Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam

NEW QUESTION 1
You have an application running in Google Kubernetes Engine. The application invokes multiple services per request but responds too slowly. You need to identify which downstream service or services are causing the delay. What should you do?

A. Analyze VPC flow logs along the path of the request.
B. Investigate the Liveness and Readiness probes for each service.
C. Create a Dataflow pipeline to analyze service metrics in real time.
D. Use a distributed tracing framework such as OpenTelemetry or Stackdriver Trace.

Answer: C

NEW QUESTION 2
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 3
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 4
You support a high-traffic web application and want to ensure that the home page loads in a timely manner. As a first step, you decide to implement a Service Level Indicator (SLI) to represent home page request latency with an acceptable page load time set to 100 ms. What is the Google-recommended way of calculating this SLI?

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

Answer: C

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

NEW QUESTION 5
You are running an application on Compute Engine and collecting logs through Stackdriver. You discover that some personally identifiable information (PII) is leaking into certain log entry fields. All PII entries begin with the text userinfo. You want to capture these log entries in a secure location for later review and prevent them from leaking to Stackdriver Logging. What should you do?

A. Create a basic log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.
B. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, and then copy the entries to a Cloud Storage bucket.
C. Create an advanced log filter matching userinfo, configure a log export in the Stackdriver console with Cloud Storage as a sink, and then configure a tog exclusion with userinfo as a filter.
D. Use a Fluentd filter plugin with the Stackdriver Agent to remove log entries containing userinfo, create an advanced log filter matching userinfo, and then configure a log export in the Stackdriver console with Cloud Storage as a sink.

Answer: B

Explanation:
https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e7

NEW QUESTION 6
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 7
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 Slackdriver 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 outpu
G. Run the script as a sidecar container with the application's po
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 8
You support a high-traffic web application with a microservice architecture. The home page of the application displays multiple widgets containing content such as the current weather, stock prices, and news headlines. The main serving thread makes a call to a dedicated microservice for each widget and then lays out the homepage for the user. The microservices occasionally fail; when that happens, the serving thread serves the homepage with some missing content. Users of the application are unhappy if this degraded mode occurs too frequently, but they would rather have some content served instead of no content at all. You want to set a Service Level Objective (SLO) to ensure that the user experience does not degrade too much. What Service Level Indicator {SLI) should you use to measure this?

A. A quality SLI: the ratio of non-degraded responses to total responses
B. An availability SLI: the ratio of healthy microservices to the total number of microservices
C. A freshness SLI: the proportion of widgets that have been updated within the last 10 minutes
D. A latency SLI: the ratio of microservice calls that complete in under 100 ms to the total number of microservice calls

Answer: B

Explanation:
https://cloud.google.com/blog/products/gcp/available-or-not-that-is-the-question-cre-life-lessons

NEW QUESTION 9
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 10
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 10
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 11
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 13
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 17
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 21
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 22
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 25
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 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 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 32
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 36
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 40
You need to reduce the cost of virtual machines (VM| for your organization. After reviewing different options, you decide to leverage preemptible VM instances. Which application is suitable for preemptible VMs?

A. A scalable in-memory caching system
B. The organization's public-facing website
C. A distributed, eventually consistent NoSQL database cluster with sufficient quorum
D. A GPU-accelerated video rendering platform that retrieves and stores videos in a storage bucket

Answer: D

Explanation:
https://cloud.google.com/compute/docs/instances/preemptible

NEW QUESTION 43
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 48
Your application services run in Google Kubernetes Engine (GKE). You want to make sure that only images from your centrally-managed Google Container Registry (GCR) image registry in the altostrat-images project can be deployed to the cluster while minimizing development time. What should you do?

A. Create a custom builder for Cloud Build that will only push images to gcr.io/altostrat-images.
B. Use a Binary Authorization policy that includes the whitelist name pattern gcr.io/attostrat-images/.
C. Add logic to the deployment pipeline to check that all manifests contain only images from gcr.io/altostrat-images.
D. Add a tag to each image in gcr.io/altostrat-images and check that this tag is present when the image is deployed.

Answer: B

NEW QUESTION 51
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 55
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 59
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 61
Your company experiences bugs, outages, and slowness in its production systems. Developers use the production environment for new feature development and bug fixes. Configuration and experiments are done in the production environment, causing outages for users. Testers use the production environment for load testing, which often slows the production systems. You need to redesign the environment to reduce the number of bugs and outages in production and to enable testers to load test new features. What should you do?

A. Create an automated testing script in production to detect failures as soon as they occur.
B. Create a development environment with smaller server capacity and give access only to developers and testers.
C. Secure the production environment to ensure that developers can't change it and set up one controlled update per year.
D. Create a development environment for writing code and a test environment for configurations, experiments, and load testing.

Answer: D

NEW QUESTION 62
You support an application running on GCP and want to configure SMS notifications to your team for the most critical alerts in Stackdriver Monitoring. You have already identified the alerting policies you want to configure this for. What should you do?

A. Download and configure a third-party integration between Stackdriver Monitoring and an SMS gateway.Ensure that your team members add their SMS/phone numbers to the external tool.
B. Select the Webhook notifications option for each alerting policy, and configure it to use a third-party integration too
C. Ensure that your team members add their SMS/phone numbers to the external tool.
D. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profil
E. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
F. Configure a Slack notification for each alerting polic
G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are receive
H. Ensure that your team members add their SMS/phone numbers to the external integration.

Answer: C

Explanation:
https://cloud.google.com/monitoring/support/notification-options#creating_channels To configure SMS notifications, do the following:
In the SMS section, click Add new and follow the instructions. Click Save. When you set up your alerting policy, select the SMS notification type and choose a verified phone number from the list.

NEW QUESTION 67
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 70
Your product is currently deployed in three Google Cloud Platform (GCP) zones with your users divided between the zones. You can fail over from one zone to another, but it causes a 10-minute service disruption for the affected users. You typically experience a database failure once per quarter and can detect it within five minutes. You are cataloging the reliability risks of a new real-time chat feature for your product. You catalog the following information for each risk:
• Mean Time to Detect (MUD} in minutes
• Mean Time to Repair (MTTR) in minutes
• Mean Time Between Failure (MTBF) in days
• User Impact Percentage
The chat feature requires a new database system that takes twice as long to successfully fail over between zones. You want to account for the risk of the new database failing in one zone. What would be the values for the risk of database failover with the new system?

A. MTTD: 5MTTR: 10MTBF: 90Impact: 33%
B. MTTD:5 MTTR: 20MTBF: 90Impact: 33%
C. MTTD:5 MTTR: 10MTBF: 90Impact 50%
D. MTTD:5 MTTR: 20MTBF: 90Impact: 50%

Answer: B

Explanation:
https://www.atlassian.com/incident-management/kpis/common-metrics https://linkedin.github.io/school-of-sre/

NEW QUESTION 75
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 77
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 78
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 82
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 86
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 91
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 92
You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem? (Choose two.)

A. An explanation of the root cause of the incident
B. A list of employees responsible for causing the incident
C. A list of action items to prevent a recurrence of the incident
D. Your opinion of the incident's severity compared to past incidents
E. Copies of the design documents for all the services impacted by the incident

Answer: AC

Explanation:
For a postmortem to be truly blameless, it must focus on identifying the contributing causes of the incident without indicting any individual or team for bad or inappropriate behavior.

NEW QUESTION 97
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 tim
B. Instruct developers to not store secrets at rest.
C. Store secrets in a separate configuration file on Gi
D. Provide select developers with access to the configuration file.
E. Store secrets in Cloud Storage encrypted with a key from Cloud KM
F. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
G. Encrypt the secrets and store them in the source code repositor
H. Store a decryption key in a separate repository and grant your pipeline access to it

Answer: C

NEW QUESTION 98
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 99
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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](https://www.certshared.com)