



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.pass4itsure.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

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-developer.html>
2024 Latest pass4itsure 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

You are designing a resource-sharing policy for applications used by different teams in a Google Kubernetes Engine cluster.

You need to ensure that all applications can access the resources needed to run. What should you do? (Choose two.)

- A. Specify the resource limits and requests in the object specifications.
- B. Create a namespace for each team, and attach resource quotas to each namespace.
- C. Create a LimitRange to specify the default compute resource requirements for each namespace.
- D. Create a Kubernetes service account (KSA) for each application, and assign each KSA to the namespace.
- E. Use the Anthos Policy Controller to enforce label annotations on all namespaces. Use taints and tolerations to allow resource sharing for namespaces.

Correct Answer: BC

<https://kubernetes.io/docs/concepts/policy/resource-quotas/>

<https://kubernetes.io/docs/concepts/policy/limit-range/>

<https://cloud.google.com/blog/products/containers-kubernetes/kubernetes-best-practices-resource-requests-and-limits>

QUESTION 2

Your code is running on Cloud Functions in project A. It is supposed to write an object in a Cloud Storage bucket owned by project B. However, the write call is failing with the error "403 Forbidden". What should you do to correct the problem?

- A. Grant your user account the roles/storage.objectCreator role for the Cloud Storage bucket.
- B. Grant your user account the roles/iam.serviceAccountUser role for the service-PROJECTA@gcf-adminrobot.iam.gserviceaccount.com service account.
- C. Grant the service-PROJECTA@gcf-admin-robot.iam.gserviceaccount.com service account the roles/storage.objectCreator role for the Cloud Storage bucket.
- D. Enable the Cloud Storage API in project B.

Correct Answer: B

QUESTION 3

You have an application that uses an HTTP Cloud Function to process user activity from both desktop browser and mobile application clients. This function will serve as the endpoint for all metric submissions using HTTP POST.



Due to legacy restrictions, the function must be mapped to a domain that is separate from the domain requested by users on web or mobile sessions. The domain for the Cloud Function is `https://fn.example.com`. Desktop and mobile clients

use the domain `https://www.example.com`. You need to add a header to the function's HTTP response so that only those browser and mobile sessions can submit metrics to the Cloud Function.

Which response header should you add?

- A. Access-Control-Allow-Origin: *
- B. Access-Control-Allow-Origin: `https://*.example.com`
- C. Access-Control-Allow-Origin: `https://fn.example.com`
- D. Access-Control-Allow-origin: `https://www.example.com`

Correct Answer: D

QUESTION 4

Your company is planning to migrate their on-premises Hadoop environment to the cloud. Increasing storage cost and maintenance of data stored in HDFS is a major concern for your company. You also want to make minimal changes to existing data analytics jobs and existing architecture. How should you proceed with the migration?

- A. Migrate your data stored in Hadoop to BigQuery. Change your jobs to source their information from BigQuery instead of the on-premises Hadoop environment.
- B. Create Compute Engine instances with HDD instead of SSD to save costs. Then perform a full migration of your existing environment into the new one in Compute Engine instances.
- C. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop environment to the new Cloud Dataproc cluster. Move your HDFS data into larger HDD disks to save on storage costs.
- D. Create a Cloud Dataproc cluster on Google Cloud Platform, and then migrate your Hadoop code objects to the new cluster. Move your data to Cloud Storage and leverage the Cloud Dataproc connector to run jobs on that data.

Correct Answer: D

QUESTION 5

Your team develops stateless services that run on Google Kubernetes Engine (GKE). You need to deploy a new service that will only be accessed by other services running in the GKE cluster. The service will need to scale as quickly as possible to respond to changing load. What should you do?

- A. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.
- B. Use a Vertical Pod Autoscaler to scale the containers, and expose them via a NodePort Service.
- C. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a ClusterIP Service.



D. Use a Horizontal Pod Autoscaler to scale the containers, and expose them via a NodePort Service.

Correct Answer: C

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

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

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

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)