

# DUMPS ARENA

## Google Certified Professional - Cloud Developer

Google Professional-Cloud-Developer

Version Demo

Total Demo Questions: 15

Total Premium Questions: 248

Buy Premium PDF

<https://dumpsarena.co>

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)  
[dumpsarena.co](https://dumpsarena.co)

**QUESTION NO: 1**

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

- 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: D E****Explanation:**

Reference: <https://cloudacademy.com/course/deploying-applications-on-gcp-compute/deploying-applications-and-services-on-compute-engine/>

**QUESTION NO: 2**

You are running an application on App Engine that you inherited. You want to find out whether the application is using insecure binaries or is vulnerable to XSS attacks.

Which service should you use?

- A. Cloud Amor
- B. Stackdriver Debugger
- C. Cloud Security Scanner
- D. Stackdriver Error Reporting

**ANSWER: C****Explanation:**

Reference: <https://cloud.google.com/security-scanner>

**QUESTION NO: 3**

You manage an application that runs in a Compute Engine instance. You also have multiple backend services executing in stand-alone Docker containers running in Compute Engine instances. The Compute Engine instances supporting the backend services are scaled by managed instance groups in multiple regions. You want your calling application to be loosely coupled. You need to be able to invoke distinct service implementations that are chosen based on the value of an HTTP header found in the request. Which Google Cloud feature should you use to invoke the backend services?

- A. Traffic Director
- B. Service Directory
- C. Anthos Service Mesh
- D. Internal HTTP(S) Load Balancing

**ANSWER: D**

#### QUESTION NO: 4

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

- 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 NO: 5

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?

- 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 NO: 6**

You recently developed a new application. You want to deploy the application on Cloud Run without a Dockerfile. Your organization requires that all container images are pushed to a centrally managed container repository. How should you build your container using Google Cloud services? (Choose two.)

- A. Push your source code to Artifact Registry.
- B. Submit a Cloud Build job to push the image.
- C. Use the pack build command with pack CLI.
- D. Include the --source flag with the gcloud run deploy CLI command.
- E. Include the --platform=kubernetes flag with the gcloud run deploy CLI command.

**ANSWER: A C****Explanation:**

<https://cloud.google.com/run/docs/deploying#images>

<https://cloud.google.com/blog/products/containers-kubernetes/google-cloud-now-supports-buildpacks>

**QUESTION NO: 7**

You are deploying a microservices application to Google Kubernetes Engine (GKE) that will broadcast livestreams. You expect unpredictable traffic patterns and large variations in the number of concurrent users. Your application must meet the following requirements:

- Scales automatically during popular events and maintains high availability
- Is resilient in the event of hardware failures

How should you configure the deployment parameters? (Choose two.)

- A. Distribute your workload evenly using a multi-zonal node pool.
- B. Distribute your workload evenly using multiple zonal node pools.
- C. Use cluster autoscaler to resize the number of nodes in the node pool, and use a Horizontal Pod Autoscaler to scale the workload.
- D. Create a managed instance group for Compute Engine with the cluster nodes. Configure autoscaling rules for the managed instance group.

E. Create alerting policies in Cloud Monitoring based on GKE CPU and memory utilization. Ask an on-duty engineer to scale the workload by executing a script when CPU and memory usage exceed predefined thresholds.

**ANSWER: A C**

#### QUESTION NO: 8

Your team is developing unit tests for Cloud Function code. The code is stored in a Cloud Source Repositories repository. You are responsible for implementing the tests. Only a specific service account has the necessary permissions to deploy the code to Cloud Functions. You want to ensure that the code cannot be deployed without first passing the tests. How should you configure the unit testing process?

- A. Configure Cloud Build to deploy the Cloud Function. If the code passes the tests, a deployment approval is sent to you.
- B. Configure Cloud Build to deploy the Cloud Function, using the specific service account as the build agent. Run the unit tests after successful deployment.
- C. Configure Cloud Build to run the unit tests. If the code passes the tests, the developer deploys the Cloud Function.
- D. Configure Cloud Build to run the unit tests, using the specific service account as the build agent. If the code passes the tests, Cloud Build deploys the Cloud Function.

**ANSWER: D**

#### QUESTION NO: 9

You are developing a microservice-based application that will run on Google Kubernetes Engine (GKE). Some of the services need to access different Google Cloud APIs. How should you set up authentication of these services in the cluster following Google-recommended best practices? (Choose two.)

- A. Use the service account attached to the GKE node.
- B. Enable Workload Identity in the cluster via the gcloud command-line tool.
- C. Access the Google service account keys from a secret management service.
- D. Store the Google service account keys in a central secret management service.
- E. Use gcloud to bind the Kubernetes service account and the Google service account using roles/iam.workloadIdentity.

**ANSWER: B E**

**Explanation:**

<https://cloud.google.com/kubernetes-engine/docs/how-to/workload-identity>

#### QUESTION NO: 10

You configured your Compute Engine instance group to scale automatically according to overall CPU usage. However, your application's response latency increases sharply before the cluster has finished adding up instances. You want to provide a more consistent latency experience for your end users by changing the configuration of the instance group autoscaler. Which two configuration changes should you make? (Choose two.)

- A. Add the label "AUTOSCALE" to the instance group template.
- B. Decrease the cool-down period for instances added to the group.
- C. Increase the target CPU usage for the instance group autoscaler.
- D. Decrease the target CPU usage for the instance group autoscaler.
- E. Remove the health-check for individual VMs in the instance group.

**ANSWER: A C**

#### QUESTION NO: 11

You are porting an existing Apache/MySQL/PHP application stack from a single machine to Google

Kubernetes Engine. You need to determine how to containerize the application. Your approach should follow Google-recommended best practices for availability. What should you do?

- A. Package each component in a separate container. Implement readiness and liveness probes.
- B. Package the application in a single container. Use a process management tool to manage each component.
- C. Package each component in a separate container. Use a script to orchestrate the launch of the components.
- D. Package the application in a single container. Use a bash script as an entrypoint to the container, and then spawn each component as a background job.

**ANSWER: D**

#### Explanation:

Reference: <https://cloud.google.com/architecture/best-practices-for-building-containers>

#### QUESTION NO: 12

You are designing an application that uses a microservices architecture. You are planning to deploy the application in the cloud and on-premises. You want to make sure the application can scale up on demand and also use managed services as much as possible. What should you do?

- A. Deploy open source Istio in a multi-cluster deployment on multiple Google Kubernetes Engine (GKE) clusters managed by Anthos.
- B. Create a GKE cluster in each environment with Anthos, and use Cloud Run for Anthos to deploy your application to each cluster.

- C. Install a GKE cluster in each environment with Anthos, and use Cloud Build to create a Deployment for your application in each cluster.
- D. Create a GKE cluster in the cloud and install open-source Kubernetes on-premises. Use an external load balancer service to distribute traffic across the two environments.

**ANSWER: B**

**Explanation:**

<https://cloud.google.com/anthos/run>

Integrated with Anthos, Cloud Run for Anthos provides a flexible serverless development platform for hybrid and multicloud environments. Cloud Run for Anthos is Google's managed and fully supported Knative offering, an open source project that enables serverless workloads on Kubernetes.

**QUESTION NO: 13**

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.

**ANSWER: A B**

**QUESTION NO: 14**

You are developing a web application that will be accessible over both HTTP and HTTPS and will run on Compute Engine instances. On occasion, you will need to SSH from your remote laptop into one of the Compute Engine instances to conduct maintenance on the app. How should you configure the instances while following Google-recommended best practices?

- A. Set up a backend with Compute Engine web server instances with a private IP address behind a TCP proxy load balancer.
- B. Configure the firewall rules to allow all ingress traffic to connect to the Compute Engine web servers, with each server having a unique external IP address.
- C. Configure Cloud Identity-Aware Proxy API for SSH access. Then configure the Compute Engine servers with private IP addresses behind an HTTP(s) load balancer for the application web traffic.

D. Set up a backend with Compute Engine web server instances with a private IP address behind an HTTP(S) load balancer. Set up a bastion host with a public IP address and open firewall ports. Connect to the web instances using the bastion host.

**ANSWER: C**

**Explanation:**

Reference: [https://cloud.google.com/compute/docs/instances/connecting-advanced#cloud\\_iap](https://cloud.google.com/compute/docs/instances/connecting-advanced#cloud_iap)

**QUESTION NO: 15**

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.

**ANSWER: C**

**Explanation:**

Reference: <https://cloud.google.com/endpoints/docs/openapi/get-started-cloud-functions>