



# PROFESSIONAL-CLOUD-DEVELOPER<sup>Q&As</sup>

Professional Cloud Developer

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

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

<https://www.pass4itsure.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



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-developer.html>  
2024 Latest pass4itsure PROFESSIONAL-CLOUD-DEVELOPER PDF and  
VCE dumps Download

---

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





### QUESTION 1

You have an application deployed in Google Kubernetes Engine (GKE). You need to update the application to make authorized requests to Google Cloud managed services. You want this to be a one-time setup, and you need to follow security best practices of auto-rotating your security keys and storing them in an encrypted store. You already created a service account with appropriate access to the Google Cloud service. What should you do next?

- A. Assign the Google Cloud service account to your GKE Pod using Workload Identity.
- B. Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.
- C. Export the Google Cloud service account, and embed it in the source code of the application.
- D. Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Correct Answer: A

<https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity>

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs.

Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

---

### QUESTION 2

You are deploying your application on a Compute Engine instance that communicates with Cloud SQL. You will use Cloud SQL Proxy to allow your application to communicate to the database using the service account associated with the application's instance. You want to follow the Google-recommended best practice of providing minimum access for the role assigned to the service account. What should you do?

- A. Assign the Project Editor role.
- B. Assign the Project Owner role.
- C. Assign the Cloud SQL Client role.
- D. Assign the Cloud SQL Editor role.

Correct Answer: C

Reference: <https://cloud.google.com/sql/docs/mysql/sql-proxy>

---

### QUESTION 3

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you



build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.

Correct Answer: AC

Explanation: <https://cloud.google.com/run/docs/deploying#images> <https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks>

---

#### QUESTION 4

Your team develops services that run on Google Kubernetes Engine. You need to standardize their log data using Google-recommended practices and make the data more useful in the fewest number of steps. What should you do? (Choose two.)

- A. Create aggregated exports on application logs to BigQuery to facilitate log analytics.
- B. Create aggregated exports on application logs to Cloud Storage to facilitate log analytics.
- C. Write log output to standard output (stdout) as single-line JSON to be ingested into Cloud Logging as structured logs.
- D. Mandate the use of the Logging API in the application code to write structured logs to Cloud Logging.
- E. Mandate the use of the Pub/Sub API to write structured data to Pub/Sub and create a Dataflow streaming pipeline to normalize logs and write them to BigQuery for analytics.

Correct Answer: AC

Explanation: [https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best\\_practices](https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices)

---

#### QUESTION 5

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

[PROFESSIONAL-CLOUD-DEVELOPER Practice Test](#)

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

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)