



VCE & PDF

GeekCert.com

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

GeekCert.com

<https://www.geekcert.com/professional-cloud-devops-engineer.html>
2024 Latest geekcert 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 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 2

You need to create a Cloud Monitoring SLO for a service that will be published soon. You want to verify that requests to the service will be addressed in fewer than 300 ms at least 90% of the time per calendar month. You need to identify the metric and evaluation method to use. What should you do?

- A. Select a latency metric for a request-based method of evaluation.
- B. Select a latency metric for a window-based method of evaluation.
- C. Select an availability metric for a request-based method of evaluation.
- D. Select an availability metric for a window-based method of evaluation.

Correct Answer: D

QUESTION 3

You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week. If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability. What should you do before



a major incident occurs?

- A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
- B. Negotiate with the product team to always prioritize service reliability over releasing new features.
- C. Negotiate with the development team to reduce the release frequency to no more than once a week.
- D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Correct Answer: A

Incident has not occurred yet, even when development team is already pushing new features multiple times a week.

The option A says, to define an error budget "policy", not to define error budget (It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.

QUESTION 4

You work for a global organization and run a service with an availability target of 99% with limited engineering resources.

For the current calendar month, you noticed that the service has 99.5% availability. You must ensure that your service meets the defined availability goals and can react to business changes, including the upcoming launch of new features.

You also need to reduce technical debt while minimizing operational costs. You want to follow Google-recommended practices. What should you do?

- A. Add N+1 redundancy to your service by adding additional compute resources to the service.
- B. Identify, measure, and eliminate toil by automating repetitive tasks.
- C. Define an error budget for your service level availability and minimize the remaining error budget.
- D. Allocate available engineers to the feature backlog while you ensure that the service remains within the availability target.

Correct Answer: C

QUESTION 5

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>

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

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

[PROFESSIONAL-CLOUD-DEVOPS-ENGINEER Braindumps](#)