



Free Questions for Professional-Cloud-Developer by certscare

Shared by Caldwell on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You need to deploy an internet-facing microservices application to Google Kubernetes Engine (GKE). You want to validate new features using the A/B testing method. You have the following requirements for deploying new container image releases

- * There is no downtime when new container images are deployed.
- * New production releases are tested and verified using a subset of production users.

What should you do?

Options:

- A-** 1 Configure your CI/CD pipeline to update the Deployment manifest file by replacing the container version with the latest version.
2 Recreate the Pods in your cluster by applying the Deployment manifest file.
3 Validate the application's performance by comparing its functionality with the previous release version and roll back if an issue arises.
- B-** 1 install the Anthos Service Mesh on your GKE cluster.
2 Create two Deployments on the GKE cluster and label them with different version names.
3 Create a VirtualService with a routing rule to send a small percentage of traffic to the Deployment that references the new version of the application.
- C-** 1 Create a second namespace on GKE for the new release version.
2 Create a Deployment configuration for the second namespace with the desired number of Pods.

3 Deploy new container versions in the second namespace.

4 Update the ingress configuration to route traffic to the namespace with the new container versions.

D- 1. Implement a rolling update pattern by replacing the Pods gradually with the new release version.

2 Validate the application's performance for the new subset of users during the rollout and roll back if an issue arises.

Answer:

C

Question 2

Question Type: MultipleChoice

You are developing a flower ordering application. Currently you have three microservices.

* Order Service (receives the orders).

* Order Fulfillment Service (processes the orders).

* Notification Service (notifies the customer when the order is filled).

You need to determine how the services will communicate with each other. You want incoming orders to be processed quickly and you need to collect order information for fulfillment. You also want to make sure orders are not lost between your services and are able to communicate asynchronously. How should the requests be processed?

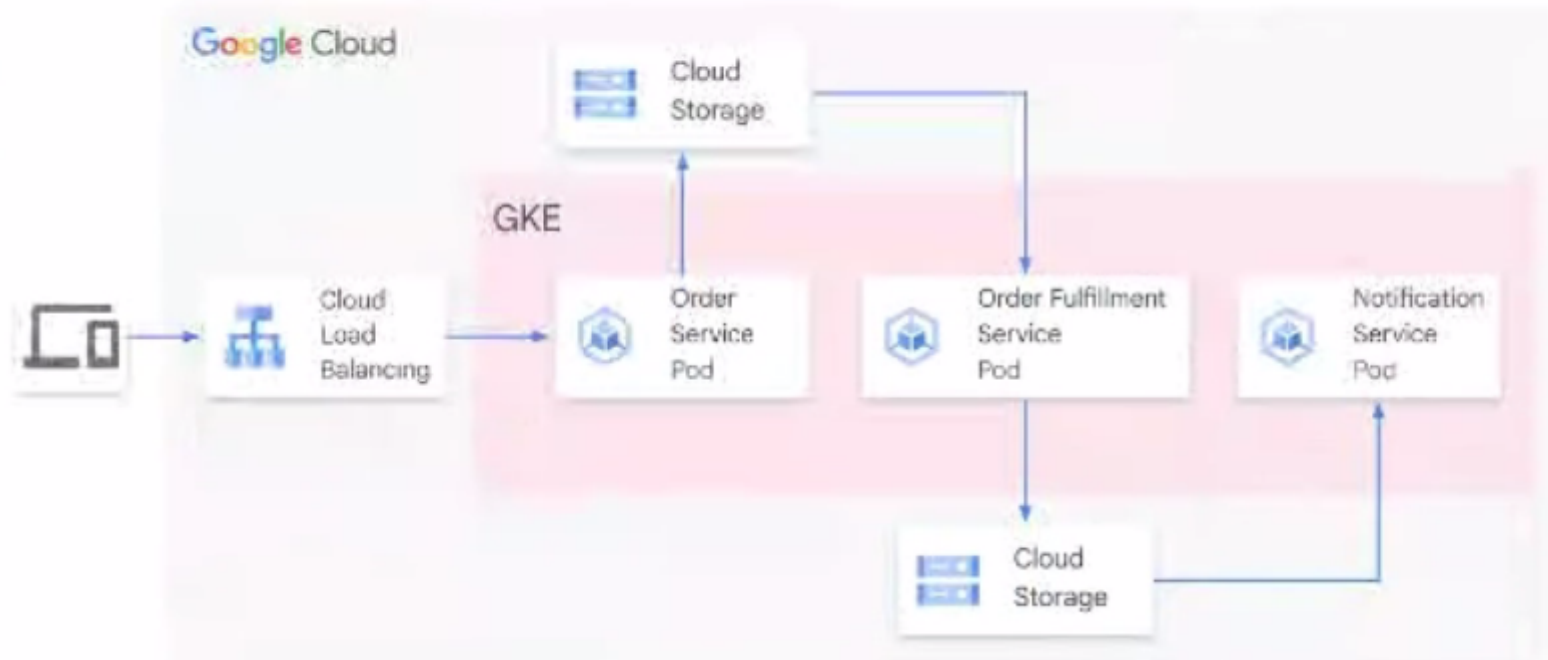
A.

Order request → Order Service → Order Fulfillment Service → Notification Service



B.

Order request → Order Service → Cloud Storage bucket → Order Fulfillment Service → Cloud Storage bucket → Notification Service



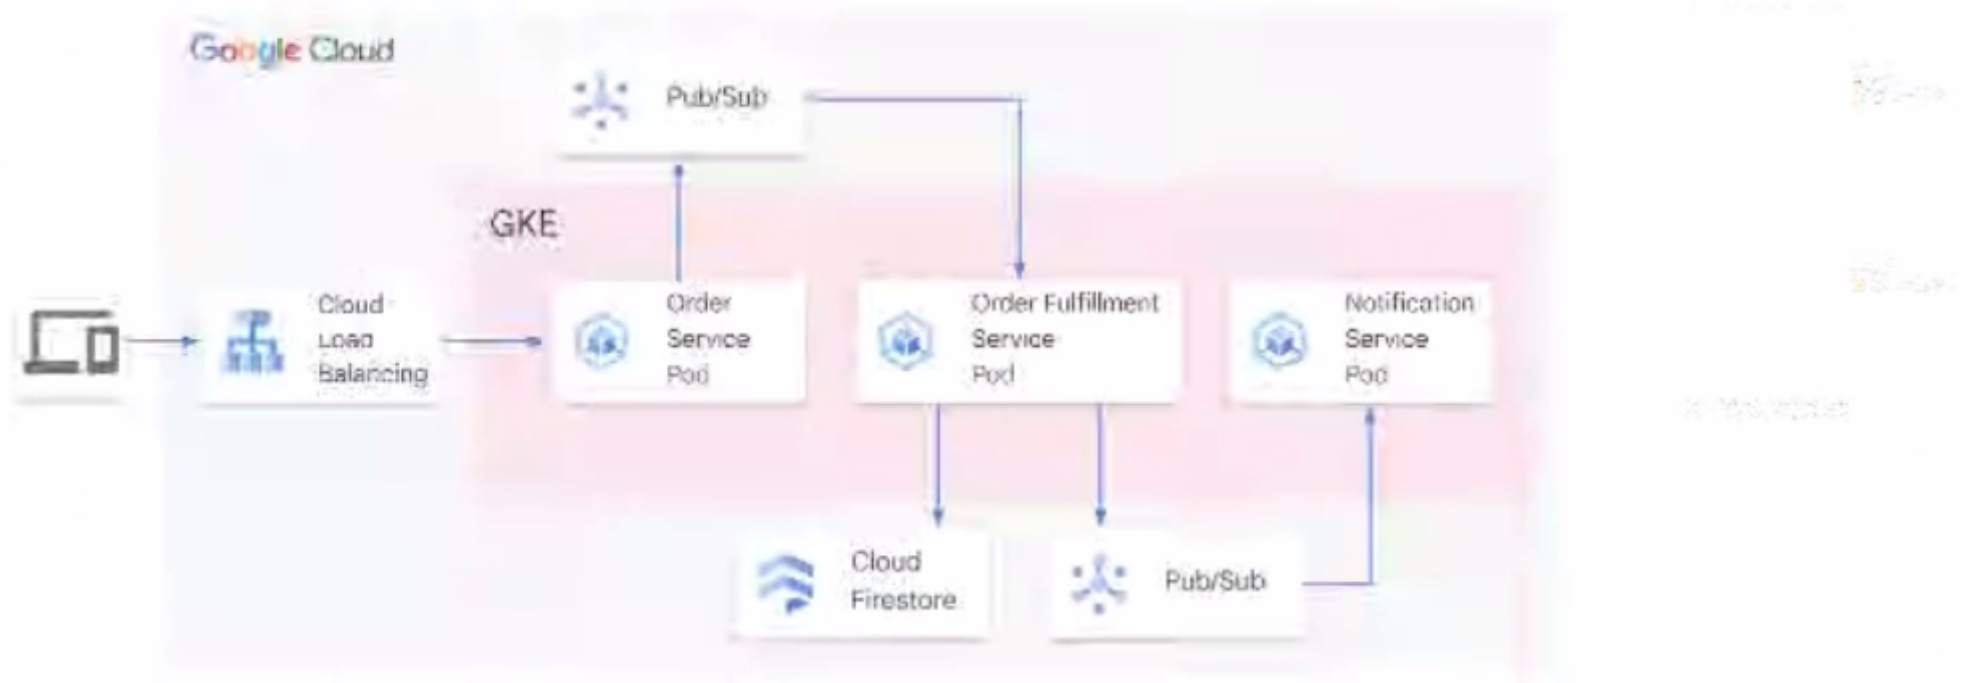
C.

Order request → Order Service → Firestore database → Order Fulfillment Service → Firestore database → Notification Service



D.

Order request → Order Service → Pub/Sub queue → Order Fulfillment Service → Firestore database → Pub/Sub queue → Notification Service



Options:

A- Option A

B- Option B

C- Option C

D- Option D

Answer:

A

Question 3

Question Type: MultipleChoice

You are developing a new application. You want the application to be triggered only when a given file is updated in your Cloud Storage bucket. Your trigger might change, so your process must support different types of triggers. You want the configuration to be simple so that multiple team members can update the triggers in the future. What should you do?

Options:

A- Create an Eventarc trigger that monitors your Cloud Storage bucket for a specific filename, and set the target as Cloud Run.

B- Configure Cloud Storage events to be sent to Pub/Sub, and use Pub/Sub events to trigger a Cloud Build job that executes your application.

- C-** Configure a Firebase function that executes your application and is triggered when an object is updated in Cloud Storage.
- D-** Configure a Cloud Function that executes your application and is triggered when an object is updated in Cloud Storage.

Answer:

C

Question 4

Question Type: MultipleChoice

Your team recently deployed an application on Google Kubernetes Engine (GKE). You are monitoring your application and want to be alerted when the average memory consumption of your containers is under 20% or above 80% How should you configure the alerts?

Options:

- A-** Create a Cloud Function that consumes the Monitoring API. Create a schedule to trigger the Cloud Function hourly and alert you if the average memory consumption is outside the defined range
- B-** In Cloud Monitoring, create an alerting policy to notify you if the average memory consumption is outside the defined range
- C-** Create a Cloud Function that runs on a schedule, executes `kubectl top` on all the workloads on the cluster, and sends an email alert

if the average memory consumption is outside the defined range

D- Write a script that pulls the memory consumption of the instance at the OS level and sends an email alert if the average memory consumption is outside the defined range

Answer:

B

Question 5

Question Type: MultipleChoice

Your team is developing a Cloud Function triggered by Cloud Storage Events. You want to accelerate testing and development of your Cloud Function while following Google-recommended best practices. What should you do?

Options:

A- Install the Functions Frameworks library, and configure the Cloud Function on localhost. Make a copy of the function, and make edits to the new version Test the new version using curl.

B- Make a copy of the Cloud Function, and rewrite the code to be HTTP-triggered Edit and test the new version by triggering the HTTP endpoint. Send mock requests to the new function to evaluate the functionality.

C- Make a copy of the Cloud Function in the Google Cloud Console Use the Cloud console's in-line editor to make source code changes to the new function Modify your web application to call the new function and test the new version in production.

D- Create a new Cloud Function that is triggered when Cloud Audit Logs detects the cloudfunctions. functions. sourceCodeSet operation in the original Cloud Function Send mock requests to the new function to evaluate the functionality.

Answer:

A

Question 6

Question Type: MultipleChoice

You have deployed a Java application to Cloud Run. Your application requires access to a database hosted on Cloud SQL Due to regulatory requirements: your connection to the Cloud SQL instance must use its internal IP address. How should you configure the connectivity while following Google-recommended best practices'?

Options:

- A- Configure your Cloud Run service with a Cloud SQL connection.
- B- Configure your Cloud Run service to use a Serverless VPC Access connector
- C- Configure your application to use the Cloud SQL Java connector
- D- Configure your application to connect to an instance of the Cloud SQL Auth proxy

Answer:

A

Question 7

Question Type: MultipleChoice

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 8

Question Type: MultipleChoice

You are developing an online gaming platform as a microservices application on Google Kubernetes Engine (GKE). Users on social media are complaining about long loading times for certain URL requests to the application. You need to investigate performance bottlenecks in the application and identify which HTTP requests have a significantly high latency span in user requests. What should you do?

Options:

- A-** Instrument your microservices by installing the OpenTelemetry tracing package Update your application code to send traces to Trace for inspection and analysis Create an analysis report on Trace to analyze user requests
- B-** Configure GKE workload metrics using kubectl Select all Pods to send their metrics to Cloud Monitoring. Create a custom dashboard of application metrics in Cloud Monitoring to determine performance bottlenecks of your GKE cluster
- C-** Install tcpdump on your GKE nodes. Run tcpdump to capture network traffic over an extended period of time to collect data Analyze the data files using Wireshark to determine the cause of high latency
- D-** Update your microservices to log HTTP request methods and URL paths to STDOUT Use the logs router to send container logs to Cloud Logging Create filters in Cloud Logging to evaluate the latency of user requests across different methods and URL paths.

Answer:

B

Question 9

Question Type: MultipleChoice

Your team is setting up a build pipeline for an application that will run in Google Kubernetes Engine (GKE). For security reasons, you only want images produced by the pipeline to be deployed to your GKE cluster. Which combination of Google Cloud services should you use?

Options:

- A- Google Cloud Deploy, Artifact Registry, and Google Cloud Armor
- B- Google Cloud Deploy, Cloud Storage and Google Cloud Armor
- C- Cloud Build, Cloud Storage, and Binary Authorization
- D- Cloud Build, Artifact Registry and Binary Authorization

Answer:

A

Question 10

Question Type: MultipleChoice

You are using Cloud Run to host a global ecommerce web application. Your company's design team is creating a new color scheme for the web app. You have been tasked with determining whether the new color scheme will increase sales. You want to conduct testing on live production traffic. How should you design the study?

Options:

- A-** Use an external HTTP(S) load balancer to route a predetermined percentage of traffic to two different color schemes of your application. Analyze the results to determine whether there is a statistically significant difference in sales.
- B-** Use an external HTTP(S) load balancer to route traffic to the original color scheme while the new deployment is created and tested. After testing is complete, reroute all traffic to the new color scheme. Analyze the results to determine whether there is a statistically significant difference in sales.
- C-** Enable a feature flag that displays the new color scheme to half of all users. Monitor sales to see whether they increase for this group of users.
- D-** Use an external HTTP(S) load balancer to mirror traffic to the new version of your application. Analyze the results to determine whether there is a statistically significant difference in sales.

Answer:

A

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>

