Google

GCP-PCDE
Google Professional Cloud DevOps Engineer

For More Information – Visit link below:
https://www.examsempire.com/

Product Version

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**

https://examsempire.com/

Latest Version: 6.0

Question: 1

You have an application deployed on Google Kubernetes Engine (GKE). The application logs are captured by Cloud Logging. You need to remove sensitive data before it reaches the Cloud Logging API. What should you do?
Response:

A. Customize the GKE clusters' Fluentd configuration with a filter rule. Update the Fluentd Config Map and Daemon Set in the GKE cluster.
B. Write the log information to the container file system. Execute a second process inside the container that will filter the sensitive information before writing to Standard Output.
C. Configure a filter in the Cloud Logging UI to exclude the logs with sensitive data.
D. Configure BigQuery as a sink for the logs from Cloud Logging, and then create a Data Loss Prevention job.

Answer: A

Question: 2

Your application runs in Google Kubernetes Engine (GKE). You want to use Spinnaker with the Kubernetes Provider to perform blue/green deployments and control which version of the application receives traffic. What should you do?
Response:

A. Use a Kubernetes Replica Set and use Spinnaker to create a new service for each new version of the application to be deployed.
B. Use a Kubernetes Replica Set and use Spinnaker to update the Replica Set for each new version of the application to be deployed.
C. Use a Kubernetes Deployment and use Spinnaker to update the deployment for each new version of the application to be deployed.
D. Use a Kubernetes Deployment and use Spinnaker to create a new deployment object for each new version of the application to be deployed.

Answer: B

Question: 3

You support a Python application running in production on Compute Engine. You want to debug some of the application code by inspecting the value of a specific variable. What should you do?
Response:

A. Create a Cloud Debugger logpoint with the variable at a specific line location in your application's source code, and view the value in the Logs Viewer.
B. Use your local development environment and code editor to set up a breakpoint in the source code, run the application locally, and then inspect the value of the variable.
C. Modify the source code of the application to log the value of the variable, deploy to the development environment, and then run the application to capture the value in Cloud Logging.
D. Create a Cloud Debugger snapshot at a specific line location in your application's source code, and view the value of the variable in the Google Cloud Console.

Answer: D

Question: 4

You work with a video rendering application that publishes small tasks as messages to a Cloud Pub/Sub topic. You need to deploy the application that will execute these tasks on multiple virtual machines (VMs).
Each task takes less than 1 hour to complete. The rendering is expected to be completed within a month. You need to minimize rendering costs. What should you do?
Response:

A. Deploy the application as a managed instance group with Preemptible VMs.
B. Deploy the application as a managed instance group. Configure a Committed Use Discount for the amount of CPU and memory required.
C. Deploy the application as a managed instance group.
D. Deploy the application as a managed instance group with Preemptible VMs. Configure a Committed Use Discount for the amount of CPU and memory required.

Answer: A

Question: 5

You have a Compute Engine instance that uses the default Debian image. The application hosted on this instance recently suffered a series of crashes that you weren't able to debug in real time: the application process died suddenly every time.
The application usually consumes 50% of the instance's memory, and normally never more than 70%, but you suspect that a memory leak was responsible for the crashes. You want to validate this hypothesis. What should you do?
Response:

A. Go to Metrics Explorer and look for the “compute.googleapis.com/guest/system/problem_count” metric for that instance. Examine its value for when the application crashed in the past.
B. In Cloud Monitoring, create an uptime check for your application. Create an alert policy for that uptime check to be notified when your application crashes. When you receive an alert, use your usual debugging tools to investigate the behavior of the application in real time.
C. Install the Cloud Monitoring agent on the instance. Go to Metrics Explorer and look for the “agent.googleapis.com/memory/percent_used” metric for that instance. Examine its value for when the application crashed in the past.
D. Install the Cloud Monitoring agent on the instance. Create an alert policy on the “agent.googleapis.com/memory/percent_used” metric for that instance to be alerted when the memory used is higher than 75%. When you receive an alert, use your usual debugging tools to investigate the behavior of the application in real time.

Answer: D

Question: 6

You have a service running on Compute Engine virtual machine instances behind a global load balancer. You need to ensure that when an instance fails, it is recovered. What should you do?
Response:

A. Set up health checks in the load balancer configuration.
B. Deploy a service to the instances to notify you when they fail.
C. Use Cloud Logging alerts to trigger a workflow to reboot the instance.
D. Set up health checks in the managed instance group configuration.

Answer: D

Question: 7

You are deploying an application to a Kubernetes cluster that requires a username and password to connect to another service. When you deploy the application, you want to ensure that the credentials are used securely in multiple environments with minimal code changes. What should you do?
Response:

A. Bundle the credentials with the code inside the container and secure the container registry.
B. Leverage a CI/CD pipeline to update the variables at build time and inject them into a templated Kubernetes application manifest.
C. Store the credentials as a Kubernetes Secret and let the application access it via environment variables at runtime.
D. Store the credentials as a Kubernetes ConfigMap and let the application access it via environment variables at runtime.

Answer: C

Question: 8

Your Site Reliability Engineering team does toil work to archive unused data in tables within your application's relational database. This toil is required to ensure that your application has a low Latency Service Level Indicator (SLI) to meet your Service Level Objective (SLO).
Toil is preventing your team from focusing on a high-priority engineering project that will improve the Availability SLI of your application. You want to: (1) reduce repetitive tasks to avoid burnout, (2) improve organizational efficiency, and (3) follow the Site Reliability Engineering recommended practices.
What should you do?
Response:

A. Identify repetitive tasks that contribute to toil and onboard additional team members for support.
B. Identify repetitive tasks that contribute to toil and automate them.
C. Change the SLO of your Latency SLI to accommodate toil being done less often. Use this capacity to work on the Availability SLI engineering project.
D. Assign the Availability SLI engineering project to the Software Engineering team.

Answer: B

Question: 9

You support a website with a global audience. The website has a frontend web service and a backend database service that runs on different clusters. All clusters are scaled to handle at least⅓ of the total user traffic. You use 4 different regions in Google Cloud and Cloud Load Balancing to direct traffic to a region closer to the user.
You are applying a critical security patch to the backend database. You successfully patch the database in the first 2 regions, but you make a configuration error while patching Region 3. The unsuccessful patching causes 50% of user requests to Region 3 to time out. You want to mitigate the impact of unsuccessful patching on users.
What should you do?
Response:

A. Add more capacity to the frontend of Region 3.
B. Revert the Region 3 backend database and run it without the patch.
C. Drain the requests to Region 3 and redirect new requests to other regions.
D. Back up the database in the backend of Region 3 and restart the database.

Answer: C

Thank You for Trying Our Product

Discount Coupon Code is: 20OFF2022

Email: support@examsempire.com

Check our Customer Testimonials and ratings available on every product page.

Visit our website.

https://examsempire.com/