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 support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google-recommended way of calculating this SLI?

A. Buckelize Ihe request latencies into ranges, and then compute the percentile at 100 ms.
B. Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.
C. Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
D. Count the number of home page requests that load in under 100 m
E. and then divide by the total number of all web application requests.

Answer: C

Explanation:
https://sre.google/workbook/implementing-slos/
In the SRE principles book, it's recommended treating the SLI as the ratio of two numbers: the number of good events divided by the total number of events. For example: Number of successful HTTP requests / total HTTP requests (success rate)

NEW QUESTION 2
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 3
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: B

NEW QUESTION 4
You have a pool of application servers running on Compute Engine. You need to provide a secure solution that requires the least amount of configuration and allows developers to easily access application logs for troubleshooting. How would you implement the solution on GCP?

A. • Deploy the Stackdriver logging agent to the application servers.• Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
B. • Deploy the Stackdriver logging agent to the application servers.• Give the developers the IAM Logs Private Logs Viewer role to access Stackdriver and view logs.
C. • Deploy the Stackdriver monitoring agent to the application servers.• Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
D. • Install the gsutil command line tool on your application servers.• Write a script using gsutil to upload your application log to a Cloud Storage bucket, and then schedule it to run via cron every 5 minutes.• Give the developers IAM Object Viewer access to view the logs in the specified bucket.

Answer: A

Explanation:
https://cloud.google.com/logging/docs/audit#access-control

NEW QUESTION 5
You need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

A. Purchase Committed Use Discounts.
B. Migrate the instances to a Managed Instance Group.
C. Convert the instances to preemptible virtual machines.
D. Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: A

NEW QUESTION 6

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 Buitd 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 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: D

Explanation:
The keywords here is "developers or operators". Option A the operators could push images to production without approval (operators could touch the cluster directly and the cluster cannot do any action against them). Rest same as francisco_guerra.

NEW QUESTION 8
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 BigOuery.* 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 (rom Stackdriver to BigQuery.* 2. From BigQuer
D. 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.
E. * 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.

Answer: A

NEW QUESTION 9
You are on-call for an infrastructure service that has a large number of dependent systems. You receive an alert indicating that the service is failing to serve most of its requests and all of its dependent systems with hundreds of thousands of users are affected. As part of your Site Reliability Engineering (SRE) incident management protocol, you declare yourself Incident Commander (IC) and pull in two experienced people from your team as Operations Lead (OLJ and Communications Lead (CL). What should you do next?

A. Look for ways to mitigate user impact and deploy the mitigations to production.
B. Contact the affected service owners and update them on the status of the incident.
C. Establish a communication channel where incident responders and leads can communicate with each other.
D. Start a postmortem, add incident information, circulate the draft internally, and ask internal stakeholders for input.

Answer: A

Explanation:
https://sre.google/sre-book/managing-incidents/

NEW QUESTION 10
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 Studi
B. Filler (he 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 BigOuery as a sink for Stackdriver Loggin
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 10
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 Slackdriver 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/REA 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 14
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 15
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 Profil
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 polic
G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are receive
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 17
You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week. If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability. What should you do before a major incident occurs?

A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
B. Negotiate with the product team to always prioritize service reliability over releasing new features.
C. Negotiate with the development team to reduce the release frequency to no more than once a week.
D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Answer: A

Explanation:
Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget(It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.
The goals of this policy are to: -- Protect customers from repeated SLO misses -- Provide an incentive to balance reliability with other features
https://sre.google/workbook/error-budget-policy/

NEW QUESTION 18
You support a web application that is hosted on Compute Engine. The application provides a booking service for thousands of users. Shortly after the release of a new feature, your monitoring dashboard shows that all users are experiencing latency at login. You want to mitigate the impact of the incident on the users of your service. What should you do first?

A. Roll back the recent release.
B. Review the Stackdriver monitoring.
C. Upsize the virtual machines running the login services.
D. Deploy a new release to see whether it fixes the problem.

Answer: C

Explanation:
Rollback to previous stable version. Then you need to find what is causing the issue.

NEW QUESTION 20
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: A

NEW QUESTION 23
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 syste
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 releas
D. Ensure that everyone works on different files.
E. Store the Terraform code in a Cloud Storage bucket using object versionin
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 compute
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 25
Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

A. Introduce the new version of the API.Announce deprecation of the old version of the AP
B. Deprecate the old version of the API.Contact remaining users of the old API.Provide best effort support to users of the old AP
C. Turn down the old version of the API.
D. Announce deprecation of the old version of the AP
E. Introduce the new version of the API.Contact remaining users on the old AP
F. Deprecate the old version of the AP
G. Turn down the old version of the API.Provide best effort support to users of the old API.
H. Announce deprecation of the old version of the AP
I. Contact remaining users on the old API.Introduce the new version of the AP
J. Deprecate the old version of the API.Provide best effort support to users of the old AP
K. Turn down the old version of the API.
L. Introduce the new version of the AP
M. Contact remaining users of the old API.Announce deprecation of the old version of the AP
N. Deprecate the old version of the API.Turn down the old version of the API.Provide best effort support to users of the old API.

Answer: A

NEW QUESTION 28
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 duster 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 29
You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. 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 2 answers

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

Answer: BE

NEW QUESTION 30
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 Slacked 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 metricKin
E. set toMETRIc_KIND_UNSPECIFIEDanda valueType set to INT64.• In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C

Explanation:
 https://sre.google/workbook/implementing-slos/ https://cloud.google.com/architecture/adopting-slos/
Latency is commonly measured as a distribution. Given a distribution, you can measure various percentiles.
For example, you might measure the number of requests that are slower than the historical 99th percentile.

NEW QUESTION 32
You are running a real-time gaming application on Compute Engine that has a production and testing environment. Each environment has their own Virtual Private Cloud (VPC) network. The application frontend and backend servers are located on different subnets in the environment's VPC. You suspect there is a malicious process communicating intermittently in your production frontend servers. You want to ensure that network traffic is captured for analysis. What should you do?

A. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 0.5.
B. Enable VPC Flow Logs on the production VPC network frontend and backend subnets only with a sample volume scale of 1.0.
C. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 0.5. Apply changes intesting before production.
D. Enable VPC Flow Logs on the testing and production VPC network frontend and backend subnets with a volume scale of 1.0. Apply changes in testing before production.

Answer: D

NEW QUESTION 33
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 BigOuer
E. aggregate logs per container, and then analyze CPU and memory consumption.

Answer: A

Explanation:
https://cloud.google.com/anthos/clusters/docs/on-prem/1.7/concepts/logging-and-monitoring

NEW QUESTION 35
You are performing a semiannual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP). using a Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

A. Verity the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verity your expected resource needs.
B. Because you are deployed on GKE and are using a cluster autoscale
C. your GKE cluster will scale automatically, regardless of growth rate.
D. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
E. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Answer: A

Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler
The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption

NEW QUESTION 38
......

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