



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 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.

QUESTION 2

You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

- A. As the I/O wait times aggregated across all report generation backends
- B. As the proportion of report generation requests that result in a successful response
- C. As the application's report generation queue size compared to a known-good threshold
- D. As the reporting backend PD throughput capacity compared to a known-good threshold

Correct Answer: B

As the proportion of report generation requests that result in a successful response is a valid availability Service Level Indicator (SLI) for the report generation feature. This indicator measures the percentage of requests for report generation that are successfully completed, and it is an indicator of the service availability. This SLI provides a clear and measurable way to track the availability of the report generation feature. It is a simple and easy to understand metric, that it can be easily monitored and reported. It can be used to track the performance of the report generation feature over time and detect any potential issues that may cause it to become unavailable. It also allows you to detect and diagnose issues in the system quickly and take appropriate action to mitigate them. Additionally, it aligns well with the customer's expectation of the report generation feature as they want to see a high percentage of successful report generation requests, which indicates that the feature is working correctly.

QUESTION 3



You support a high-traffic web application with a microservice architecture. The home page of the application displays multiple widgets containing content such as the current weather, stock prices, and news headlines. The main serving thread makes a call to a dedicated microservice for each widget and then lays out the homepage for the user. The microservices occasionally fail; when that happens, the serving thread serves the homepage with some missing content. Users of the application are unhappy if this degraded mode occurs too frequently, but they would rather have some content served instead of no content at all. You want to set a Service Level Objective (SLO) to ensure that the user experience does not degrade too much. What Service Level Indicator (SLI) should you use to measure this?

- A. A quality SLI: the ratio of non-degraded responses to total responses.
- B. An availability SLI: the ratio of healthy microservices to the total number of microservices.
- C. A freshness SLI: the proportion of widgets that have been updated within the last 10 minutes.
- D. A latency SLI: the ratio of microservice calls that complete in under 100 ms to the total number of microservice calls.

Correct Answer: A

<https://cloud.google.com/architecture/adopting-slos>

QUESTION 4

You are building an application that runs on Cloud Run. The application needs to access a third-party API by using an API key. You need to determine a secure way to store and use the API key in your application by following Google-recommended practices. What should you do?

- A. Save the API key in Secret Manager as a secret. Reference the secret as an environment variable in the Cloud Run application.
- B. Save the API key in Secret Manager as a secret key. Mount the secret key under the `/sys/api_key` directory, and decrypt the key in the Cloud Run application.
- C. Save the API key in Cloud Key Management Service (Cloud KMS) as a key. Reference the key as an environment variable in the Cloud Run application.
- D. Encrypt the API key by using Cloud Key Management Service (Cloud KMS), and pass the key to Cloud Run as an environment variable. Decrypt and use the key in Cloud Run.

Correct Answer: A

QUESTION 5

Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

- A. Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
- B. Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.
- C. Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by



Stackdriver.

D. Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.

Correct Answer: A

<https://cloud.google.com/profiler/docs/about-profiler>

Cloud Profiler is a statistical, low-overhead profiler that continuously gathers CPU usage and memory-allocation information from your production applications.

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

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

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Study Guide](#)