



## 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.passapply.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



VCE & PDF

PassApply.com

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

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

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





## QUESTION 1

Mountkirk Games wants to set up a real-time analytics platform for their new game. The new platform must meet their technical requirements.

Which combination of Google technologies will meet all of their requirements?

- A. Kubernetes Engine, Cloud Pub/Sub, and Cloud SQL
- B. Cloud Dataflow, Cloud Storage, Cloud Pub/Sub, and BigQuery
- C. Cloud SQL, Cloud Storage, Cloud Pub/Sub, and Cloud Dataflow
- D. Cloud Dataproc, Cloud Pub/Sub, Cloud SQL, and Cloud Dataflow
- E. Cloud Pub/Sub, Compute Engine, Cloud Storage, and Cloud Dataproc

Correct Answer: B

Ingest millions of streaming events per second from anywhere in the world with Cloud Pub/Sub, powered by Google's unique, high-speed private network. Process the streams with Cloud Dataflow to ensure reliable, exactly-once, low-latency data transformation. Stream the transformed data into BigQuery, the cloud-native data warehousing service, for immediate analysis via SQL or popular visualization tools.

From scenario: They plan to deploy the game's backend on Google Compute Engine so they can capture streaming metrics, run intensive analytics.

Requirements for Game Analytics Platform

1.  
Dynamically scale up or down based on game activity
2.  
Process incoming data on the fly directly from the game servers
3.  
Process data that arrives late because of slow mobile networks
4.  
Allow SQL queries to access at least 10 TB of historical data
5.  
Process files that are regularly uploaded by users' mobile devices
6.  
Use only fully managed services

References: <https://cloud.google.com/solutions/big-data/stream-analytics/>



## QUESTION 2

You are configuring the cloud network architecture for a newly created project in Google Cloud that will host applications in Compute Engine. Compute Engine virtual machine instances will be created in two different subnets (sub-a and sub-b) within a single region.

1.

Instances in sub-a will have public IP addresses.

2.

Instances in sub-b will have only private IP addresses.

To download updated packages, instances must connect to a public repository outside the boundaries of Google Cloud. You need to allow sub-b to access the external repository.

What should you do?

- A. Enable Private Google Access on sub-b
- B. Configure Cloud NAT and select sub-b in the NAT mapping section
- C. Configure a bastion host instance in sub-a to connect to instances in sub-b
- D. Enable Identity Aware Proxy for TCP forwarding for instances in sub-b

Correct Answer: B

Cloud NAT (network address translation) lets Google Cloud virtual machine (VM) instances without external IP addresses and private Google Kubernetes Engine (GKE) clusters send outbound packets to the internet and receive any corresponding established inbound response packets. By configuring Cloud NAT and selecting sub-b in the NAT mapping section, you can allow instances in sub-b to access the external repository without exposing them to the internet<sup>1</sup>.

---

## QUESTION 3

Your team is developing a web application that will be deployed on Google Kubernetes Engine (GKE). Your CTO expects a successful launch and you need to ensure your application can handle the expected load of tens of thousands of users. You want to test the current deployment to ensure the latency of your application stays below a certain threshold. What should you do?

- A. Use a load testing tool to simulate the expected number of concurrent users and total requests to your application, and inspect the results.
- B. Enable autoscaling on the GKE cluster and enable horizontal pod autoscaling on your application deployments. Send curl requests to your application, and validate if the auto scaling works.
- C. Replicate the application over multiple GKE clusters in every Google Cloud region. Configure a global HTTP(S) load balancer to expose the different clusters over a single global IP address.
- D. Use Cloud Debugger in the development environment to understand the latency between the different microservices.



Correct Answer: B

---

#### QUESTION 4

You are analyzing and defining business processes to support your startup's trial usage of GCP, and you don't yet know what consumer demand for your product will be.

Your manager requires you to minimize GCP service costs and adhere to Google best practices.

What should you do?

- A. Utilize free tier and sustained use discounts. Provision a staff position for service cost management.
- B. Utilize free tier and sustained use discounts. Provide training to the team about service cost management.
- C. Utilize free tier and committed use discounts. Provision a staff position for service cost management.
- D. Utilize free tier and committed use discounts. Provide training to the team about service cost management.

Correct Answer: D

[https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing\\_and\\_management](https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#billing_and_management)

---

#### QUESTION 5

Your company has an enterprise application running on Compute Engine that requires high availability and high performance. The application has been deployed on two instances in two zones in the same region in active passive mode. The application writes data to a persistent disk in the case of a single zone outage that data should be immediately made available to the other instance in the other zone. You want to maximize performance while minimizing downtime and data loss. What should you do?

- A. 1. Attach a persistent SSD disk to the first instance

2.

Create a snapshot every hour

3.

In case of a zone outage, recreate a persistent SSD disk in the second instance where data is coming from the created snapshot

- B. 1 Create a Cloud Storage bucket

2.

Mount the bucket into the first instance with gcs-fuse

3.

In case of a zone outage, mount the Cloud Storage bucket to the second instance with gcs-fuse



C. 1. Attach a local SSD to the first instance disk

2.

Execute an rsync command every hour where the target is a persistent SSD disk attached to the second instance

3.

In case of a zone outage, use the second instance

D. 1. Attach a regional SSD persistent disk to the first instance

2. In case of a zone outage, force-attach the disk to the other instance

Correct Answer: D

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)