



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

The current Dress4win system architecture has high latency to some customers because it is located in one data center.

As of a future evaluation and optimizing for performance in the cloud, Dress4win wants to distribute its system architecture to multiple locations when Google cloud platform.

Which approach should they use?

- A. Use regional managed instance groups and a global load balancer to increase performance because the regional managed instance group can grow instances in each region separately based on traffic.
- B. Use a global load balancer with a set of virtual machines that forward the requests to a closer group of virtual machines managed by your operations team.
- C. Use regional managed instance groups and a global load balancer to increase reliability by providing automatic failover between zones in different regions.
- D. Use a global load balancer with a set of virtual machines that forward the requests to a closer group of virtual machines as part of a separate managed instance groups.

Correct Answer: A

QUESTION 3

Your company has a Google Workspace account and Google Cloud Organization Some developers in the company have created Google Cloud projects outside of the Google Cloud Organization

You want to create an Organization structure that allows developers to create projects, but prevents them from modifying production projects You want to manage policies for all projects centrally and be able to set more restrictive policies for production projects

You want to minimize disruption to users and developers when business needs change in the future You want to follow Google-recommended practices How should you design the Organization structure?

- A. 1 Create a second Google Workspace account and Organization 2 Grant all developers the Project Creator IAM role on the new Organization 3 Move the developer projects into the new Organization 4 Set the policies for all projects on both Organizations. 5 Additionally set the production policies on the original Organization
- B. 1 Create a folder under the Organization resource named "Production" 2 Grant all developers the Project Creator IAM role on the Organization 3. Move the developer projects into the Organization 4 Set the policies for all projects on the Organization 5 Additionally set the production policies on the "Production" folder
- C. 1 Create folders under the Organization resource named "Development" and "Production" 2 Grant all developers the Project Creator IAM role on the "Development" folder 3. Move the developer projects into the "Development" folder 4 Set the policies for all projects on the Organization 5 Additionally set the production policies on the "Production" folder
- D. 1 Designate the Organization for production projects only 2 Ensure that developers do not have the Project Creator IAM role on the Organization 3 Create development projects outside of the Organization using the developer Google Workspace accounts 4 Set the policies for all projects on the Organization 5 Additionally set the production policies on



the individual production projects

Correct Answer: C

This option can help create an organization structure that allows developers to create projects, but prevents them from modifying production projects. Folders are containers for projects and other folders within Google Cloud organizations. Folders allow resources to be structured hierarchically and inherit policies from their parent resources. By creating folders under the organization resource named "Development" and "Production", you can organize your projects by environment and apply different policies to them. By granting all developers the Project Creator IAM role on the "Development" folder, you can allow them to create projects under that folder, but not under the "Production" folder. By moving the developer projects into the "Development" folder, you can ensure that they are subject to the policies set on that folder. By setting the policies for all projects on the organization, you can manage policies centrally and efficiently. By additionally setting the production policies on the "Production" folder, you can enforce more restrictive policies for production projects and prevent developers from modifying them. The other options are not optimal for this scenario, because they either create a second Google Workspace account and organization, which increases complexity and cost (A), or do not use folders to organize projects by environment, which makes it harder to manage policies and permissions (B, D). References: <https://cloud.google.com/resource-manager/docs/creating-managing-folders>
<https://cloud.google.com/architecture/framework/system-design>

QUESTION 4

You want your Google Kubernetes Engine cluster to automatically add or remove nodes based on CPUload. What should you do?

- A. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable the Cluster Autoscaler from the GCP Console.
- B. Configure a HorizontalPodAutoscaler with a target CPU usage. Enable autoscaling on the managed instance group for the cluster using the gcloud command.
- C. Create a deployment and set the maxUnavailable and maxSurge properties. Enable the Cluster Autoscaler using the gcloud command.
- D. Create a deployment and set the maxUnavailable and maxSurge properties. Enable autoscaling on the cluster managed instance group from the GCP Console.

Correct Answer: B

QUESTION 5

Your company uses Google Kubernetes Engine (GKE) as a platform for all workloads. Your company has a single large GKE cluster that contains batch, stateful, and stateless workloads. The GKE cluster is configured with a single node pool with 200 nodes. Your company needs to reduce the cost of this cluster but does not want to compromise availability. What should you do?

- A. Create a second GKE cluster for the batch workloads only. Allocate the 200 original nodes across both clusters.
- B. Configure a HorizontalPodAutoscaler for all stateless workloads and for all compatible stateful workloads. Configure the cluster to use node auto scaling.
- C. Configure CPU and memory limits on the namespaces in the cluster. Configure all Pods to have a CPU and memory



limits.

D. Change the node pool to use spot VMs.

Correct Answer: C

One way to reduce the cost of a Google Kubernetes Engine (GKE) cluster without compromising availability is to use horizontal pod autoscalers (HPA) and node auto scaling. HPA allows you to automatically scale the number of Pods in a deployment based on the resource usage of the Pods. By configuring HPA for stateless workloads and for compatible stateful workloads, you can ensure that the number of Pods is automatically adjusted based on the actual resource usage, which can help to reduce costs. Node auto scaling allows you to automatically add or remove nodes from the node pool based on the resource usage of the cluster. By configuring node auto scaling, you can ensure that the cluster has the minimum number of nodes needed to meet the resource requirements of the workloads, which can also help to reduce costs.

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)