



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.passapply.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

PassApply.com

<https://www.passapply.com/professional-cloud-developer.html>
2024 Latest passapply 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

For this question, refer to the HipLocal case study.

Which Google Cloud product addresses HipLocal's business requirements for service level indicators and objectives?

- A. Cloud Profiler
- B. Cloud Monitoring
- C. Cloud Trace
- D. Cloud Logging

Correct Answer: B

<https://cloud.google.com/stackdriver/docs/solutions/slo-monitoring#defn-sli>

QUESTION 2

You are developing a microservice-based application that will be deployed on a Google Kubernetes Engine cluster. The application needs to read and write to a Spanner database. You want to follow security best practices while minimizing code changes.

How should you configure your application to retrieve Spanner credentials?

- A. Configure the appropriate service accounts, and use Workload Identity to run the pods.
- B. Store the application credentials as Kubernetes Secrets, and expose them as environment variables.
- C. Configure the appropriate routing rules, and use a VPC-native cluster to directly connect to the database.
- D. Store the application credentials using Cloud Key Management Service, and retrieve them whenever a database connection is made.

Correct Answer: A

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

QUESTION 3

You are building a highly available and globally accessible application that will serve static content to users. You need to configure the storage and serving components. You want to minimize management overhead and latency while maximizing reliability for users. What should you do?

- A. 1) Create a managed instance group. Replicate the static content across the virtual machines (VMs) 2) Create an external HTTP(S) load balancer. 3) Enable Cloud CDN, and send traffic to the managed instance group.
- B. 1) Create an unmanaged instance group. Replicate the static content across the VMs. 2) Create an external HTTP(S) load balancer 3) Enable Cloud CDN, and send traffic to the unmanaged instance group.



C. 1) Create a Standard storage class, regional Cloud Storage bucket. Put the static content in the bucket 2) Reserve an external IP address, and create an external HTTP(S) load balancer 3) Enable Cloud CDN, and send traffic to your backend bucket

D. 1) Create a Standard storage class, multi-regional Cloud Storage bucket. Put the static content in the bucket. 2) Reserve an external IP address, and create an external HTTP(S) load balancer. 3) Enable Cloud CDN, and send traffic to your backend bucket.

Correct Answer: D

QUESTION 4

You are deploying your applications on Compute Engine. One of your Compute Engine instances failed to launch. What should you do? (Choose two.)

- A. Determine whether your file system is corrupted.
- B. Access Compute Engine as a different SSH user.
- C. Troubleshoot firewall rules or routes on an instance.
- D. Check whether your instance boot disk is completely full.
- E. Check whether network traffic to or from your instance is being dropped.

Correct Answer: AD

<https://cloud.google.com/compute/docs/troubleshooting/vm-startup>

QUESTION 5

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

https://cloud.google.com/stackdriver/docs/solutions/gke/managing-logs#best_practices



VCE & PDF

PassApply.com

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

2024 Latest passapply PROFESSIONAL-CLOUD-DEVELOPER PDF and VCE dumps Download

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

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

[PROFESSIONAL-CLOUD-DEVELOPER Exam Questions](#)