

## Q&As

Professional Cloud Architect on Google Cloud Platform

# Pass Google PROFESSIONAL-CLOUD-ARCHITECT Exam with 100% Guarantee

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

<https://www.pass2lead.com/professional-cloud-architect.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 wants you to build a highly reliable web application with a few public APIs as the backend. You don't expect a lot of user traffic, but traffic could spike occasionally. You want to leverage Cloud Load Balancing, and the solution must be cost-effective for users. What should you do?

- A. Store static content such as HTML and images in Cloud CDN. Host the APIs on App Engine and store the user data in Cloud SQL.
- B. Store static content such as HTML and images in a Cloud Storage bucket. Host the APIs on a zonal Google Kubernetes Engine cluster with worker nodes in multiple zones, and save the user data in Cloud Spanner.
- C. Store static content such as HTML and images in Cloud CDN. Use Cloud Run to host the APIs and save the user data in Cloud SQL.
- D. Store static content such as HTML and images in a Cloud Storage bucket. Use Cloud Functions to host the APIs and save the user data in Firestore.

Correct Answer: D

---

### QUESTION 2

Your company has a support ticketing solution that uses App Engine Standard. The project that contains the App Engine application already has a Virtual Private Cloud (VPC) network fully connected to the company's on-premises environment through a Cloud VPN tunnel. You want to enable the App Engine application to communicate with a database that is running in the company's on-premises environment. What should you do?

- A. Configure private Google access for on-premises hosts only.
- B. Configure private Google access.
- C. Configure private services access.
- D. Configure serverless VPC access.

Correct Answer: D

<https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc>

[https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on\\_premises](https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises)

---

### QUESTION 3

You are designing a large distributed application with 30 microservices. Each of your distributed microservices needs to connect to a database back-end. You want to store the credentials securely.

Where should you store the credentials?

- A. In the source code
- B. In an environment variable

- C. In a secret management system
- D. In a config file that has restricted access through ACLs

Correct Answer: C

References: <https://cloud.google.com/kms/docs/secret-management>

---

#### QUESTION 4

Your company has an application that is running on multiple instances of Compute Engine. It generates 1 TB per day of logs. For compliance reasons, the logs need to be kept for at least two years. The logs need to be available for active query for 30 days. After that, they just need to be retained for audit purposes. You want to implement a storage solution that is compliant, minimizes costs, and follows Google-recommended practices. What should you do?

- A. 1. Install a Cloud Logging agent on all instances.  
2.  
Create a sink to export logs into a regional Cloud Storage bucket.  
3.  
Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.  
4.  
Configure a retention policy at the bucket level using bucket lock.
- B. 1. Write a daily cron job, running on all instances, that uploads logs into a Cloud Storage bucket.  
2.  
Create a sink to export logs into a regional Cloud Storage bucket.  
3.  
Create an Object Lifecycle rule to move files into a Coldline Cloud Storage bucket after one month.
- C. 1. Install a Cloud Logging agent on all instances.  
2.  
Create a sink to export logs into a partitioned BigQuery table.  
3.  
Set a `time_partitioning_expiration` of 30 days.
- D. 1. Create a daily cron job, running on all instances, that uploads logs into a partitioned BigQuery table.  
2. Set a `time_partitioning_expiration` of 30 days.

Correct Answer: A

---

**QUESTION 5**

Your development teams release new versions of games running on Google Kubernetes Engine (GKE) daily. You want to create service level indicators (SLIs) to evaluate the quality of the new versions from the user's perspective. What should you do?

- A. Create CPU Utilization and Request Latency as service level indicators.
- B. Create GKE CPU Utilization and Memory Utilization as service level indicators.
- C. Create Request Latency and Error Rate as service level indicators.
- D. Create Server Uptime and Error Rate as service level indicators.

Correct Answer: C

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

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

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