



# Free Questions for Professional-Cloud-Developer

Shared by Farrell on 29-01-2024

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



## Question 1

---

Question Type: MultipleChoice

---

Case Study: Mix Questions

---

### Mix Questions

#### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your development team has built several Cloud Functions using Java along with corresponding integration and service tests. You are building and deploying the functions and launching the tests using Cloud Build. Your Cloud Build job is reporting deployment failures immediately after successfully validating the code. What should you do?

#### Options:

---

- A- Check the maximum number of Cloud Function instances.
- B- Verify that your Cloud Build trigger has the correct build parameters.
- C- Retry the tests using the truncated exponential backoff polling strategy.
- D- Verify that the Cloud Build service account is assigned the Cloud Functions Developer role.

#### Answer:

---

D

#### Explanation:

---

<https://cloud.google.com/build/docs/securing-builds/configure-access-for-cloud-build-service-account>

## Question 2

---

Question Type: MultipleChoice

---

Case Study: Mix Questions

---

### Mix Questions

#### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You recently deployed a Go application on Google Kubernetes Engine (GKE). The operations team has noticed that the application's CPU usage is high even when there is low production traffic.

The operations team has asked you to optimize your application's CPU resource consumption. You want to determine which Go functions consume the largest amount of CPU. What should you do?

### Options:

---

- A- Deploy a Fluent Bit daemonset on the GKE cluster to log data in Cloud Logging. Analyze the logs to get insights into your application code's performance.
- B- Create a custom dashboard in Cloud Monitoring to evaluate the CPU performance metrics of your application.
- C- Connect to your GKE nodes using SSH. Run the top command on the shell to extract the CPU utilization of your application.
- D- Modify your Go application to capture profiling data. Analyze the CPU metrics of your application in flame graphs in Profiler.

### Answer:

---

D

### Explanation:

---

<https://cloud.google.com/profiler/docs/about-profiler>

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation information from your production applications. It attributes that information to the source code that generated it, helping you identify the parts of your application that are consuming the most resources, and otherwise illuminating your applications performance characteristics.

<https://cloud.google.com/profiler/docs>

## Question 3

---

**Question Type:** MultipleChoice

---

**Case Study:** Mix Questions

---

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You are working on a new application that is deployed on Cloud Run and uses Cloud Functions. Each time new features are added, new Cloud Functions and Cloud Run services are deployed.

You use ENV variables to keep track of the services and enable interservice communication but the maintenance of the ENV variables has become difficult. You want to implement dynamic discovery in a scalable way. What should you do?

### Options:

---

- A- Create a Service Directory Namespace Use API calls to register the services during deployment, and query during runtime.
- B- Configure your microservices to use the Cloud Run Admin and Cloud Functions APIs to query for deployed Cloud Run services and Cloud Functions in the Google Cloud project.
- C- Deploy Hashicorp Consul on a single Compute Engine Instance Register the services with Consul during deployment and query during runtime
- D- Rename the Cloud Functions and Cloud Run services endpoints using a well-documented naming convention

### Answer:

---

B

## Question 4

---

Question Type: MultipleChoice

---

Case Study: Mix Questions

---

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your application stores customers' content in a Cloud Storage bucket, with each object being encrypted with the customer's encryption key. The key for each object in Cloud Storage is entered into your application by the customer. You discover that your application is receiving an HTTP 4xx error when reading the object from Cloud Storage What is a possible cause of this error?

### Options:

---

- A- You attempted the read operation without the base64-encoded SHA256 hash of the encryption key.
- B- You entered the same encryption algorithm specified by the customer when attempting the read operation.
- C- You attempted the read operation on the object with the base64-encoded SHA256 hash of the

customer's key.

D- You attempted the read operation on the object with the customers base64-encoded key.

Answer:

D

## Question 5

Question Type: MultipleChoice

Case Study: Mix Questions

### Mix Questions

#### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your team manages a Google Kubernetes Engine (GKE) cluster where an application is running. A different team is planning to integrate with this application. Before they start the integration, you need to ensure that the other team cannot make changes to your application, but they can deploy the integration on GKE. What should you do?

Options:

- A- Using Identity and Access Management (IAM), grant the Viewer IAM role on the cluster project to the other team.
- B- Create a new GKE cluster. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- C- Create a new namespace in the existing cluster. Using Identity and Access Management (IAM), grant the Editor role on the cluster project to the other team.
- D- Create a new namespace in the existing cluster. Using Kubernetes role-based access control (RBAC), grant the Admin role on the new namespace to the other team.

Answer:

D

## Question 6

Question Type: MultipleChoice

Case Study: Mix Questions

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your application is deployed in a Google Kubernetes Engine (GKE) cluster. You want to expose this application publicly behind a Cloud Load Balancing HTTP(S) load balancer. What should you do?

#### Options:

---

- A- Configure a GKE Ingress resource.
- B- Configure a GKE Service resource.
- C- Configure a GKE Ingress resource with type: LoadBalancer.
- D- Configure a GKE Service resource with type: LoadBalancer.

#### Answer:

---

A

## Question 7

---

**Question Type:** MultipleChoice

---

**Case Study:** Mix Questions

---

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You have a web application that publishes messages to Pub/Sub. You plan to build new versions of the application locally and need to quickly test Pub/Sub integration for each new build. How should you configure local testing?

#### Options:

---

- A- Run the `gcloud config set api_endpoint_overrides/pubsub https://pubsubemulator.googleapis.com` command to change the Pub/Sub endpoint prior to starting the application
- B- In the Google Cloud console, navigate to the API Library and enable the Pub/Sub API. When developing locally, configure your application to call `pubsub.googleapis.com`
- C- Install Cloud Code on the integrated development environment (IDE). Navigate to Cloud APIs, and enable Pub/Sub against a valid Google Project ID. When developing locally, configure your application to call `pubsub.googleapis.com`

D- Install the Pub/Sub emulator using gcloud and start the emulator with a valid Google Project ID. When developing locally, configure your application to use the local emulator by exporting the fuhsb emulator Host variable

Answer:

---

C

## Question 8

Question Type: MultipleChoice

Case Study: Mix Questions

### Mix Questions

#### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

Your team develops services that run on Google Cloud. You want to process messages sent to a Pub/Sub topic, and then store them. Each message must be processed exactly once to avoid duplication of data and any data conflicts. You need to use the cheapest and most simple solution. What should you do?

Options:

---

A- Process the messages with a Dataproc job, and write the output to storage.

B- Process the messages with a Dataflow streaming pipeline using Apache Beam's PubSubIO package, and write the output to storage.

C- Process the messages with a Cloud Function, and write the results to a BigQuery location where you can run a job to deduplicate the data.

D- Retrieve the messages with a Dataflow streaming pipeline, store them in Cloud Bigtable, and use another Dataflow streaming pipeline to deduplicate messages.

Answer:

---

B

Explanation:

---

<https://cloud.google.com/dataflow/docs/concepts/streaming-with-cloud-pubsub>

## Question 9

---

Question Type: MultipleChoice

---

Case Study: Mix Questions

---

### Mix Questions

#### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You are in the final stage of migrating an on-premises data center to Google Cloud. You are quickly approaching your deadline, and discover that a web API is running on a server slated for decommissioning. You need to recommend a solution to modernize this API while migrating to Google Cloud. The modernized web API must meet the following requirements:

- \* Autoscales during high traffic periods at the end of each month
- \* Written in Python 3.x
- \* Developers must be able to rapidly deploy new versions in response to frequent code changes

You want to minimize cost, effort, and operational overhead of this migration. What should you do?

#### Options:

---

- A- Modernize and deploy the code on App Engine flexible environment.
- B- Modernize and deploy the code on App Engine standard environment.
- C- Deploy the modernized application to an n1-standard-1 Compute Engine instance.
- D- Ask the development team to re-write the application to run as a Docker container on Google Kubernetes Engine.

#### Answer:

---

B

#### Explanation:

---

<https://cloud.google.com/appengine/docs/standard>

## Question 10

---

Question Type: MultipleChoice

---

Case Study: Mix Questions

---

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You want to migrate an on-premises container running in Knative to Google Cloud. You need to make sure that the migration doesn't affect your application's deployment strategy, and you want to use a fully managed service. Which Google Cloud service should you use to deploy your container?

Options:

- A- Cloud Run
- B- Compute Engine
- C- Google Kubernetes Engine
- D- App Engine flexible environment

Answer:

A

Explanation:

<https://cloud.google.com/blog/products/serverless/knative-based-cloud-run-services-are-ga>

## Question 11

Question Type: MultipleChoice

Case Study: Mix Questions

---

## Mix Questions

### Professional-Cloud-Developer Mix Questions IN THIS CASE STUDY

You are deploying a microservices application to Google Kubernetes Engine (GKE). The application will receive daily updates. You expect to deploy a large number of distinct containers that will run on the Linux operating system (OS). You want to be alerted to any known OS vulnerabilities in the new containers. You want to follow Google-recommended best practices. What should you do?

### Options:

---

- A- Use the gcloud CLI to call Container Analysis to scan new container images. Review the vulnerability results before each deployment.
- B- Enable Container Analysis, and upload new container images to Artifact Registry. Review the vulnerability results before each deployment.
- C- Enable Container Analysis, and upload new container images to Artifact Registry. Review the critical vulnerability results before each deployment.
- D- Use the Container Analysis REST API to call Container Analysis to scan new container images. Review the vulnerability results before each deployment.

### Answer:

---

D



### Explanation:

---

<https://cloud.google.com/container-analysis/docs/automated-scanning-howto>

<https://cloud.google.com/container-analysis/docs/os-overview> says: The Container Scanning API allows you to automate OS vulnerability detection, scanning each time you push an image to Container Registry or Artifact Registry. Enabling this API also triggers language package scans for Go and Java vulnerabilities (Preview).



To Get Premium Files for Professional-Cloud-Developer Visit

<https://www.p2pexams.com/products/professional-cloud-developer>



For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-developer>

