

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.leads4pass.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

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



QUESTION 1

You are developing an application that will be launched on Compute Engine instances into multiple distinct projects, each corresponding to the environments in your software development process (development, QA, staging, and production). The instances in each project have the same application code but a different configuration. During deployment, each instance should receive the application's configuration based on the environment it serves. You want to minimize the number of steps to configure this flow.

What should you do?

- A. When creating your instances, configure a startup script using the `gcloud` command to determine the project name that indicates the correct environment.
- B. In each project, configure a metadata key "environment" whose value is the environment it serves. Use your deployment tool to query the instance metadata and configure the application based on the "environment" value.
- C. Deploy your chosen deployment tool on an instance in each project. Use a deployment job to retrieve the appropriate configuration file from your version control system, and apply the configuration when deploying the application on each instance.
- D. During each instance launch, configure an instance custom-metadata key named "environment" whose value is the environment the instance serves. Use your deployment tool to query the instance metadata, and configure the application based on the "environment" value.

Correct Answer: B

Reference: <https://cloud.google.com/compute/docs/metadata/overview>

QUESTION 2

For this question, refer to the HipLocal case study.

How should HipLocal increase their API development speed while continuing to provide the QA team with a stable testing environment that meets feature requirements?

- A. Include unit tests in their code, and prevent deployments to QA until all tests have a passing status.
- B. Include performance tests in their code, and prevent deployments to QA until all tests have a passing status.
- C. Create health checks for the QA environment, and redeploy the APIs at a later time if the environment is unhealthy.
- D. Redeploy the APIs to App Engine using Traffic Splitting. Do not move QA traffic to the new versions if errors are found.

Correct Answer: B

QUESTION 3

You recently migrated a monolithic application to Google Cloud by breaking it down into microservices. One of the

microservices is deployed using Cloud Functions. As you modernize the application, you make a change to the API of the service that is backward-incompatible. You need to support both existing callers who use the original API and new callers who use the new API. What should you do?

- A. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use a load balancer to distribute calls between the versions.
- B. Leave the original Cloud Function as-is and deploy a second Cloud Function that includes only the changed API. Calls are automatically routed to the correct function.
- C. Leave the original Cloud Function as-is and deploy a second Cloud Function with the new API. Use Cloud Endpoints to provide an API gateway that exposes a versioned API.
- D. Re-deploy the Cloud Function after making code changes to support the new API. Requests for both versions of the API are fulfilled based on a version identifier included in the call.

Correct Answer: D

Reference: <https://cloud.google.com/endpoints/docs/openapi/versioning-an-api>

QUESTION 4

This architectural diagram depicts a system that streams data from thousands of devices. You want to ingest data into a pipeline, store the data, and analyze the data using SQL statements. Which Google Cloud services should you use for steps 1, 2, 3, and 4?

- A.
- ```
livenessProbe:
 httpGet:
 path: /healthz
 port: 80
```
- B.
- ```
readinessProbe:  
  httpGet:  
    path: /healthz  
    port: 80
```
- C.
- ```
loadbalancerHealthCheck:
 httpGet:
 path: /healthz
 port: 80
```
- D.
- ```
healthCheck:  
  httpGet:  
    path: /healthz  
    port: 80
```

A. 1) App Engine 2) Pub/Sub 3) BigQuery 4) Firestore

B. 1) Dataflow 2) Pub/Sub 3) Firestore 4) BigQuery

C. 1) Pub/Sub 2) Dataflow 3) BigQuery 4) Firestore

D. 1) Pub/Sub 2) Dataflow 3) Firestore 4) BigQuery

Correct Answer: D

QUESTION 5

Your company has a BigQuery data mart that provides analytics information to hundreds of employees. One user of wants to run jobs without interrupting important workloads. This user isn't concerned about the time it takes to run

these jobs. You want to fulfill this request while minimizing cost to the company and the effort required on your part.

What should you do?

- A. Ask the user to run the jobs as batch jobs.
- B. Create a separate project for the user to run jobs.
- C. Add the user as a job.user role in the existing project.
- D. Allow the user to run jobs when important workloads are not running.

Correct Answer: B

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

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

[PROFESSIONAL-CLOUD-DEVELOPER Braindumps](#)