Professional-Cloud-DevOps-Engineer Dumps

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

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

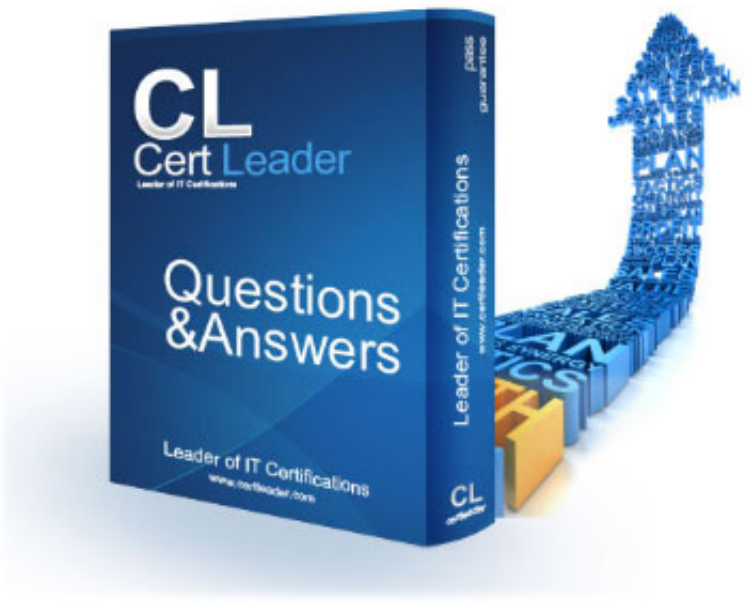

NEW QUESTION 1
You use Spinnaker to deploy your application and have created a canary deployment stage in the pipeline. Your application has an in-memory cache that loads objects at start time. You want to automate the comparison of the canary version against the production version. How should you configure 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 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 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 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 6
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 I grep fluentd

Answer: D

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

NEW QUESTION 7
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 lo 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 da
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 8
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 9
You support a Node.js application running on Google Kubernetes Engine (GKE) in production. The application makes several HTTP requests to dependent applications. You want to anticipate which dependent applications might cause performance issues. What should you do?

A. Instrument all applications with Stackdriver Profiler.
B. Instrument all applications with Stackdriver Trace and review inter-service HTTP requests.
C. Use Stackdriver Debugger to review the execution of logic within each application to instrument all applications.
D. Modify the Node.js application to log HTTP request and response times to dependent application
E. Use Stackdriver Logging to find dependent applications that are performing poorly.

Answer: B

NEW QUESTION 10
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 10
Your company follows Site Reliability Engineering practices. You are the person in charge of Communications for a large, ongoing incident affecting your customer-

facing applications. There is still no estimated time for a resolution of the outage. You are receiving emails from internal stakeholders who want updates on the outage, as well as emails from customers who want to know what is happening. You want to efficiently provide updates to everyone affected by the outage. What should you do?

A. Focus on responding to internal stakeholders at least every 30 minute
B. Commit to "next update" times.
C. Provide periodic updates to all stakeholders in a timely manne
D. Commit to a "next update" time in all communications.
E. Delegate the responding to internal stakeholder emails to another member of the Incident Response Tea
F. Focus on providing responses directly to customers.
G. Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communication
H. Focus on providing responses directly to customers.

Answer: B

Explanation:
When disaster strikes, the person who declares the incident typically steps into the IC role and directs the high-level state of the incident. The IC concentrates on the 3Cs and does the following: Commands and
coordinates the incident response, delegating roles as needed. By default, the IC assumes all roles that have not been delegated yet. Communicates effectively. Stays in control of the incident response. Works with other responders to resolve the incident. https://sre.google/workbook/incident-response/

NEW QUESTION 13
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 14
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 15
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
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 25
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 30
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 35
You manage several production systems that run on Compute Engine in the same Google Cloud Platform (GCP) project. Each system has its own set of dedicated Compute Engine instances. You want to know how must it costs to run each of the systems. What should you do?

A. In the Google Cloud Platform Console, use the Cost Breakdown section to visualize the costs per system.
B. Assign all instances a label specific to the system they ru
C. Configure BigQuery billing export and query costs per label.
D. Enrich all instances with metadata specific to the system they ru
E. Configure Stackdriver Logging to export to BigQuery, and query costs based on the metadata.
F. Name each virtual machine (VM) after the system it run
G. Set up a usage report export to a Cloud Storage bucke
H. Configure the bucket as a source in BigQuery to query costs based on VM name.

Answer: B

Explanation:
https://cloud.google.com/billing/docs/how-to/export-data-bigquery

NEW QUESTION 36
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 37
You manage an application that is writing logs to Stackdriver Logging. You need to give some team members the ability to export logs. What should you do?

A. Grant the team members the IAM role of logging.configWriter on Cloud IAM.
B. Configure Access Context Manager to allow only these members to export logs.
C. Create and grant a custom IAM role with the permissions logging.sinks.list and logging.sink.get.
D. Create an Organizational Policy in Cloud IAM to allow only these members to create log exports.

Answer: A

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

NEW QUESTION 41
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 involve
D. Determine who approved and pushed the new release to production.
E. Use the Git history to find the related code commi
F. Prevent the engineer who made that commit from working on production services.

Answer: B

NEW QUESTION 45
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 46
.....

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