Google Professional-Cloud-DevOps-Engineer Practice Questions

Shared by Gallagher on 24-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Your company wants to implement a CD pipeline in Cloud Deploy for a web service deployed to GKE. The web service currently does not have any automated testing. The Quality Assurance team must manually verify any new releases of the web service before any production traffic is processed. You need to design the CD pipeline. What should you do?

Options:

A- Create two pipeline stages, and use a canary deployment strategy.
B- Create a single pipeline stage, and use a standard deployment strategy.
C- Create a single pipeline stage, and use a canary deployment strategy.
D- Create two pipeline stages, and use a standard deployment strategy.

Answer:

D

Explanation:

To allow manual QA approval, you need separate pipeline stages (e.g., staging production) and manual promotion between them.

'Cloud Deploy supports multi-stage pipelines where you can manually approve releases before promotion.'

--- Cloud Deploy Promotion

'Use a standard strategy when there are no automated tests or SLOs to support rollback decisions.'

--- Deployment Strategies

Canary is best used when automated health checks and metrics exist, which is not the case here.

Question 2

Question Type: MultipleChoice

[Building and implementing CI/CD pipelines for a service]

Your company runs applications in Google Kubernetes Engine (GKE) that are deployed following a GitOps methodology.

Application developers frequently create cloud resources to support their applications. You want to give developers the ability to manage infrastructure as code, while ensuring that you follow Google-recommended practices. You need to ensure that infrastructure as code reconciles periodically to avoid configuration drift. What should you do?

Options:

A- Install and configure Config Connector in Google Kubernetes Engine (GKE).
B- Configure Cloud Build with a Terraform builder to execute plan and apply commands.
C- Create a Pod resource with a Terraform docker image to execute terraform plan and terraform apply commands.
D- Create a Job resource with a Terraform docker image to execute terraforrm plan and terraform apply commands.

Answer:

A

Question 3

Question Type: MultipleChoice

[Applying Site Reliability Engineering (SRE) principles to a service]

You manage your company's main revenue-generating application. You have an error budget policy in place that freezes production deployments when the application is close to breaching its SLO. A number of issues have recently occurred, and the application has exhausted its error budget. You need to deploy a new release to the application that includes a feature urgently required by your largest customer. You have been told that the release has passed all unit tests. What should you do?

Options:

A- Start the deployment of the feature immediately.
B- Delay the deployment of the feature until the error budget is replenished.
C- Re-run the unit tests, and start the deployment of the feature if the tests pass.
D- Deploy the feature to a subset of users, and gradually roll out to all users if there are no errors reported.

Answer:

D

Question 4

Question Type: MultipleChoice

[Building and implementing CI/CD pipelines for a service]

Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat-images project can be deployed to the cluster while minimizing development time. What should you do?

Options:

A- Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.
B- Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/attostrat-images/.
C- Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.
D- Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.

Answer:

B

Question 5

Question Type: MultipleChoice

[Building and implementing CI/CD pipelines for a service]

You are building and deploying a microservice on Cloud Run for your organization Your service is used by many applications internally You are deploying a new release, and you need to test the new version extensively in the staging and production environments You must minimize user and developer impact. What should you do?

Options:

A- Deploy the new version of the service to the staging environment Split the traffic, and allow 1 % of traffic through to the latest version Test the latest version If the test passes gradually roll out the latest version to the staging and production environments
B- Deploy the new version of the service to the staging environment Split the traffic, and allow 50% of traffic through to the latest version Test the latest version If the test passes, send all traffic to the latest version Repeat for the production environment
C- Deploy the new version of the service to the staging environment with a new-release tag without serving traffic Test the new-release version If the test passes; gradually roll out this tagged version Repeat for the production environment
D- Deploy a new environment with the green tag to use as the staging environment Deploy the new version of the service to the green environment and test the new version If the tests pass, send all traffic to the green environment and delete the existing staging environment Repeat for the production environment

Answer:

C

Question 6

Question Type: MultipleChoice

[Building and implementing service monitoring strategies]

You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

Options:

A- * Deploy the Stackdriver logging agent to the application servers.* Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
B- * Deploy the Stackdriver logging agent to the application servers.* Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
C- * Deploy the Stackdriver monitoring agent to the application servers.* Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
D- * Install the gsutil command line tool on your application servers.* Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes.* Give the developers IAM Object Viewer access to view the logs in the specified bucket.

Answer:

A

Question 7

Question Type: MultipleChoice

[Building and implementing service monitoring strategies]

Your organization is running multiple Google Kubernetes Engine (GKE) clusters in a project. You need to design a highly-available solution to collect and query both domain-specific workload metrics and GKE default metrics across all clusters, while minimizing operational overhead. What should you do?

Options:

A- Use Prometheus Operator to install Prometheus in every cluster and scrape the metrics. Ensure that a Thanos sidecar is enabled on every Prometheus instance. Configure Thanos in the central cluster. Query the central Thanos instance.
B- Use Prometheus Operator to install Prometheus in every cluster and scrape the metrics. Configure remote-write to one central Prometheus. Query the central Prometheus instance.
C- Enable managed collection on every GKE cluster. Query the metrics in Cloud Monitoring.
D- Enable managed collection on every GKE cluster. Query the metrics in BigQuery.

Answer:

C

Question 8

Question Type: MultipleChoice

[Building and implementing CI/CD pipelines for a service]

You are configuring the frontend tier of an application deployed in Google Cloud The frontend tier is hosted in ngmx and deployed using a managed instance group with an Envoy-based external HTTP(S) load balancer in front The application is deployed entirely within the europe-west2 region: and only serves users based in the United Kingdom. You need to Select the most cost-effective network tier and load balancing configuration What should you use?

Options:

A- Premium Tier with a global load balancer
B- Premium Tier with a regional load balancer
C- Standard Tier with a global load balancer
D- Standard Tier with a regional load balancer

Answer:

B

Question 9

Question Type: MultipleChoice

[Applying Site Reliability Engineering (SRE) principles to a service]

You have an application deployed to Cloud Run. A new version of the application has recently been deployed using the canary deployment strategy. Your Site Reliability Engineering (SRE) teammate informs you that an SLO has been exceeded for this application. You need to make the application healthy as quickly as possible. What should you do first?

Options:

A- Configure traffic splitting to send 100% of the traffic to the latest revision.
B- Configure traffic splitting to send 100% of the traffic to the previous revision.
C- Create a new revision using the last known good version of the application.
D- Identify the cause of the latency by using Cloud Trace.

Answer:

B

Question 10

Question Type: MultipleChoice

[Building and implementing CI/CD pipelines for a service]

You manage an application that runs in Google Kubernetes Engine (GKE) and uses the blue/green deployment methodology Extracts of the Kubernetes manifests are shown below:

```
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: app-green
  labels:
    app: my-app
    version: green
<other fields snipped>
---
apiVersion: v1
kind: Service
metadata:
  name: app-svc
spec:
  selector:
    app: my-app
    version: green
<other fields snipped>
---
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: app-ingress
spec:
  defaultBackend:
    service:
      name: app-svc
<other fields snipped>
```

The Deployment app-green was updated to use the new version of the application During post-deployment monitoring you notice that the majority of user requests are failing You did not observe this behavior in the testing environment You need to mitigate the incident impact on

users and enable the developers to troubleshoot the issue What should you do?

Options:

A- Update the Deployment app-blue to use the new version of the application
B- Update the Deployment ape-green to use the previous version of the application
C- Change the selector on the Service app-2vc to app: my-app.
D- Change the selector on the Service app-svc to app: my-app, version: blue

Answer:

C

P2P exams

Question 11

Question Type: MultipleChoice

[Building and implementing service monitoring strategies]

Your company is creating a new cloud-native Google Cloud organization. You expect this Google Cloud organization to first be used by a small number of departments and then expand to be used by a large number of departments. Each department has a large number of applications varying in size. You need to design the VPC network architecture. Your solution must minimize the amount of management required while remaining flexible enough for development teams to quickly adapt to their evolving needs. What should you do?

Options:

A- Create a separate VPC for each department and connect the VPCs with VPC Network Peering.
B- Create a separate VPC for each department and use Private Service Connect to connect the VPCs.
C- Create a separate VPC for each application and use Private Service Connect to connect the VPCs.
D- Create a separate VPC for each department and connect the VPCs with Cloud VPN.

Answer:

A

To Get Premium Files for Professional-Cloud-DevOps-Engineer Visit

https://www.p2pexams.com/products/professional-cloud-devops-engineer

For More Free Questions Visit

https://www.p2pexams.com/google/pdf/professional-cloud-devops-engineer