

Google GCP-PCD

**Google Professional Cloud Developer
Certification Questions & Answers**

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

GCP-PCD

[Google Cloud Platform - Professional Cloud Developer \(GCP-PCD\)](#)

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



Table of Contents:

Discover More about the GCP-PCD Certification	2
Google GCP-PCD Professional Cloud Developer Certification Details:	2
GCP-PCD Syllabus:	2
Designing highly scalable, available, and reliable cloud-native applications	2
Building and Testing Applications	4
Deploying applications	4
Integrating Google Cloud Services	5
Managing Application Performance Monitoring	6
Broaden Your Knowledge with Google GCP-PCD Sample Questions:	7
Avail the Study Guide to Pass Google GCP-PCD Professional Cloud Developer Exam:	10
Career Benefits:	11

Discover More about the GCP-PCD Certification

Are you interested in passing the Google GCP-PCD exam? First discover, who benefits from the GCP-PCD certification. The GCP-PCD is suitable for a candidate if he wants to learn about Cloud. Passing the GCP-PCD exam earns you the Google Cloud Platform - Professional Cloud Developer (GCP-PCD) title.

While preparing for the GCP-PCD exam, many candidates struggle to get the necessary materials. But do not worry; your struggling days are over. The GCP-PCD PDF contains some of the most valuable preparation tips and the details and instant access to useful [GCP-PCD study materials just at one click](#).

Google GCP-PCD Professional Cloud Developer Certification Details:

Exam Name	Google Professional Cloud Developer
Exam Code	GCP-PCD
Exam Price	\$200 USD
Duration	120 minutes
Number of Questions	50
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-PCD Sample Questions
Recommended Practice	Google Cloud Platform - Professional Cloud Developer (GCP-PCD) Practice Test

GCP-PCD Syllabus:

Section	Objectives
Designing highly scalable, available, and reliable cloud-native applications	
Designing high-performing applications and APIs.	<ul style="list-style-type: none"> - Microservices - Scaling velocity characteristics/trade-offs of IaaS (infrastructure as a service) vs. CaaS (container as a service) vs. PaaS (platform as a service)

Section	Objectives
Considerations include:	<ul style="list-style-type: none"> - Geographic distribution of Google Cloud services (e.g., latency, regional services, zonal services) - Defining a key structure for high-write applications using Cloud Storage, Cloud Bigtable, Cloud Spanner, or Cloud SQL - User session management - Caching solutions - Deploying and securing API services - Loosely coupled asynchronous applications (e.g., Apache Kafka, Pub/Sub) - Graceful shutdown on platform termination - Google-recommended practices and documentation
Designing secure applications. Considerations include:	<ul style="list-style-type: none"> - Implementing requirements that are relevant for applicable regulations (e.g., data wipeout) - Security mechanisms that protect services and resources - Security mechanisms that secure/scan application binaries and manifests - Storing and rotating application secrets and keys (e.g., Cloud KMS, HashiCorp Vault) - Authenticating to Google services (e.g., application default credentials, JSON Web Token (JWT), OAuth 2.0) - IAM roles for users/groups/service accounts - Securing service-to-service communications (e.g., service mesh, Kubernetes network policies, and Kubernetes namespaces) - Running services with least privileged access (e.g., Workload Identity) - Certificate-based authentication (e.g., SSL, mTLS) - Google-recommended practices and documentation
Managing application data. Considerations include:	<ul style="list-style-type: none"> - Defining database schemas for Google-managed databases (e.g., Firestore, Cloud Spanner, Cloud Bigtable, Cloud SQL) - Choosing data storage options based on use case considerations, such as: <ul style="list-style-type: none"> • Time-limited access to objects • Data retention requirements • Structured vs. unstructured data • Strong vs. eventual consistency • Data volume • Frequency of data access in Cloud Storage - Google-recommended practices and documentation

Section	Objectives
Application modernization. Considerations include:	<ul style="list-style-type: none"> - Using managed services - Refactoring a monolith to microservices - Designing stateless, horizontally scalable services - Google-recommended practices and documentation
<h3 style="color: #0070C0;">Building and Testing Applications</h3>	
Setting up your local development environment. Considerations include:	<ul style="list-style-type: none"> - Emulating Google Cloud services for local application development - Creating Google Cloud projects - Using the command-line interface (CLI), Google Cloud Console, and Cloud Shell tools - Using developer tooling (e.g., Cloud Code, Skaffold)
Writing efficient code. Considerations include:	<ul style="list-style-type: none"> - Algorithm design - Modern application patterns - Software development methodologies - Debugging and profiling code
Testing. Considerations include:	<ul style="list-style-type: none"> - Unit testing - Integration testing - Performance testing - Load testing
Building. Considerations include:	<ul style="list-style-type: none"> - Source control management - Creating secure container images from code - Developing a continuous integration pipeline using services (e.g., Cloud Build, Container Registry) that construct deployment artifacts - Reviewing and improving continuous integration pipeline efficiency
<h3 style="color: #0070C0;">Deploying applications</h3>	
Recommend appropriate deployment strategies using the appropriate tools (e.g., Cloud Build, Spinnaker, Tekton, Anthos Configuration Manager) for the target compute environment (e.g., Compute Engine, Google Kubernetes Engine).	<ul style="list-style-type: none"> - Blue/green deployments - Traffic-splitting deployments - Rolling deployments - Canary deployments

Section	Objectives
Considerations include:	
Deploying applications and services on Compute Engine. Considerations include:	<ul style="list-style-type: none"> - Installing an application into a virtual machine (VM) - Managing service accounts for VMs - Bootstrapping applications - Exporting application logs and metrics - Managing Compute Engine VM images and binaries
Deploying applications and services to Google Kubernetes Engine (GKE). Considerations include:	<ul style="list-style-type: none"> - Deploying a containerized application to GKE - Managing Kubernetes RBAC and Google Cloud IAM relationships - Configuring Kubernetes namespaces - Defining workload specifications (e.g., resource requirements) - Building a container image using Cloud Build - Configuring application accessibility to user traffic and other services - Managing container life cycle - Define Kubernetes resources and configurations
Deploying a Cloud Function. Considerations include:	<ul style="list-style-type: none"> - Cloud Functions that are triggered via an event from Google Cloud services (e.g., Pub/Sub, Cloud Storage objects) - Cloud Functions that are invoked via HTTP - Securing Cloud Functions
Using service accounts. Considerations include:	<ul style="list-style-type: none"> - Creating a service account according to the principle of least privilege - Downloading and using a service account private key file
Integrating Google Cloud Services	
Integrating an application with Data and Storage services. Considerations include:	<ul style="list-style-type: none"> - Read/write data to/from various databases (e.g., SQL) - Connecting to a data store (e.g., Cloud SQL, Cloud Spanner, Firestore, Cloud Bigtable) - Writing an application that publishes/consumes data asynchronously (e.g., from Pub/Sub) - Storing and retrieving objects from Cloud Storage
Integrating an application with compute services. Considerations include:	<ul style="list-style-type: none"> - Implementing service discovery in GKE and Compute Engine - Reading instance metadata to obtain application configuration - Authenticating users by using OAuth2.0 Web Flow and Identity-Aware Proxy - Authenticating to Cloud APIs with Workload Identity

Section	Objectives
Integrating Cloud APIs with applications. Considerations include:	<ul style="list-style-type: none"> - Enabling a Cloud API - Making API calls using supported options (e.g., Cloud Client Library, REST API or gRPC, APIs Explorer) taking into consideration: <ul style="list-style-type: none"> • Batching requests • Restricting return data • Paginating results • Caching results • Error handling (e.g., exponential backoff) - Using service accounts to make Cloud API calls
Managing Application Performance Monitoring	
Managing Compute Engine VMs. Considerations include:	<ul style="list-style-type: none"> - Debugging a custom VM image using the serial port - Diagnosing a failed Compute Engine VM startup - Sending logs from a VM to Cloud Logging - Viewing and analyzing logs - Inspecting resource utilization over time
Managing Google Kubernetes Engine workloads. Considerations include:	<ul style="list-style-type: none"> - Configuring logging and monitoring - Analyzing container life cycle events (e.g., CrashLoopBackOff, ImagePullErr) - Viewing and Analyzing logs - Writing and exporting custom metrics - Using external metrics and corresponding alerts - Configuring workload autoscaling
Troubleshooting application performance. Considerations include:	<ul style="list-style-type: none"> - Creating a monitoring dashboard - Writing custom metrics and creating log-based metrics - Using Cloud Debugger - Reviewing stack traces for error analysis - Exporting logs from Google Cloud - Viewing logs in the Google Cloud Console - Reviewing application performance (e.g., Cloud Trace, Prometheus, OpenTelemetry) - Monitoring and profiling a running application - Using documentation, forums, and Google Cloud support

Broaden Your Knowledge with Google GCP-PCD Sample Questions:

Question: 1

Your company has a successful multi-player game that has become popular in the US. Now, it wants to expand to other regions. It is launching a new feature that allows users to trade points. This feature will work for users across the globe.

Your company's current MySQL backend is reaching the limit of the Compute Engine instance that hosts the game. Your company wants to migrate to a different database that will provide global consistency and high availability across the regions.

Which database should they choose?

- a) BigQuery
- b) Cloud Spanner
- c) Cloud SQL
- d) Cloud Bigtable

Answer: b

Question: 2

You are building a storage layer for an analytics Hadoop cluster for your company. This cluster will run multiple jobs on a nightly basis, and you need to access the data frequently.

You want to use Cloud Storage for this purpose. Which storage option should you choose?

- a) Multi-regional storage
- b) Regional storage
- c) Nearline storage
- d) Coldline storage

Answer: b

Question: 3

As part of their expansion, HipLocal is creating new projects in order to separate resources. They want to build a system to automate enabling of their APIs. What should they do?

- a) Copy existing persistent disks to the new project.
- b) Use the service management API to define a new service.
- c) Use the service management API to enable the Compute API.
- d) Use the service management API to enable the Cloud Storage API.

Answer: c

Question: 4

You are capturing important audit activity in Stackdriver Logging. You need to read the information from Stackdriver Logging to perform real-time analysis of the logs.

You will have multiple processes performing different types of analysis on the logging data. What should you do?

- a) Read the logs directly from the Stackdriver Logging API.
- b) Set up a Stackdriver Logging sync to BigQuery, and read the logs from the BigQuery table.
- c) Set up a Stackdriver Logging sync to Cloud Pub/Sub, and read the logs from a Cloud Pub/Sub topic.
- d) Set up a Stackdriver Logging sync to Cloud Storage, and read the logs from a Cloud Storage bucket.

Answer: c

Question: 5

Your organization has grown, and new teams need access to manage network connectivity within and across projects. You are now seeing intermittent timeout errors in your application.

You want to find the cause of the problem. What should you do?

- a) Set up wireshark on each Google Cloud Virtual Machine instance.
- b) Configure VPC flow logs for each of the subnets in your VPC.
- c) Review the instance admin activity logs in Stackdriver for the application instances.
- d) Configure firewall rules logging for each of the firewalls in your VPC.

Answer: b

Question: 6

You have an application that accepts inputs from users. The application needs to kick off different background tasks based on these inputs.

You want to allow for automated asynchronous execution of these tasks as soon as input is submitted by the user.

Which product should you use?

- a) Cloud Tasks
- b) Cloud Bigtable
- c) Cloud Pub/Sub
- d) Cloud Composer

Answer: a

Question: 7

Which architecture should HipLocal use for log analysis?

- a) Use Cloud Spanner to store each event.
- b) Start storing key metrics in Cloud Memorystore.
- c) Use Stackdriver Logging with a BigQuery sink.
- d) Use Stackdriver Logging with a Cloud Storage sink.

Answer: c

Question: 8

You have a service running on Compute Engine virtual machine instances behind a global load balancer. You need to ensure that when the instance fails, it is recovered. What should you do?

- a) Set up health checks in the load balancer configuration.
- b) Deploy a service to the instances to notify you when they fail.
- c) Use Stackdriver alerting to trigger a workflow to reboot the instance.
- d) Set up health checks in the managed instance group configuration.

Answer: d

Question: 9

Your company plans to expand their analytics use cases. One of the new use cases requires your data analysts to analyze events using SQL on a near real-time basis.

You expect rapid growth and want to use managed services as much as possible. What should you do?

- a) Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into BigQuery.
- b) Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into Cloud Storage.
- c) Create a Kafka instance on a large Compute Engine instance. Stream your events from the source into a Kafka pipeline. Leverage Cloud Dataflow to ingest these events into Cloud Storage.
- d) Create a Cloud Pub/Sub topic and a subscription. Stream your events from the source into the Pub/Sub topic. Leverage Cloud Dataflow to ingest these events into Cloud Datastore.

Answer: a

Question: 10

Your application starts on the VM as a systemd service. Your application outputs its log information to stdout.

You need to send the application logs to Stackdriver without changing the application. What should you do?

- a) Review the application logs from the Compute Engine VM Instance activity logs in Stackdriver.
- b) Review the application logs from the Compute Engine VM Instance data access logs in Stackdriver.
- c) Install Stackdriver Logging Agent. Review the application logs from the Compute Engine VM Instance syslog logs in Stackdriver.
- d) Install Stackdriver Logging Agent. Review the application logs from the Compute Engine VM Instance system event logs in Stackdriver.

Answer: c

Avail the Study Guide to Pass Google GCP-PCD Professional Cloud Developer Exam:

- Find out about the GCP-PCD syllabus topics. Visiting the official site offers an idea about the exam structure and other important study resources. Going through the syllabus topics help to plan the exam in an organized manner.
- Once you are done exploring the [GCP-PCD syllabus](#), it is time to plan for studying and covering the syllabus topics from the core. Chalk out the best plan for yourself to cover each part of the syllabus in a hassle-free manner.
- A study schedule helps you to stay calm throughout your exam preparation. It should contain your materials and thoughts like study hours, number of topics for daily studying mentioned on it. The best bet to clear the exam is to follow your schedule rigorously.
- The candidate should not miss out on the scope to learn from the GCP-PCD training. Joining the Google provided training for GCP-PCD exam helps a candidate to strengthen his practical knowledge base from the certification.
- Learning about the probable questions and gaining knowledge regarding the exam structure helps a lot. Go through the [GCP-PCD sample questions](#) and boost your knowledge

- Make yourself a pro through online practicing the syllabus topics. GCP-PCD practice tests would guide you on your strengths and weaknesses regarding the syllabus topics. Through rigorous practicing, you can improve the weaker sections too. Learn well about time management during exam and become confident gradually with practice tests.

Career Benefits:

- Passing the GCP-PCD exam, helps a candidate to prosper highly in his career. Having the certification on the resume adds to the candidate's benefit and helps to get the best opportunities.

Here Is the Trusted Practice Test for the GCP-PCD Certification

VMExam.Com is here with all the necessary details regarding the GCP-PCD exam. We provide authentic practice tests for the GCP-PCD 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-PCD 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 Cloud Developer (GCP-PCD).

Start Online practice of GCP-PCD Exam by visiting URL

<https://www.vmexam.com/google/gcp-pcd-google-professional-cloud-developer>