



VCE & PDF

PassApply.com

<https://www.passapply.com/professional-cloud-devops-engineer.html>
2024 Latest passapply 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.passapply.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

PassApply.com

<https://www.passapply.com/professional-cloud-devops-engineer.html>
2024 Latest passapply 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 support a web application that is hosted on Compute Engine. The application provides a booking service for thousands of users. Shortly after the release of a new feature, your monitoring dashboard shows that all users are experiencing latency at login. You want to mitigate the impact of the incident on the users of your service. What should you do first?

- A. Roll back the recent release.
- B. Review the Stackdriver monitoring.
- C. Upsize the virtual machines running the login services.
- D. Deploy a new release to see whether it fixes the problem.

Correct Answer: A

Roll back the recent release. This would be the quickest way to remove the new feature that is causing the latency and restore the application to its previous state. This would immediately mitigate the impact on users, while you continue to investigate the issue with the new feature and identify a long-term solution.

QUESTION 2

You are using Terraform to manage infrastructure as code within a CI/CD pipeline. You notice that multiple copies of the entire infrastructure stack exist in your Google Cloud project, and a new copy is created each time a change to the existing infrastructure is made. You need to optimize your cloud spend by ensuring that only a single instance of your infrastructure stack exists at a time. You want to follow Google-recommended practices. What should you do?

- A. Create a new pipeline to delete old infrastructure stacks when they are no longer needed.
- B. Confirm that the pipeline is storing and retrieving the terraform.tfstate file from Cloud Storage with the Terraform gcs backend.
- C. Verify that the pipeline is storing and retrieving the terraform.tfstate file from a source control.
- D. Update the pipeline to remove any existing infrastructure before you apply the latest configuration.

Correct Answer: B

QUESTION 3

You are developing reusable infrastructure as code modules. Each module contains integration tests that launch the module in a test project. You are using GitHub for source control. You need to continuously test your feature branch and ensure that all code is tested before changes are accepted. You need to implement a solution to automate the integration tests. What should you do?

- A. Use a Jenkins server for CI/CD pipelines. Periodically run all tests in the feature branch.
- B. Ask the pull request reviewers to run the integration tests before approving the code.
- C. Use Cloud Build to run the tests. Trigger all tests to run after a pull request is merged.



D. Use Cloud Build to run tests in a specific folder. Trigger Cloud Build for every GitHub pull request.

Correct Answer: C

<https://cloud.google.com/build/docs/automating-builds/create-manage-triggers>

QUESTION 4

You use a multiple step Cloud Build pipeline to build and deploy your application to Google Kubernetes Engine (GKE). You want to integrate with a third-party monitoring platform by performing a HTTP POST of the build information to a webhook. You want to minimize the development effort. What should you do?

- A. Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- B. Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- C. Use Stackdriver Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.
- D. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

Correct Answer: D

<https://cloud.google.com/pubsub/docs/push> <https://cloud.google.com/build/docs/subscribe-build-notifications>

QUESTION 5

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Stackdriver Logging. What should you do?

- A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.
- B. Deploy a Fluentd daemonset to GKE. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- C. Install Kubernetes on Google Compute Engine (GCE) and redeploy your applications. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Write a script to tail the log file within the pod and write entries to standard output. Run the script as a sidecar container with the application's pod. Configure a shared volume between the containers to allow the script to have read access to `/var/log` in the application container.

Correct Answer: B

Reference: <https://cloud.google.com/solutions/customizing-stackdriver-logs-fluentd>

QUESTION 6



Your Cloud Run application writes unstructured logs as text strings to Cloud Logging. You want to convert the unstructured logs to JSON-based structured logs. What should you do?

- A. Modify the application to use Cloud Logging software development kit (SDK), and send log entries with a jsonPayload field.
- B. Install a Fluent Bit sidecar container, and use a JSON parser.
- C. Install the log agent in the Cloud Run container image, and use the log agent to forward logs to Cloud Logging.
- D. Configure the log agent to convert log text payload to JSON payload.

Correct Answer: D

<https://cloud.google.com/logging/docs/agent/logging/configuration#process-payload>

QUESTION 7

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 8

Your team has recently deployed an NGINX-based application into Google Kubernetes Engine (GKE) and has exposed it to the public via an HTTP Google Cloud Load Balancer (GCLB) ingress. You want to scale the deployment of the application's frontend using an appropriate Service Level Indicator (SLI). What should you do?

- A. Configure the horizontal pod autoscaler to use the average response time from the Liveness and Readiness probes.
- B. Configure the vertical pod autoscaler in GKE and enable the cluster autoscaler to scale the cluster as pods expand.
- C. Install the Stackdriver custom metrics adapter and configure a horizontal pod autoscaler to use the number of requests provided by the GCLB.
- D. Expose the NGINX stats endpoint and configure the horizontal pod autoscaler to use the request metrics exposed by the NGINX deployment.

Correct Answer: C



<https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics>

QUESTION 9

You recently migrated an ecommerce application to Google Cloud. You now need to prepare the application for the upcoming peak traffic season. You want to follow Google-recommended practices. What should you do first to prepare for the busy season?

- A. Migrate the application to Cloud Run, and use autoscaling.
- B. Create a Terraform configuration for the application's underlying infrastructure to quickly deploy to additional regions.
- C. Load test the application to profile its performance for scaling.
- D. Pre-provision the additional compute power that was used last season, and expect growth.

Correct Answer: C

QUESTION 10

You are the Site Reliability Engineer responsible for managing your company's data services and products. You regularly navigate operational challenges, such as unpredictable data volume and high cost, with your company's data ingestion processes. You recently learned that a new data ingestion product will be developed in Google Cloud. You need to collaborate with the product development team to provide operational input on the new product. What should you do?

- A. Deploy the prototype product in a test environment, run a load test, and share the results with the product development team.
- B. When the initial product version passes the quality assurance phase and compliance assessments, deploy the product to a staging environment. Share error logs and performance metrics with the product development team.
- C. When the new product is used by at least one internal customer in production, share error logs and monitoring metrics with the product development team.
- D. Review the design of the product with the product development team to provide feedback early in the design phase.

Correct Answer: D

QUESTION 11

Your team deploys applications to three Google Kubernetes Engine (GKE) environments: development, staging, and production. You use GitHub repositories as your source of truth. You need to ensure that the three environments are consistent. You want to follow Google-recommended practices to enforce and install network policies and a logging DaemonSet on all the GKE clusters in those environments. What should you do?

- A. Use Google Cloud Deploy to deploy the network policies and the DaemonSet. Use Cloud Monitoring to trigger an alert if the network policies and DaemonSet drift from your source in the repository.
- B. Use Google Cloud Deploy to deploy the DaemonSet and use Policy Controller to configure the network policies. Use



Cloud Monitoring to detect drifts from the source in the repository and Cloud Functions to correct the drifts.

C. Use Cloud Build to render and deploy the network policies and the DaemonSet. Set up Config Sync to sync the configurations for the three environments.

D. Use Cloud Build to render and deploy the network policies and the DaemonSet. Set up a Policy Controller to enforce the configurations for the three environments.

Correct Answer: C

QUESTION 12

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the 90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

- A. 90th percentile – 100ms 95th percentile – 250ms
- B. 90th percentile – 120ms 95th percentile – 275ms
- C. 90th percentile – 150ms 95th percentile – 300ms
- D. 90th percentile – 250ms 95th percentile – 400ms

Correct Answer: C

<https://sre.google/sre-book/service-level-objectives/>

QUESTION 13

A third-party application needs to have a service account key to work properly. When you try to export the key from your cloud project, you receive an error: “The organization policy constraint iam.disableServiceAccountKeyCreation is enforced.” You need to make the third-party application work while following Google-recommended security practices.

What should you do?

- A. Enable the default service account key, and download the key.
- B. Remove the iam.disableServiceAccountKeyCreation policy at the organization level, and create a key.
- C. Disable the service account key creation policy at the project's folder, and download the default key.
- D. Add a rule to set the iam.disableServiceAccountKeyCreation policy to off in your project, and create a key.

Correct Answer: B

QUESTION 14

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use



Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Correct Answer: D

Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

This option allows you to compare the previous and current Cloud Build Configuration files and find what has been changed and understand what could be causing the issue. After identifying the problem, you can fix it and ensure that the

pipeline is working correctly again. This approach is based on the SRE practices of identifying and resolving issues quickly and effectively.

QUESTION 15

You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.

You need to troubleshoot the issue. What should you do?

- A. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
- B. Confirm that your account has the proper permissions to use the Stackdriver dashboard.
- C. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
- D. Confirm that the application is using the required client library and the service account key has proper permissions.

Correct Answer: A

Using a version-control system such as Git, allows you to store different versions of the code and track changes that are made to it. This allows you to easily roll back to a previous version if necessary, and ensures that all team members are working on the same version of the code. It also allows team members to collaborate on the code and merge changes with the master version. This is the most effective way to share code and maintain versioning, as it provides a centralized location for all code and allows for easy collaboration and management of code changes.

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

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

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Exam Questions](#)