DUMPSOARENA

Professional Cloud DevOps Engineer

Google Google-Professional-Cloud-DevOps-Engineer

Version Demo

Total Demo Questions: 10

Total Premium Questions: 81

Buy Premium PDF

https://dumpsarena.com

sales@dumpsarena.com

dumpsarena.com

QUESTION NO: 1

You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem?

(Choose two.)

A. An explanation of the root cause of the incident.

B. A list of employees responsible for causing the incident

C. A list of action items to prevent a recurrence of the incident

D. Your opinion of the incident's severity compared to past incidents

E. Copies of the design documents for all the services impacted by the incident

ANSWER: A B

Explanation:

Reference: https://cloud.google.com/blog/products/gcp/fearless-shared-postmortems-cre-life-lessons

QUESTION NO: 2

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.

ANSWER: B

Explanation:

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

QUESTION NO: 3

You support a high-traffic web application that runs on Google Cloud Platform (GCP). You need to measure application reliability from a user perspective without making any engineering changes to it.

What should you do? (Choose two.)

A. Review current application metrics and add new ones as needed.

B. Modify the code to capture additional information for user interaction.

C. Analyze the web proxy logs only and capture response time of each request.

D. Create new synthetic clients to simulate a user journey using the application.

E. Use current and historic Request Logs to trace customer interaction with the application.

ANSWER: B D

QUESTION NO: 4

You deploy a new release of an internal application during a weekend maintenance window when there is minimal user tragic. After the window ends, you learn that one of the new features isn't working as expected in the production environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to recovery so you can avoid extended outages in the future. What should you do? (Choose two.)

A. Before merging new code, require 2 different peers to review the code changes.

B. Adopt the blue/green deployment strategy when releasing new code via a CD server.

C. Integrate a code linting tool to validate coding standards before any code is accepted into the repository.

D. Require developers to run automated integration tests on their local development environments before release.

E. Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

ANSWER: A C

QUESTION NO: 5

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP/S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN), which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration, CDN failure, or other global networking catastrophe. Where should you measure this new SLI? (Choose two.)

A. Your application servers' logs.

B. Instrumentation coded directly in the client.

C. Metrics exported from the application servers.

D. GKE health checks for your application servers.

E. A synthetic client that periodically sends simulated user requests.

ANSWER: C D

QUESTION NO: 6

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application has never consumed more than 5% of its error budget in any given time window. You hold a Service Level Objective (SLO) review with business stakeholders and confirm that the SLO is set appropriately. You want your application's SLO to more closely reflect its observed reliability. What steps can you take to further that goal while balancing velocity, reliability, and business needs? (Choose two.)

A. Add more serving capacity to all of your application's zones.

B. Have more frequent or potentially risky application releases.

C. Tighten the SLO match the application's observed reliability.

D. Implement and measure additional Service Level Indicators (SLIs) fro the application.

E. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

ANSWER: A D

QUESTION NO: 7

Your organization recently adopted a container-based workflow for application development. Your team develops numerous applications that are deployed continuously through an automated build pipeline to a Kubernetes cluster in the production environment. The security auditor is concerned that developers or operators could circumvent automated testing and push code changes to production without approval. What should you do to enforce approvals?

A. Configure the build system with protected branches that require pull request approval.

B. Use an Admission Controller to verify that incoming requests originate from approved sources.

C. Leverage Kubernetes Role-Based Access Control (RBAC) to restrict access to only approved users.

D. Enable binary authorization inside the Kubernetes cluster and configure the build pipeline as an attestor.

ANSWER: C

QUESTION NO: 8

Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes. Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to toad test new features. What should you do?

A. Create an automated testing script in production to detect failures as soon as they occur.

B. Create a development environment with smaller server capacity and give access only to developers and testers.

C. Secure the production environment to ensure that developers can't change it and set up one controlled update per year.

D. Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

ANSWER: A

QUESTION NO: 9

You are managing an application that exposes an HTTP endpoint without using a load balancer. The latency of the HTTP responses is important for the user experience. You want to understand what HTTP latencies all of your users are experiencing. You use Stackdriver Monitoring. What should you do?

A. • In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE.
• In Stackdriver's Metrics Explorer, use a Stacked Bar graph to visualize the metric.

B. • In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE.
• In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.

C. • In your application, create a metric with a metricKind set to GAUGE and a valueType set to DISTRIBUTION.
• In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.

D. • In your application, create a metric with a metricKind set to METRIC_KIND_UNSPECIFIED and a valueType set to INT64.
• In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

ANSWER: A

QUESTION NO: 10

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.

ANSWER: A

Explanation:

Reference: https://cloud.google.com/kubernetes-engine/docs/how-to/updating-apps