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

QUESTION 1

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 2

Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk: Mean Time to Detect (MTTD) in minutes Mean Time to Repair (MTTR) in minutes Mean Time Between Failure (MTBF) in days User Impact Percentage

The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

A. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 33%

B. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 33%

C. MTTD: 5 MTTR: 10 MTBF: 90 Impact: 50%

D. MTTD: 5 MTTR: 20 MTBF: 90 Impact: 50%

Correct Answer: B

QUESTION 3

You currently store the virtual machine (VM) utilization logs in Stackdriver. You need to provide an easy-to-share interactive VM utilization dashboard that is updated in real time and contains information aggregated on a quarterly

basis. You want to use Google Cloud Platform solutions. What should you do?

A. 1. Export VM utilization logs from Stackdriver to BigQuery.

2.

Create a dashboard in Data Studio.

3.

Share the dashboard with your stakeholders.

B. 1. Export VM utilization logs from Stackdriver to Cloud Pub/Sub.

2.

From Cloud Pub/Sub, send the logs to a Security Information and Event Management (SIEM) system.

3.

Build the dashboards in the SIEM system and share with your stakeholders.

C. 1. Export VM utilization logs from Stackdriver to BigQuery.

2.

From BigQuery, export the logs to a CSV file.

3.

Import the CSV file into Google Sheets.

4.

Build a dashboard in Google Sheets and share it with your stakeholders.

D. 1. Export VM utilization logs from Stackdriver to a Cloud Storage bucket.

2.

Enable the Cloud Storage API to pull the logs programmatically.

3.

Build a custom data visualization application.

4.

Display the pulled logs in a custom dashboard.

Correct Answer: A

QUESTION 4

You built a serverless application by using Cloud Run and deployed the application to your production environment. You

want to identify the resource utilization of the application for cost optimization. What should you do?

A. Use Cloud Trace with distributed tracing to monitor the resource utilization of the application.

B. Use Cloud Profiler with Ops Agent to monitor the CPU and memory utilization of the application.

C. Use Cloud Monitoring to monitor the container CPU and memory utilization of the application.

D. Use Cloud Ops to create logs-based metrics to monitor the resource utilization of the application.

Correct Answer: C

QUESTION 5

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods.

What should you do?

A. Use a partitioned rolling update.

B. Use Node taints with NoExecute.

C. Use a replica set in the deployment specification.

D. Use a stateful set with parallel pod management policy.

Correct Answer: A