

Professional-Cloud-DevOps-Engineer Dumps

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

<https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html>



NEW QUESTION 1

You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

- A. Analyze VPC flow logs along the path of the request.
- B. Investigate the Liveness and Readiness probes for each service.
- C. Create a Dataflow pipeline to analyze service metrics in real time.
- D. Use a distributed tracing framework such as OpenTelemetry or Stackdriver Trace.

Answer: C

NEW QUESTION 2

You encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

- A. Call individual stakeholders to explain what happened.
- B. Develop a post-mortem to be distributed to stakeholders.
- C. Send the Incident State Document to all the stakeholders.
- D. Require the engineer responsible to write an apology email to all stakeholders.

Answer: B

NEW QUESTION 3

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: A

Explanation:

https://cloud.google.com/monitoring/api/metrics_gcp#gcp-appengine

NEW QUESTION 4

You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

- A. Reroute the user traffic from the affected region to other regions that don't report issues.
- B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
- C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
- D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Answer: A

Explanation:

Google always aims to first stop the impact of an incident, and then find the root cause (unless the root cause just happens to be identified early on).

NEW QUESTION 5

Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?

- A. Implement Jenkins on local workstations.
- B. Implement Jenkins on Kubernetes on-premises
- C. Implement Jenkins on Google Cloud Functions.
- D. Implement Jenkins on Compute Engine virtual machines.

Answer: D

Explanation:

Your application runs on Google Cloud Platform (GCP). You need to implement Jenkins for deploying application releases to GCP. You want to streamline the release process, lower operational toil, and keep user data secure. What should you do?

<https://plugins.jenkins.io/google-compute-engine/>

NEW QUESTION 6

You support a trading application written in Python and hosted on App Engine flexible environment. You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

- A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
- B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
- C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.

D. Use the Stackdriver Error Reporting API to write errors from your application to ReportedErrorEvent, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Answer: D

Explanation:

<https://cloud.google.com/error-reporting/docs/formatting-error-messages> <https://cloud.google.com/error-reporting/docs/reference/libraries#client-libraries-install-python> no need to install error reporting library on App Engine Flex.

NEW QUESTION 7

You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.

You need to troubleshoot the issue. What should you do?

- A. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.
- B. Confirm that your account has the proper permissions to use the Stackdriver dashboard.
- C. Confirm that port 25 has been opened in the firewall to allow messages through to Stackdriver.
- D. Confirm that the application is using the required client library and the service account key has proper permissions.

Answer: A

Explanation:

<https://cloud.google.com/monitoring/agent/monitoring/troubleshooting#checklist>

NEW QUESTION 8

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 log
- D. Create an Alert with a Webhook notification type.
- E. Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

Answer: D

NEW QUESTION 9

You have a CI/CD pipeline that uses Cloud Build to build new Docker images and push them to Docker Hub. You use Git for code versioning. After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline. You need to resolve the issue following Site Reliability Engineering practices. What should you do?

- A. Disable the CI pipeline and revert to manually building and pushing the artifacts.
- B. Change the CI pipeline to push the artifacts to Container Registry instead of Docker Hub.
- C. Upload the configuration YAML file to Cloud Storage and use Error Reporting to identify and fix the issue.
- D. Run a Git compare between the previous and current Cloud Build Configuration files to find and fix the bug.

Answer: D

Explanation:

"After making a change in the Cloud Build YAML configuration, you notice that no new artifacts are being built by the pipeline"- means something wrong on the recent change not with the image registry.

NEW QUESTION 10

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. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- C. Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- D. Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

Answer: A

NEW QUESTION 10

You support a service with a well-defined Service Level Objective (SLO). Over the previous 6 months, your service has consistently met its SLO and customer satisfaction has been consistently high. Most of your service's operations tasks are automated and few repetitive tasks occur frequently. You want to optimize the balance between reliability and deployment velocity while following site reliability engineering best practices. What should you do? (Choose two.)

- A. Make the service's SLO more strict.
- B. Increase the service's deployment velocity and/or risk.
- C. Shift engineering time to other services that need more reliability.
- D. Get the product team to prioritize reliability work over new features.
- E. Change the implementation of your Service Level Indicators (SLIs) to increase coverage.

Answer: BC

Explanation:

(<https://sre.google/workbook/implementing-slos/#slo-decision-matrix>)

NEW QUESTION 14

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: D

Explanation:

<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 19

You are responsible for creating and modifying the Terraform templates that define your Infrastructure. Because two new engineers will also be working on the same code, you need to define a process and adopt a tool that will prevent you from overwriting each other's code. You also want to ensure that you capture all updates in the latest version. What should you do?

- A. • Store your code in a Git-based version control system. • Establish a process that allows developers to merge their own changes at the end of each day. • Package and upload code to a versioned Cloud Storage bucket as the latest master version.
- B. • Store your code in a Git-based version control system. • Establish a process that includes code reviews by peers and unit testing to ensure integrity and functionality before integration of code. • Establish a process where the fully integrated code in the repository becomes the latest master version.
- C. • Store your code as text files in Google Drive in a defined folder structure that organizes the files. • At the end of each day, confirm that all changes have been captured in the files within the folder structure.
- D. confirm that all changes have been captured in the files within the folder structure. • Rename the folder structure with a predefined naming convention that increments the version.
- E. • Store your code as text files in Google Drive in a defined folder structure that organizes the files. • At the end of each day, confirm that all changes have been captured in the files within the folder structure and create a new .zip archive with a predefined naming convention. • Upload the .zip archive to a versioned Cloud Storage bucket and accept it as the latest version.

Answer: B

NEW QUESTION 24

You support an application that stores product information in cached memory. For every cache miss, an entry is logged in Stackdriver Logging. You want to visualize how often a cache miss happens over time. What should you do?

- A. Link Stackdriver Logging as a source in Google Data Studio
- B. Filter the logs on the cache misses.
- C. Configure Stackdriver Profiler to identify and visualize when the cache misses occur based on the logs.
- D. Create a logs-based metric in Stackdriver Logging and a dashboard for that metric in Stackdriver Monitoring.
- E. Configure BigQuery as a sink for Stackdriver Logging
- F. Create a scheduled query to filter the cache miss logs and write them to a separate table

Answer: C

Explanation:

<https://cloud.google.com/logging/docs/logs-based-metrics#counter-metric>

NEW QUESTION 26

Your team is designing a new application for deployment into Google Kubernetes Engine (GKE). You need to set up monitoring to collect and aggregate various application-level metrics in a centralized location. You want to use Google Cloud Platform services while minimizing the amount of work required to set up monitoring. What should you do?

- A. Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.
- B. Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.
- C. Install the OpenTelemetry client libraries in the application, configure Stackdriver as the export destination for the metrics, and then observe the application's metrics in Stackdriver.
- D. Emit all metrics in the form of application-specific log messages, pass these messages from the containers to the Stackdriver logging collector, and then observe metrics in Stackdriver.

Answer: A

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics#custom_metrics <https://github.com/GoogleCloudPlatform/k8s-stackdriver/blob/master/custom-metrics-stackdriver-adapter/README.md> Your application can report a custom metric to Cloud Monitoring. You can configure Kubernetes to respond to these metrics and scale your workload automatically. For example, you can scale your application based on metrics such as queries per second, writes per second, network performance, latency when communicating with a different application, or other metrics that make sense for your workload.
<https://cloud.google.com/kubernetes-engine/docs/concepts/custom-and-external-metrics>

NEW QUESTION 30

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: AC

Explanation:

<https://sre.google/workbook/incident-response/>

"The main roles in incident response are the Incident Commander (IC), Communications Lead (CL), and Operations or Ops Lead (OL)."

NEW QUESTION 31

You support an application running on GCP and want to configure SMS notifications to your team for the most critical alerts in Stackdriver Monitoring. You have already identified the alerting policies you want to configure this for. What should you do?

- A. Download and configure a third-party integration between Stackdriver Monitoring and an SMS gateway. Ensure that your team members add their SMS/phone numbers to the external tool.
- B. Select the Webhook notifications option for each alerting policy, and configure it to use a third-party integration too
- C. Ensure that your team members add their SMS/phone numbers to the external tool.
- D. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile
- E. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
- F. Configure a Slack notification for each alerting policy
- G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received
- H. Ensure that your team members add their SMS/phone numbers to the external integration.

Answer: C

Explanation:

https://cloud.google.com/monitoring/support/notification-options#creating_channels To configure SMS notifications, do the following:

In the SMS section, click Add new and follow the instructions. Click Save. When you set up your alerting policy, select the SMS notification type and choose a verified phone number from the list.

NEW QUESTION 36

You are creating and assigning action items in a postmortem for an outage. The outage is over, but you need to address the root causes. You want to ensure that your team handles the action items quickly and efficiently. How should you assign owners and collaborators to action items?

- A. Assign one owner for each action item and any necessary collaborators.
- B. Assign multiple owners for each item to guarantee that the team addresses items quickly
- C. Assign collaborators but no individual owners to the items to keep the postmortem blameless.
- D. Assign the team lead as the owner for all action items because they are in charge of the SRE team.

Answer: A

Explanation:

<https://devops.com/when-it-disaster-strikes-part-3-conducting-a-blameless-post-mortem/>

NEW QUESTION 38

Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

- A. Store the Terraform code in a version-control system
- B. Establish procedures for pushing new versions and merging with the master.
- C. Store the Terraform code in a network shared folder with child folders for each version release
- D. Ensure that everyone works on different files.
- E. Store the Terraform code in a Cloud Storage bucket using object versioning
- F. Give access to the bucket to every team member so they can download the files.
- G. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member's computer
- H. Organize files with a naming convention that identifies each new version.

Answer: A

Explanation:

<https://www.terraform.io/docs/cloud/guides/recommended-practices/part3.3.html>

NEW QUESTION 42

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: D

Explanation:

Costs associated with a load balancer are charged to the project containing the load balancer components. Because of these benefits, container-native load balancing is the recommended solution for load balancing through Ingress. When NEGs are used with GKE Ingress, the Ingress controller facilitates the creation of

all aspects of the L7 load balancer. This includes creating the virtual IP address, forwarding rules, health checks, firewall rules, and more.
<https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke>

NEW QUESTION 44

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: C

Explanation:

https://cloud.google.com/compute/quotas#understanding_quotas <https://cloud.google.com/compute/quotas>

NEW QUESTION 48

You support a service that recently had an outage. The outage was caused by a new release that exhausted the service memory resources. You rolled back the release successfully to mitigate the impact on users. You are now in charge of the post-mortem for the outage. You want to follow Site Reliability Engineering practices when developing the post-mortem. What should you do?

- A. Focus on developing new features rather than avoiding the outages from recurring.
- B. Focus on identifying the contributing causes of the incident rather than the individual responsible for the cause.
- C. Plan individual meetings with all the engineers involved
- D. Determine who approved and pushed the new release to production.
- E. Use the Git history to find the related code commit
- F. Prevent the engineer who made that commit from working on production services.

Answer: B

NEW QUESTION 49

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Professional-Cloud-DevOps-Engineer Exam with Our Prep Materials Via below:

<https://www.certleader.com/Professional-Cloud-DevOps-Engineer-dumps.html>