



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-devops-engineer.html>
2024 Latest pass4itsure PROFESSIONAL-CLOUD-DEVOPS-ENGINEER PDF
and VCE dumps Download

PROFESSIONAL-CLOUD-DEVOPS-ENGINEER^{Q&As}

Professional Cloud DevOps Engineer

Pass Google PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/professional-cloud-devops-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-devops-engineer.html>
2024 Latest pass4itsure PROFESSIONAL-CLOUD-DEVOPS-ENGINEER PDF
and VCE dumps Download

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (OL) and Communications Lead (CL). What should you do next?

- A. Look for ways to mitigate user impact and deploy the mitigations to production.
- B. Contact the affected service owners and update them on the status of the incident.
- C. Establish a communication channel where incident responders and leads can communicate with each other.
- D. Start a postmortem, add incident information, circulate the draft internally, and ask internal stakeholders for input.

Correct Answer: C

<https://sre.google/workbook/incident-response/>

QUESTION 2

You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how much it costs to run each of the systems. What should you do?

- A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
- B. Assign all instances a label specific to the system they run. Configure BigQuery billing export and query costs per label.
- C. Enrich all instances with metadata specific to the system they run. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
- D. Name each virtual machine (VM) after the system it runs. Set up a usage report export to a Cloud Storage bucket. Configure the bucket as a source in BigQuery to query costs based on VM name.

Correct Answer: B

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

QUESTION 3

You are creating a CI/CD pipeline in Cloud Build to build an application container image. The application code is stored in GitHub. Your company requires that production image builds are only run against the main branch and that the change control team approves all pushes to the main branch. You want the image build to be as automated as possible. What should you do? (Choose two.)

- A. Create a trigger on the Cloud Build job. Set the repository event setting to 'Pull request'.
- B. Add the OWNERS file to the Included files filter on the trigger.



- C. Create a trigger on the Cloud Build job. Set the repository event setting to 'Push to a branch'
- D. Configure a branch protection rule for the main branch on the repository.
- E. Enable the Approval option on the trigger.

Correct Answer: CD

Option C: Setting the repository event setting to 'Push to a branch' will trigger the Cloud Build job whenever a push is made to any branch in the repository. This is necessary because you want the image build to be triggered when a push is

made to the main branch.

Option D: Configuring a branch protection rule for the main branch on the repository will require that all pushes to the main branch be approved by the change control team. This is necessary to ensure that only approved changes are made to

the main branch, which will then trigger the image build.

QUESTION 4

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

- A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.
- D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Correct Answer: D

<https://cloud.google.com/logging/docs/routing/overview>

QUESTION 5

You are building the CI/CD pipeline for an application deployed to Google Kubernetes Engine (GKE). The application is deployed by using a Kubernetes Deployment, Service, and Ingress. The application team asked you to deploy the application by using the blue/green deployment methodology. You need to implement the rollback actions. What should you do?

- A. Run the kubectl rollout undo command.
- B. Delete the new container image, and delete the running Pods.
- C. Update the Kubernetes Service to point to the previous Kubernetes Deployment.
- D. Scale the new Kubernetes Deployment to zero.



Correct Answer: C

The point of using blue/green deployment is to have both blue and green versions deployed, and we chose the active one at service level, so rollback should only consist of updating the service to serve the other version.

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER PDF Dumps](#)

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER VCE Dumps](#)

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Practice Test](#)