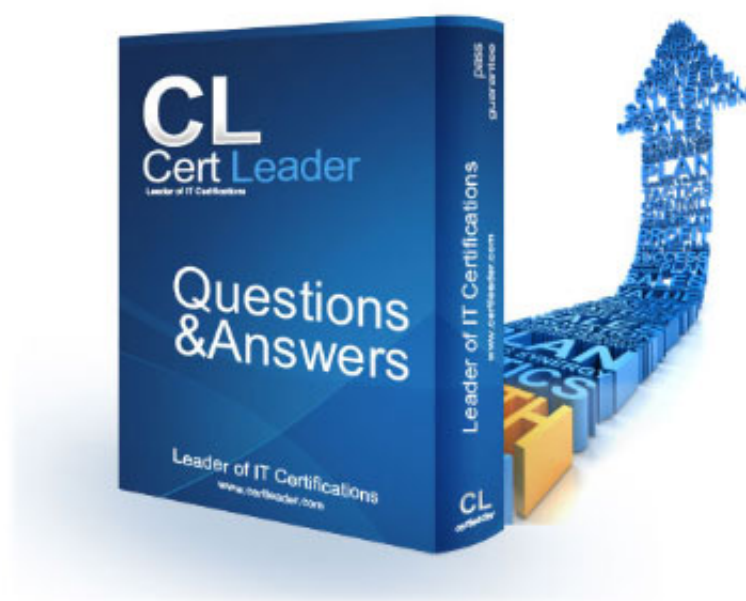


Professional-Cloud-DevOps-Engineer Dumps

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. Bucketize the 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 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 are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods. What should you do?

- A. Use a partitioned rolling update.
- B. Use Node taints with NoExecute.
- C. Use a replica set in the deployment specification.
- D. Use a stateful set with parallel pod management policy.

Answer: A

Explanation:

<https://medium.com/velotio-perspectives/exploring-upgrade-strategies-for-stateful-sets-in-kubernetes-c02b8286f>

NEW QUESTION 4

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 5

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 6

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the

90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

- A. 90th percentile – 100ms 95th percentile – 250ms
- B. 90th percentile – 120ms 95th percentile – 275ms
- C. 90th percentile – 150ms 95th percentile – 300ms
- D. 90th percentile – 250ms 95th percentile – 400ms

Answer: C

Explanation:

<https://sre.google/sre-book/service-level-objectives/>

NEW QUESTION 7

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

- A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.
- D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Answer: D

Explanation:

<https://cloud.google.com/logging/docs/routing/overview>

NEW QUESTION 8

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: CE

Explanation:

<https://cloud.google.com/architecture/adopting-slos?hl=en>

NEW QUESTION 9

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 10

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 BigQuery.* 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 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 15

You are running an application in a virtual machine (VM) using a custom Debian image. The image has the Stackdriver Logging agent installed. The VM has the cloud-platform scope. The application is logging information via syslog. You want to use Stackdriver Logging in the Google Cloud Platform Console to visualize the logs. You notice that syslog is not showing up in the "All logs" dropdown list of the Logs Viewer. What is the first thing you should do?

- A. Look for the agent's test log entry in the Logs Viewer.
- B. Install the most recent version of the Stackdriver agent.
- C. Verify the VM service account access scope includes the monitoring.write scope.
- D. SSH to the VM and execute the following commands on your VM: `ps ax | grep fluentd`

Answer: D

Explanation:

https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION 20

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: A

Explanation:

<https://cloud.google.com/build/docs/deploying-builds/deploy-gke> <https://cloud.google.com/build/docs/securing-builds/configure-user-specified-service-accounts>

NEW QUESTION 21

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 26

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 28

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 the production environment. A recent security audit alerted your team that the code pushed to production could contain vulnerabilities and that the existing tooling around virtual machine (VM) vulnerabilities no longer applies to the containerized environment. You need to ensure the security and patch level of all code running through the pipeline. What should you do?

- A. Set up Container Analysis to scan and report Common Vulnerabilities and Exposures.
- B. Configure the containers in the build pipeline to always update themselves before release.
- C. Reconfigure the existing operating system vulnerability software to exist inside the container.
- D. Implement static code analysis tooling against the Docker files used to create the containers.

Answer: D

Explanation:

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

Binary Authorization is a deploy-time security control that ensures only trusted container images are deployed on Google Kubernetes Engine (GKE) or Cloud Run. With Binary Authorization, you can require images to be signed by trusted authorities during the development process and then enforce signature validation when deploying. By enforcing validation, you can gain tighter control over your container environment by ensuring only verified images are integrated into the build-and-release process.

NEW QUESTION 32

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 36

You use Cloud Build to build and deploy your application. You want to securely incorporate database credentials and other application secrets into the build pipeline. You also want to minimize the development effort. What should you do?

- A. Create a Cloud Storage bucket and use the built-in encryption at res
- B. Store the secrets in the bucket and grant Cloud Build access to the bucket.
- C. Encrypt the secrets and store them in the application repositor
- D. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
- E. Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucke
- F. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
- G. Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuratio
- H. Grant Cloud Build access to the KeyRing.

Answer: D

Explanation:

<https://cloud.google.com/build/docs/securing-builds/use-encrypted-credentials>

NEW QUESTION 38

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 43

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 NEG's 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 48

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP'S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN). which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration. CDN failure, or other global networking catastrophe. Where should you measure this new SLI?

Choose 2 answers

- A. Your application servers' logs
- B. Instrumentation coded directly in the client
- C. Metrics exported from the application servers
- D. GKE health checks for your application servers
- E. A synthetic client that periodically sends simulated user requests

Answer: BE

NEW QUESTION 53

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 in testing 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 54

Your application artifacts are being built and deployed via a CI/CD pipeline. You want the CI/CD pipeline to securely access application secrets. You also want to more easily rotate secrets in case of a security breach. What should you do?

- A. Prompt developers for secrets at build time
- B. Instruct developers to not store secrets at rest.
- C. Store secrets in a separate configuration file on Git
- D. Provide select developers with access to the configuration file.
- E. Store secrets in Cloud Storage encrypted with a key from Cloud KMS
- F. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
- G. Encrypt the secrets and store them in the source code repository
- H. Store a decryption key in a separate repository and grant your pipeline access to it

Answer: C

NEW QUESTION 56

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 58

You are developing a strategy for monitoring your Google Cloud Platform (GCP) projects in production using Stackdriver Workspaces. One of the requirements is to be able to quickly identify and react to production environment issues without false alerts from development and staging projects. You want to ensure that you adhere to the principle of least privilege when providing relevant team members with access to Stackdriver Workspaces. What should you do?

- A. Grant relevant team members read access to all GCP production project
- B. Create Stackdriver workspaces inside each project.
- C. Grant relevant team members the Project Viewer IAM role on all GCP production project
- D. Create Stackdriver workspaces inside each project.
- E. Choose an existing GCP production project to host the monitoring workspace
- F. Attach the production projects to this workspace

- G. Grant relevant team members read access to the Stackdriver Workspace.
- H. Create a new GCP monitoring project, and create a Stackdriver Workspace inside i
- I. Attach the production projects to this workspac
- J. Grant relevant team members read access to the Stackdriver Workspace.

Answer: D

Explanation:

"A Project can host many Projects and appear in many Projects, but it can only be used as the scoping project once. We recommend that you create a new Project for the purpose of having multiple Projects in the same scope."

NEW QUESTION 59

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application has never consumed more than 5% of its error budget in any given time window. You hold a Service Level Objective (SLO) review with business stakeholders and confirm that the SLO is set appropriately. You want your application's SLO to more closely reflect its observed reliability. What steps can you take to further that goal while balancing velocity, reliability, and business needs? (Choose two.)

- A. Add more serving capacity to all of your application's zones.
- B. Have more frequent or potentially risky application releases.
- C. Tighten the SLO match the application's observed reliability.
- D. Implement and measure additional Service Level Indicators (SLIs) fro the application.
- E. Announce planned downtime to consume more error budget, and ensure that users are not depending on a tighter SLO.

Answer: DE

Explanation:

<https://sre.google/sre-book/service-level-objectives/>

You want the application's SLO to more closely reflect it's observed reliability. The key here is error budget never goes over 5%. This means they can have additional downtime and still stay within their budget.

NEW QUESTION 60

You support a web application that runs on App Engine and uses CloudSQL and Cloud Storage for data storage. After a short spike in website traffic, you notice a big increase in latency for all user requests, increase in CPU use, and the number of processes running the application. Initial troubleshooting reveals:

After the initial spike in traffic, load levels returned to normal but users still experience high latency. Requests for content from the CloudSQL database and images from Cloud Storage show the same high latency.

No changes were made to the website around the time the latency increased. There is no increase in the number of errors to the users.

You expect another spike in website traffic in the coming days and want to make sure users don't experience latency. What should you do?

- A. Upgrade the GCS buckets to Multi-Regional.
- B. Enable high availability on the CloudSQL instances.
- C. Move the application from App Engine to Compute Engine.
- D. Modify the App Engine configuration to have additional idle instances.

Answer: D

Explanation:

Scaling App Engine scales the number of instances automatically in response to processing volume. This scaling factors in the automatic_scaling settings that are provided on a per-version basis in the configuration file. A service with basic scaling is configured by setting the maximum number of instances in the max_instances parameter of the basic_scaling setting. The number of live instances scales with the processing volume. You configure the number of instances of each version in that service's configuration file. The number of instances usually corresponds to the size of a dataset being held in memory or the desired throughput for offline work. You can adjust the number of instances of a manually-scaled version very quickly, without stopping instances that are currently running, using the Modules API set_num_instances function. <https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed>

<https://cloud.google.com/appengine/docs/standard/python/config/appref>

max_idle_instances Optional. The maximum number of idle instances that App Engine should maintain for this version. Specify a value from 1 to 1000. If not specified, the default value is automatic, which means App Engine will manage the number of idle instances. Keep the following in mind: A high maximum reduces the number of idle instances more gradually when load levels return to normal after a spike. This helps your application maintain steady performance through fluctuations in request load, but also raises the number of idle instances (and consequent running costs) during such periods of heavy load.

NEW QUESTION 63

.....

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>