



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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-architect.html>

2024 Latest pass4itsure 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

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version.

What should you do?

- A. Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.
- B. Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
- C. Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
- D. Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Correct Answer: B

<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

QUESTION 2

Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asia. You want to reduce the latency for the users in Asia. What should you do?

- A. Use a global HTTP(s) load balancer with Cloud CDN enabled
- B. Create a second GKE cluster in asia-southeast1, and expose both APIs using a Service of type Load Balancer. Add the public IPs to the Cloud DNS zone
- C. Increase the memory and CPU allocated to the application in the cluster
- D. Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Correct Answer: D

https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works

<https://github.com/GoogleCloudPlatform/k8s-multicloud-ingress> <https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci>

QUESTION 3

You have an application that will run on Compute Engine. You need to design an architecture that takes into account a disaster recovery plan that requires your application to fail over to another region in case of a regional outage. What should you do?

- A. Deploy the application on two Compute Engine instances in the same project but in a different region. Use the first instance to serve traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a



disaster.

B. Deploy the application on a Compute Engine instance. Use the instance to serve traffic, and use the HTTP load balancing service to fail over to an instance on your premises in case of a disaster.

C. Deploy the application on two Compute Engine instance groups, each in the same project but in a different region. Use the first instance group to serve traffic, and use the HTTP load balancing service to fail over to the standby instance group in case of a disaster.

D. Deploy the application on two Compute Engine instance groups, each in separate project and a different region. Use the first instance group to server traffic, and use the HTTP load balancing service to fail over to the standby instance in case of a disaster.

Correct Answer: C

QUESTION 4

You are developing your microservices application on Google Kubernetes Engine. During testing, you want to validate the behavior of your application in case a specific microservice should suddenly crash. What should you do?

A. Add a taint to one of the nodes of the Kubernetes cluster. For the specific microservice, configure a pod anti-affinity label that has the name of the tainted node as a value.

B. Use Istio's fault injection on the particular microservice whose faulty behavior you want to simulate.

C. Destroy one of the nodes of the Kubernetes cluster to observe the behavior.

D. Configure Istio's traffic management features to steer the traffic away from a crashing microservice.

Correct Answer: B

Microservice runs on all nodes. The Micro service runs on Pod, Pod runs on Nodes. Nodes is nothing but Virtual machines. Once deployed the application microservices will get deployed across all Nodes. Destroying one node may not mimic the behaviour of microservice crashing as it may be running in other nodes. link:

<https://istio.io/latest/docs/tasks/traffic-management/fault-injection/>

QUESTION 5

Your applications will be writing their logs to BigQuery for analysis. Each application should have its own table.

Any logs older than 45 days should be removed. You want to optimize storage and follow Google recommended practices. What should you do?

A. Configure the expiration time for your tables at 45 days

B. Make the tables time-partitioned, and configure the partition expiration at 45 days

C. Rely on BigQuery's default behavior to prune application logs older than 45 days

D. Create a script that uses the BigQuery command line tool (bq) to remove records older than 45 days

Correct Answer: B



<https://cloud.google.com/bigquery/docs/managing-partitioned-tables>

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

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

[PROFESSIONAL-CLOUD-ARCHITECT Braindumps](#)