

PROFESSIONAL-CLOUD- DEVELOPER^{Q&As}

Professional Cloud Developer

**Pass Google PROFESSIONAL-CLOUD-DEVELOPER
Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/professional-cloud-developer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

Your company needs a database solution that stores customer purchase history and meets the following requirements:

Customers can query their purchase immediately after submission.

Purchases can be sorted on a variety of fields.

Distinct record formats can be stored at the same time.

Which storage option satisfies these requirements?

- A. Firestore in Native mode
- B. Cloud Storage using an object read
- C. Cloud SQL using a SQL SELECT statement
- D. Firestore in Datastore mode using a global query

Correct Answer: A

QUESTION 2

For this question, refer to the HipLocal case study.

HipLocal's application uses Cloud Client Libraries to interact with Google Cloud. HipLocal needs to configure authentication and authorization in the Cloud Client Libraries to implement least privileged access for the application. What should they do?

- A. Create an API key. Use the API key to interact with Google Cloud.
- B. Use the default compute service account to interact with Google Cloud.
- C. Create a service account for the application. Export and deploy the private key for the application. Use the service account to interact with Google Cloud.
- D. Create a service account for the application and for each Google Cloud API used by the application. Export and deploy the private keys used by the application. Use the service account with one Google Cloud API to interact with Google Cloud.

Correct Answer: A

QUESTION 3

You are monitoring a web application that is written in Go and deployed in Google Kubernetes Engine. You notice an increase in CPU and memory utilization. You need to determine which source code is consuming the most CPU and memory resources. What should you do?

- A. Download, install, and start the Snapshot Debugger agent in your VM. Take debug snapshots of the functions that take the longest time. Review the call stack frame, and identify the local variables at that level in the stack.

B. Import the Cloud Profiler package into your application, and initialize the Profiler agent. Review the generated flame graph in the Google Cloud console to identify time-intensive functions.

C. Import OpenTelemetry and Trace export packages into your application, and create the trace provider. Review the latency data for your application on the Trace overview page, and identify where bottlenecks are occurring.

D. Create a Cloud Logging query that gathers the web application's logs. Write a Python script that calculates the difference between the timestamps from the beginning and the end of the application's longest functions to identify time-intensive functions.

Correct Answer: B

QUESTION 4

Your data is stored in Cloud Storage buckets. Fellow developers have reported that data downloaded from Cloud Storage is resulting in slow API performance. You want to research the issue to provide details to the GCP support team. Which command should you run?

A. `gsutil test -o output.json gs://my-bucket`

B. `gsutil perfdiag -o output.json gs://my-bucket`

C. `gcloud compute scp example-instance:~/test-data -o output.json gs://my-bucket`

D. `gcloud services test -o output.json gs://my-bucket`

Correct Answer: B

Reference: <https://groups.google.com/forum/#!topic/gce-discussion/xBI9Jq5HDsY>

QUESTION 5

You are configuring a continuous integration pipeline using Cloud Build to automate the deployment of new container images to Google Kubernetes Engine (GKE). The pipeline builds the application from its source code, runs unit and integration tests in separate steps, and pushes the container to Container Registry. The application runs on a Python web server. The Dockerfile is as follows: FROM python:3.7-alpine COPY . /app WORKDIR /app RUN pip install -r requirements.txt CMD ["unicorn", "-w 4", "main:app"] You notice that Cloud Build runs are taking longer than expected to complete. You want to decrease the build time. What should you do? (Choose two.)

A. Select a virtual machine (VM) size with higher CPU for Cloud Build runs.

B. Deploy a Container Registry on a Compute Engine VM in a VPC, and use it to store the final images.

C. Cache the Docker image for subsequent builds using the `--cache-from` argument in your build config file.

D. Change the base image in the Dockerfile to `ubuntu:latest`, and install Python 3.7 using a package manager utility.

E. Store application source code on Cloud Storage, and configure the pipeline to use `gsutil` to download the source code.

Correct Answer: AC

<https://cloud.google.com/build/docs/optimize-builds/increase-vcpu-for-builds>

By default, Cloud Build runs your builds on a standard virtual machine (VM). In addition to the standard VM, Cloud Build provides several high-CPU VM types to run builds. To increase the speed of your build, select a machine with a higher

vCPU to run builds. Keep in mind that although selecting a high vCPU machine increases your build speed, it may also increase the startup time of your build as Cloud Build only starts non-standard machines on demand.

https://cloud.google.com/build/docs/optimize-builds/speeding-up-builds#using_a_cached_docker_image

The easiest way to increase the speed of your Docker image build is by specifying a cached image that can be used for subsequent builds. You can specify the cached image by adding the `--cache-from` argument in your build config file,

which will instruct Docker to build using that image as a cache source.

[Latest PROFESSIONAL-CLOUD-DEVELOPER Dumps](#)

[PROFESSIONAL-CLOUD-DEVELOPER PDF Dumps](#)

[PROFESSIONAL-CLOUD-DEVELOPER Study Guide](#)