



Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

NEW QUESTION 1

You use Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure the canary analysis?

- A. Compare the canary with a new deployment of the current production version.
- B. Compare the canary with a new deployment of the previous production version.
- C. Compare the canary with the existing deployment of the current production version.
- D. Compare the canary with the average performance of a sliding window of previous production versions.

Answer: A

Explanation:

<https://cloud.google.com/architecture/automated-canary-analysis-kubernetes-engine-spinnaker> <https://spinnaker.io/guides/user/canary/best-practices/#compare-canary-against-baseline-not-against-production>

NEW QUESTION 2

You encounter a large number of outages in the production systems you support. You receive alerts for all the outages that wake you up at night. The alerts are due to unhealthy systems that are automatically restarted within a minute. You want to set up a process that would prevent staff burnout while following Site Reliability Engineering practices. What should you do?

- A. Eliminate unactionable alerts.
- B. Create an incident report for each of the alerts.
- C. Distribute the alerts to engineers in different time zones.
- D. Redefine the related Service Level Objective so that the error budget is not exhausted.

Answer: A

Explanation:

Eliminate bad monitoring : Unactionable alerts (i.e., spam) <https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles>

agree with kyubiblaze about having to remove unactionable items aka spam: "good monitoring alerts on actionable problems" @ <https://cloud.google.com/blog/products/management-tools/meeting-reliability-challenges-with-sre-principles>

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 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 GK
- C. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Install Kubernetes on Google Compute Engine (GCE) and redeploy your application
- E. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- F. Write a script to tail the log file within the pod and write entries to standard output
- G. Run the script as a sidecar container with the application's pod
- H. 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:

<https://cloud.google.com/architecture/customizing-stackdriver-logs-fluentd>

Besides the list of default logs that the Logging agent streams by default, you can customize the Logging agent to send additional logs to Logging or to adjust agent settings by adding input configurations. The configuration definitions in these sections apply to the fluent-plugin-google-cloud output plugin only and specify how logs are transformed and ingested into Cloud Logging. <https://cloud.google.com/logging/docs/agent/logging/configuration#configure>

NEW QUESTION 5

Some of your production services are running in Google Kubernetes Engine (GKE) in the eu-west-1 region. Your build system runs in the us-west-1 region. You want to push the container images from your build system to a scalable registry to maximize the bandwidth for transferring the images to the cluster. What should you do?

- A. Push the images to Google Container Registry (GCR) using the gcr.io hostname.
- B. Push the images to Google Container Registry (GCR) using the us.gcr.io hostname.

- C. Push the images to Google Container Registry (GCR) using the eu.gcr.io hostname.
- D. Push the images to a private image registry running on a Compute Engine instance in the eu-west-1 region.

Answer: C

Explanation:

Hostname Storage location gcr.io Stores images in data centers in the United States asia.gcr.io Stores images in data centers in Asia eu.gcr.io Stores images in data centers within member states of the European Union us.gcr.io Stores images in data centers in the United States

NEW QUESTION 6

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

- A. Check the serial port logs of the Compute Engine instance.
- B. Use Stackdriver Profiler to visualize the resources utilization throughout the application.
- C. Determine whether there is an increased number of connections to the Cloud SQL instance.
- D. Use Cloud Security Scanner to see whether your Cloud SQL is under a Distributed Denial of Service (DDoS) attack.

Answer: B

NEW QUESTION 7

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 8

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 9

Your company is developing applications that are deployed on Google Kubernetes Engine (GKE). Each team manages a different application. You need to create the development and production environments for each team, while minimizing costs. Different teams should not be able to access other teams' environments. What should you do?

- A. Create one GCP Project per tea
- B. In each project, create a cluster for Development and one for Productio
- C. Grant the teams IAM access to their respective clusters.
- D. Create one GCP Project per tea
- E. In each project, create a cluster with a Kubernetes namespace for Development and one for Productio
- F. Grant the teams IAM access to their respective clusters.
- G. Create a Development and a Production GKE cluster in separate project
- H. In each cluster, create a Kubernetes namespace per team, and then configure Identity Aware Proxy so that each team can only access its own namespace.
- I. Create a Development and a Production GKE cluster in separate project
- J. In each cluster, create a Kubernetes namespace per team, and then configure Kubernetes Role-based access control (RBAC) so that each team can only access its own namespace.

Answer: D

Explanation:

https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups

NEW QUESTION 10

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 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 12

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 14

You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality – a data intensive reporting feature – is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

- A. As the I/O wait times aggregated across all report generation backends
- B. As the proportion of report generation requests that result in a successful response
- C. As the application's report generation queue size compared to a known-good threshold
- D. As the reporting backend PD throughput capacity compared to a known-good threshold

Answer: B

Explanation:

According to SRE Workbook, one of potential SLI is as below:

* Type of service: Request-driven

* Type of SLI: Availability

* Description: The proportion of requests that resulted in a successful response. <https://sre.google/workbook/implementing-slos/>

NEW QUESTION 15

You are creating and assigning action items in a postmodern 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 19

You are running an experiment to see whether your users like a new feature of a web application. Shortly after deploying the feature as a canary release, you receive a spike in the number of 500 errors sent to users, and your monitoring reports show increased latency. You want to quickly minimize the negative impact on users.

What should you do first?

- A. Roll back the experimental canary release.
- B. Start monitoring latency, traffic, errors, and saturation.
- C. Record data for the postmortem document of the incident.
- D. Trace the origin of 500 errors and the root cause of increased latency.

Answer: A

NEW QUESTION 21

You created a Stackdriver chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?

- A. Share the workspace Project ID with the SRE tea
- B. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- C. Share the workspace Project ID with the SRE tea
- D. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
- E. Click "Share chart by URL" and provide the URL to the SRE tea
- F. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- G. Click "Share chart by URL" and provide the URL to the SRE tea
- H. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Answer: C

Explanation:

<https://cloud.google.com/monitoring/access-control>

NEW QUESTION 26

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: B

Explanation:

<https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud> <https://spinnaker.io/guides/user/pipeline/triggers/pubsub/>

NEW QUESTION 30

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: D

Explanation:

<https://cloud.google.com/binary-authorization/docs/overview>

NEW QUESTION 33

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: C

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/autoscaling-metrics>

NEW QUESTION 37

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 39

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 server
- 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 41

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
- 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 42

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-DevOps-Engineer Practice Exam Features:

- * Professional-Cloud-DevOps-Engineer Questions and Answers Updated Frequently
- * Professional-Cloud-DevOps-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-DevOps-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-DevOps-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-DevOps-Engineer Practice Test Here](#)