numerous candidates have successfully streamlined their path to achieving the Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE).

### Final Remarks:

Preparing for the GCP-PMLE examination demands commitment, strategic planning, and efficient study methods. Implementing these study suggestions can enrich your preparation, elevate your self-assurance, and increase your likelihood of excelling in the exam. Keep your focus sharp, maintain organization, and believe in your abilities. Best of luck!

### Here Is the Trusted Practice Test for the GCP-PMLE Certification

VMExam.Com is here with all the necessary details regarding the GCP-PMLE exam. We provide authentic practice tests for the GCP-PMLE exam. What do you gain from these practice tests? You get to experience the real exam-like questions made by industry experts and get a scope to improve your performance in the actual exam. Rely on VMExam.Com for rigorous, unlimited two-month attempts on the [GCP-PMLE practice tests](#), and gradually build your confidence. Rigorous practice made many aspirants successful and made their journey easy towards grabbing the Google Cloud Platform - Professional Machine Learning Engineer (GCP-PMLE).

**Start Online Practice of GCP-PMLE Exam by Visiting URL**

<https://www.vmexam.com/google/gcp-pmle-google-professional-machine-learning-engineer>