Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

https://www.2passeasy.com/dumps/Professional-Cloud-DevOps-Engineer/

NEW QUESTION 1
Your company follows Site Reliability Engineering principles. You are writing a postmortem for an incident, triggered by a software change, that severely affected users. You want to prevent severe incidents from happening in the future. What should you do?

A. Identify engineers responsible for the incident and escalate to their senior management.
B. Ensure that test cases that catch errors of this type are run successfully before new software releases.
C. Follow up with the employees who reviewed the changes and prescribe practices they should follow in the future.
D. Design a policy that will require on-call teams to immediately call engineers and management to discuss a plan of action if an incident occurs.

Answer: B

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 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 lo 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 4
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 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 a trading application written in Python and hosted on App Engine flexible environment. You want to customize the error information being sent to Stackdriver Error Reporting. What should you do?

A. Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
B. Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
C. Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.
D. Use the Stackdriver Error Reporting API to write errors from your application to ReportedErrorEvent, and then generate log entries with properly formatted error messages in Stackdriver Logging.

Answer: D

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

NEW QUESTION 7
You are using Stackdriver to monitor applications hosted on Google Cloud Platform (GCP). You recently deployed a new application, but its logs are not appearing on the Stackdriver dashboard.
You need to troubleshoot the issue. What should you do?

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

Answer: A

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

NEW QUESTION 8
You 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 9
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 10
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 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 14
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 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 I grep fluentd

Answer: D

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

NEW QUESTION 19
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 21
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 26
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 throughout 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 30
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 35
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 38
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 42
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 43
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 48
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 50
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 51
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 56
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual Professional-Cloud-DevOps-Engineer Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the Professional-Cloud-DevOps-Engineer Product From:

https://www.2passeasy.com/dumps/Professional-Cloud-DevOps-Engineer/

Money Back Guarantee

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