Google Professional-Cloud-Developer Practice Questions

Shared by Ballard on 24-08-2026

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

You want to view the memory usage of your application deployed on Compute Engine. What should you do?

Options:

A- Install the Stackdriver Client Library.
B- Install the Stackdriver Monitoring Agent.
C- Use the Stackdriver Metrics Explorer.
D- Use the Google Cloud Platform Console.

Answer:

C

Question 2

Question Type: MultipleChoice

You are deploying your applications on Compute Engine. One of your Compute Engine instances failed to launch. What should you do? (Choose two.)

Options:

A- Determine whether your file system is corrupted.
B- Access Compute Engine as a different SSH user.
C- Troubleshoot firewall rules or routes on an instance.
D- Check whether your instance boot disk is completely full.
E- Check whether network traffic to or from your instance is being dropped.

Answer:

A, D

Explanation:

https://cloud.google.com/compute/docs/troubleshooting/vm-startup

Question 3

Question Type: MultipleChoice

You are reviewing and updating your Cloud Build steps to adhere to Google-recommended practices. Currently, your build steps include:

1. Pull the source code from a source repository.

2. Build a container image

3. Upload the built image to Artifact Registry.

You need to add a step to perform a vulnerability scan of the built container image, and you want the results of the scan to be available to your deployment pipeline running in Google Cloud. You want to minimize changes that could disrupt other teams' processes What should you do?

Options:

A- Enable Binary Authorization, and configure it to attest that no vulnerabilities exist in a container image.
B- Enable the Container Scanning API in Artifact Registry, and scan the built container images for vulnerabilities.
C- Upload the built container images to your Docker Hub instance, and scan them for vulnerabilities.
D- Add Artifact Registry to your Aqua Security instance, and scan the built container images for vulnerabilities

Answer:

B

Question 4

Question Type: MultipleChoice

You are building a new API. You want to minimize the cost of storing and reduce the latency of serving

images.

Which architecture should you use?

Options:

A- App Engine backed by Cloud Storage
B- Compute Engine backed by Persistent Disk
C- Transfer Appliance backed by Cloud Filestore
D- Cloud Content Delivery Network (CDN) backed by Cloud Storage

Answer:

B

Question 5

Question Type: MultipleChoice

You are developing a web application that contains private images and videos stored in a Cloud Storage bucket. Your users are anonymous and do not have Google Accounts. You want to use your application-specific logic to control access to the images and videos. How should you configure access?

Options:

A- Cache each web application user's IP address to create a named IP table using Google Cloud Armor. Create a Google Cloud Armor security policy that allows users to access the backend bucket.
B- Grant the Storage Object Viewer IAM role to allUsers. Allow users to access the bucket after authenticating through your web application.
C- Configure Identity-Aware Proxy (IAP) to authenticate users into the web application. Allow users to access the bucket after authenticating through IAP.
D- Generate a signed URL that grants read access to the bucket. Allow users to access the URL after authenticating through your web application.

Answer:

D

Explanation:

https://cloud.google.com/storage/docs/access-control/signed-urls#should-you-use

In some scenarios, you might not want to require your users to have a Google account in order to access Cloud Storage, but you still want to control access using your application-specific logic. The typical way to address this use case is to provide a signed URL to a user, which gives the user read, write, or delete access to that resource for a limited time. You specify an expiration time when you create the signed URL. Anyone who knows the URL can access the resource until the expiration time for the URL is reached or the key used to sign the URL is rotated.

Question 6

Question Type: MultipleChoice

Your App Engine standard configuration is as follows:

service: production

instance_class: B1

You want to limit the application to 5 instances. Which code snippet should you include in your configuration?

Options:

A- manual_scaling:instances: 5min_pending_latency: 30ms
B- manual_scaling:max_instances: 5idle_timeout: 10m
C- basic_scaling:instances: 5min_pending_latency: 30ms
D- basic_scaling:max_instances: 5idle_timeout: 10m

Answer:

C

Question 7

Question Type: MultipleChoice

HipLocal wants to improve the resilience of their MySQL deployment, while also meeting their

business and technical requirements.

Which configuration should they choose?

Options:

A- Use the current single instance MySQL on Compute Engine and several read-only MySQL servers on
Compute Engine.
B- Use the current single instance MySQL on Compute Engine, and replicate the data to Cloud SQL in an
external master configuration.
C- Replace the current single instance MySQL instance with Cloud SQL, and configure high availability.
D- Replace the current single instance MySQL instance with Cloud SQL, and Google provides redundancy
without further configuration.

Answer:

B

Question 8

Question Type: MultipleChoice

You have an application deployed in Google Kubernetes Engine (GKE). You need to update the application to make authorized requests to Google Cloud managed services. You want this to be a one-time setup, and you need to follow security best practices of auto-rotating your security keys and storing them in an encrypted store. You already created a service account with appropriate access to the Google Cloud service. What should you do next?

Options:

A- Assign the Google Cloud service account to your GKE Pod using Workload Identity.
B- Export the Google Cloud service account, and share it with the Pod as a Kubernetes Secret.
C- Export the Google Cloud service account, and embed it in the source code of the application.
D- Export the Google Cloud service account, and upload it to HashiCorp Vault to generate a dynamic service account for your application.

Answer:

A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity

Applications running on GKE might need access to Google Cloud APIs such as Compute Engine API, BigQuery Storage API, or Machine Learning APIs.

Workload Identity allows a Kubernetes service account in your GKE cluster to act as an IAM service account. Pods that use the configured Kubernetes service account automatically authenticate as the IAM service account when accessing Google Cloud APIs. Using Workload Identity allows you to assign distinct, fine-grained identities and authorization for each application in your cluster.

Question 9

Question Type: MultipleChoice

You are developing a JPEG image-resizing API hosted on Google Kubernetes Engine (GKE). Callers of the service will exist within the same GKE cluster. You want clients to be able to get the IP address of the service.

What should you do?

Options:

A- Define a GKE Service. Clients should use the name of the A record in Cloud DNS to find the service's
cluster IP address.
B- Define a GKE Service. Clients should use the service name in the URL to connect to the service.
C- Define a GKE Endpoint. Clients should get the endpoint name from the appropriate environment variable in
the client container.
D- Define a GKE Endpoint. Clients should get the endpoint name from Cloud DNS.

Answer:

C

Question 10

Question Type: MultipleChoice

One of your deployed applications in Google Kubernetes Engine (GKE) is having intermittent performance issues. Your team uses a third-party logging solution. You want to install this solution on each node in your GKE cluster so you can view the logs. What should you do?

Options:

A- Deploy the third-party solution as a DaemonSet
B- Modify your container image to include the monitoring software
C- Use SSH to connect to the GKE node, and install the software manually
D- Deploy the third-party solution using Terraform and deploy the logging Pod as a Kubernetes Deployment

Answer:

A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns
DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluent-bit, and node monitoring daemons like collectd.

Question 11

Question Type: MultipleChoice

In order for HipLocal to store application state and meet their stated business requirements, which database service should they migrate to?

Options:

A- Cloud Spanner
B- Cloud Datastore
C- Cloud Memorystore as a cache
D- Separate Cloud SQL clusters for each region

Answer:

D

Question 12

Question Type: MultipleChoice

You are creating and running containers across different projects in Google Cloud. The application you are developing needs to access Google Cloud services from within Google Kubernetes Engine (GKE).

What should you do?

Options:

A- Assign a Google service account to the GKE nodes.
B- Use a Google service account to run the Pod with Workload Identity.
C- Store the Google service account credentials as a Kubernetes Secret.
D- Use a Google service account with GKE role-based access control (RBAC).

Answer:

B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/workload-identity