

Google Cloud DevOps Engineer

1. You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs.

What should you do?

A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.

B. Configure Access Context Manager to allow only these members to export logs.

C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.

D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Answer(s): A

2. You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster. The application is serving HTTP requests from users. You need to implement a solution that will reduce the network cost.

What should you do?

A. Configure the VPC as a Shared VPC Host project.

B. Configure your network services on the Standard Tier.

C. Configure your Kubernetes cluster as a Private Cluster.

D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Answer(s): B

3. 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.

Answer(s): C

4. 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 load 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(s): D

5. Your company follows Site Reliability Engineering practices. You are the Incident Commander for a new, customer-impacting incident. You need to immediately assign two incident management roles to assist you in an effective incident response.

What roles should you assign? Choose 2 answers

A. Operations Lead

B. Engineering Lead

C. Communications Lead

D. Customer Impact Assessor

E. External Customer Communications Lead

Answer(s): A C

6. Your company follows Site Reliability Engineering principles. You are writing a postmortem for an incident, triggered by a software change, that severely affected users. You want to prevent severe incidents from happening in the future.

What should you do?

A. Identify engineers responsible for the incident and escalate to their senior management.

B. Ensure that test cases that catch errors of this type are run successfully before new software releases.

C. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.

D. Design a policy that will require on-call teams to immediately call engineers and management to discuss a plan of action if an incident occurs.

Answer(s): B

7. 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.

Answer(s): A

8. Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort.

What should you do?

A. Use Cloud Build to trigger a Spinnaker pipeline.

B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.

C. Use a custom builder in Cloud Build to trigger a Jenkins pipeline.

D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine (GKE).

Answer(s): B

9. 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 2 answers

- 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(s): D E

10. You support an application deployed on Compute Engine. The application connects to a Cloud SQL instance to store and retrieve data

- A. After an update to the application, users report errors showing database timeout messages. The number of concurrent active users remained stable. You need to find the most probable cause of the database timeout. What should you do?
- B. Check the serial port logs of the Compute Engine instance.
- C. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
- D. Determine whether there is an increased number of connections to the Cloud SQL instance.
- E. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Answer(s): C

11. 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(s): D

12. You support an application running on App Engine. The application is used globally and accessed from various device types. You want to know the number of connections. You are using Stackdriver Monitoring for App Engine.

What metric should you use?

A. flex/connections/current

B. tcp_ssl_proxy/new_connections

C. tcp_ssl_proxy/open_connections

D. flex/instance/connections/current

Answer(s): A

13. You support a production service that runs on a single Compute Engine instance. You regularly need to spend time on recreating the service by deleting the crashing instance and creating a new instance based on the relevant image. You want to reduce the time spent performing manual operations while following Site Reliability Engineering principles.

What should you do?

A. File a bug with the development team so they can find the root cause of the crashing instance.

B. Create a Managed Instance Group with a single instance and use health checks to determine the system status.

C. Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.

D. Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.

Answer(s): B

14. You are managing the production deployment to a set of Google Kubernetes Engine (GKE) clusters. You want to make sure only images which are successfully built by your trusted CI/CD pipeline are deployed to production.

What should you do?

A. Enable Cloud Security Scanner on the clusters.

B. Enable Vulnerability Analysis on the Container Registry.

C. Set up the Kubernetes Engine clusters as private clusters.

D. Set up the Kubernetes Engine clusters with Binary Authorization.

Answer(s): D

15. 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

Answer(s): B

16. You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. All PII entries begin with the text `userinfo`. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging.

What should you do?

A. Create a basic log filter matching `userinfo`, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing `userinfo`, and then copy the entries to a Cloud Storage bucket.

C. Create an advanced log filter matching `userinfo`, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a log exclusion with `userinfo` as a filter.

D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing `userinfo`, create an advanced log filter matching `userinfo`, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer(s): B

17. Your team uses Cloud Build for all CI/CO pipelines. You want to use the `kubectl` builder for Cloud Build to deploy new images to Google Kubernetes Engine (GKE). You need to authenticate to GKE while minimizing development effort.

What should you do?

A. Assign the Container Developer role to the Cloud Build service account.

B. Specify the Container Developer role for Cloud Build in the `cloudbuild.yaml` file.

C. Create a new service account with the Container Developer role and use it to run Cloud Build.

D. Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

Answer(s): A

18. You need to reduce the cost of virtual machines (VM) for your organization. After reviewing different options, you decide to leverage preemptible VM instances.

Which application is suitable for preemptible VMs?

A. A scalable in-memory caching system

B. The organization's public-facing website

C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum

D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket

Answer(s): D

19. You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity.

What should you do?

A. Use the n1-highcpu-96 machine type in the configuration of the MIG.

B. Monitor results of Stackdriver Trace to determine the required amount of resources.

C. Validate that the resource requirements are within the available quota limits of each region.

D. Deploy the service in one region and use a global load balancer to route traffic to this region.

Answer(s): C

20. You support an e-commerce application that runs on a large Google Kubernetes Engine (GKE) cluster deployed on-premises and on Google Cloud Platform. The application consists of microservices that run in containers. You want to identify containers that are using the most CPU

and memory.

What should you do?

A. Use Stackdriver Kubernetes Engine Monitoring.

B. Use Prometheus to collect and aggregate logs per container, and then analyze the results in Grafana.

C. Use the Stackdriver Monitoring API to create custom metrics, and then organize your containers using groups.

D. Use Stackdriver Logging to export application logs to BigQuery, aggregate logs per container, and then analyze CPU and memory consumption.

Answer(s): A
