

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

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 log 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 3

You support a popular mobile game application deployed on Google Kubernetes Engine (GKE) across several Google Cloud regions. Each region has multiple Kubernetes clusters. You receive a report that none of the users in a specific region can connect to the application. You want to resolve the incident while following Site Reliability Engineering practices. What should you do first?

- A. Reroute the user traffic from the affected region to other regions that don't report issues.
- B. Use Stackdriver Monitoring to check for a spike in CPU or memory usage for the affected region.
- C. Add an extra node pool that consists of high memory and high CPU machine type instances to the cluster.
- D. Use Stackdriver Logging to filter on the clusters in the affected region, and inspect error messages in the logs.

Answer: A

Explanation:

Google always aims to first stop the impact of an incident, and then find the root cause (unless the root cause just happens to be identified early on).

NEW QUESTION 4

You have a set of applications running on a Google Kubernetes Engine (GKE) cluster, and you are using Stackdriver Kubernetes Engine Monitoring. You are bringing a new containerized application required by your company into production. This application is written by a third party and cannot be modified or reconfigured. The application writes its log information to `/var/log/app_messages.log`, and you want to send these log entries to Stackdriver Logging. What should you do?

- A. Use the default Stackdriver Kubernetes Engine Monitoring agent configuration.
- B. Deploy a Fluentd daemonset to GK
- C. Then create a customized input and output configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- D. Install Kubernetes on Google Compute Engine (GCE) and redeploy your application
- E. Then customize the built-in Stackdriver Logging configuration to tail the log file in the application's pods and write to Stackdriver Logging.
- F. Write a script to tail the log file within the pod and write entries to standard output
- G. Run the script as a sidecar container with the application's pod
- H. Configure a shared volume between the containers to allow the script to have read access to `/var/log` in the application container.

Answer: B

Explanation:

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

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

NEW QUESTION 5

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 6

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 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 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 9

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 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 15

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 19

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 22

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 27

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 28

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 32

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 35

You are creating and assigning action items in a postmodern for an outage. The outage is over, but you need to address the root causes. You want to ensure that your team handles the action items quickly and efficiently. How should you assign owners and collaborators to action items?

- A. Assign one owner for each action item and any necessary collaborators.
- B. Assign multiple owners for each item to guarantee that the team addresses items quickly
- C. Assign collaborators but no individual owners to the items to keep the postmortem blameless.
- D. Assign the team lead as the owner for all action items because they are in charge of the SRE team.

Answer: A

Explanation:

<https://devops.com/when-it-disaster-strikes-part-3-conducting-a-blameless-post-mortem/>

NEW QUESTION 36

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 37

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 40

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 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 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 44

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

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 49

You support the backend of a mobile phone game that runs on a Google Kubernetes Engine (GKE) cluster. The application is serving HTTP requests from users. You need to implement a solution that will reduce the network cost. What should you do?

- A. Configure the VPC as a Shared VPC Host project.
- B. Configure your network services on the Standard Tier.
- C. Configure your Kubernetes cluster as a Private Cluster.
- D. Configure a Google Cloud HTTP Load Balancer as Ingress.

Answer: D

Explanation:

Costs associated with a load balancer are charged to the project containing the load balancer components. Because of these benefits, container-native load balancing is the recommended solution for load balancing through Ingress. When NEGs are used with GKE Ingress, the Ingress controller facilitates the creation of all aspects of the L7 load balancer. This includes creating the virtual IP address, forwarding rules, health checks, firewall rules, and more.

<https://cloud.google.com/architecture/best-practices-for-running-cost-effective-kubernetes-applications-on-gke>

NEW QUESTION 54

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 Stacked 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 to METRIC_KIND_UNSPECIFIED and a 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

You need to deploy a new service to production. The service needs to automatically scale using a Managed Instance Group (MIG) and should be deployed over multiple regions. The service needs a large number of resources for each instance and you need to plan for capacity. What should you do?

- A. Use the n1-highcpu-96 machine type in the configuration of the MIG.
- B. Monitor results of Stackdriver Trace to determine the required amount of resources.
- C. Validate that the resource requirements are within the available quota limits of each region.
- D. Deploy the service in one region and use a global load balancer to route traffic to this region.

Answer: C

Explanation:

https://cloud.google.com/compute/quotas#understanding_quotas <https://cloud.google.com/compute/quotas>

NEW QUESTION 58

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

Answer: B

NEW QUESTION 59

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 62

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 it
- I. Attach the production projects to this workspace
- 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 66

You are performing a semiannual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP). using a Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

- A. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
- B. Because you are deployed on GKE and are using a cluster autoscaler
- C. your GKE cluster will scale automatically, regardless of growth rate.
- D. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
- E. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough capacity.

Answer: A

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler>

The Horizontal Pod Autoscaler changes the shape of your Kubernetes workload by automatically increasing or decreasing the number of Pods in response to the workload's CPU or memory consumption

NEW QUESTION 69

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 74

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Professional-Cloud-DevOps-Engineer Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-DevOps-Engineer Practice Test Here](https://www.surepassexam.com/Professional-Cloud-DevOps-Engineer-exam-dumps.html)