

Google.Professional-Cloud-DevOps-Engineer.v2022-01-25.q29

Exam Code:	Professional-Cloud-DevOps-Engineer
Exam Name:	Google Cloud Certified - Professional Cloud DevOps Engineer Exam
Certification Provider:	Google
Free Question Number:	29
Version:	v2022-01-25
# of views:	957
# of Questions views:	6695
https://www.freecram.net/torrent/Google.Professional-Cloud-DevOps-Engineer.v2022-01-25.q29.html	

NEW QUESTION: 1

Your application images are built using Cloud Build and pushed to Google Container Registry (GCR). You want to be able to specify a particular version of your application for deployment based on the release version tagged in source control. What should you do when you push the image?

- A. Reference the image digest in the source control tag.
- B. Supply the source control tag as a parameter within the image name.
- C. Use Cloud Build to include the release version tag in the application image.
- D. Use GCR digest versioning to match the image to the tag in source control.

Answer: (SHOW ANSWER)

<https://cloud.google.com/container-registry/docs/pushing-and-pulling>

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: ([SHOW ANSWER](#))

<https://medium.com/google-cloud/fluentd-filter-plugin-for-google-cloud-data-loss-prevention-api-42bbb1308e76>

NEW QUESTION: 3

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.** Review the Stackdriver monitoring.
- B.** Upsize the virtual machines running the login services.
- C.** Roll back the recent release.
- D.** Deploy a new release to see whether it fixes the problem.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 4

Your team is designing a new application for deployment both inside and outside Google Cloud Platform (GCP). You need to collect detailed metrics such as system resource utilization. You want to use centralized GCP services while minimizing the amount of work required to set up this collection system. What should you do?

- A.** Import the Stackdriver Debugger package, and configure the application to emit debug messages with timing information.
- B.** Import the Stackdriver Profiler package, and configure it to relay function timing data to Stackdriver for further analysis.
- C.** Install an Application Performance Monitoring (APM) tool in both locations, and configure an export to a central data storage location for analysis.
- D.** Instrument the code using a timing library, and publish the metrics via a health check endpoint that is scraped by Stackdriver.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 5

You support a multi-region web service running on Google Kubernetes Engine (GKE) behind a Global HTTP'S Cloud Load Balancer (CLB). For legacy reasons, user requests first go through a third-party Content Delivery Network (CDN). which then routes traffic to the CLB. You have already implemented an availability Service Level Indicator (SLI) at the

CLB level. However, you want to increase coverage in case of a potential load balancer misconfiguration, CDN failure, or other global networking catastrophe. Where should you measure this new SLI?

Choose 2 answers

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

Answer: (SHOW ANSWER)

NEW QUESTION: 6

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

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

Answer: D (LEAVE A REPLY)

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

NEW QUESTION: 7

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

Answer: B,E (LEAVE A REPLY)

NEW QUESTION: 8

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.** Create an Unmanaged Instance Group for the instances used to run the workload.
- D.** Convert the instances to preemptible virtual machines.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 9

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.** Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
- B.** File a bug with the development team so they can find the root cause of the crashing instance.
- C.** Create a Managed Instance Group with a single 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: ([SHOW ANSWER](#))

NEW QUESTION: 10

- A.** Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- B.** Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- 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: ([SHOW ANSWER](#))

NEW QUESTION: 11

You deploy a new release of an internal application during a weekend maintenance window when there is minimal user traffic. After the window ends, you learn that one of the new features isn't working as expected in the production environment. After an extended outage, you roll back the new release and deploy a fix. You want to modify your release process to reduce the mean time to recovery so you can avoid extended outages in the future. What should you do?

Choose 2 answers

- A.** Before merging new code, require 2 different peers to review the code changes.
- B.** Require developers to run automated integration tests on their local development environments before release.
- C.** Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
- D.** Adopt the blue/green deployment strategy when releasing new code via a CD server.
- E.** Configure a CI server. Add a suite of unit tests to your code and have your CI server run them on commit and verify any changes.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 12

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 13

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 a Development and a Production GKE cluster in separate projects. 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.
- B.** Create a Development and a Production GKE cluster in separate projects. 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.

C. Create one GCP Project per team. In each project, create a cluster for Development and one for Production. Grant the teams IAM access to their respective clusters.

D. Create one GCP Project per team. In each project, create a cluster with a Kubernetes namespace for Development and one for Production. Grant the teams IAM access to their respective clusters.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 14

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 port 25 has been opened in the firewall to allow messages through to Stackdriver.

B. Confirm that the Stackdriver agent has been installed in the hosting virtual machine.

C. Confirm that your account has the proper permissions to use the Stackdriver dashboard.

D. Confirm that the application is using the required client library and the service account key has proper permissions.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 15

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. Enable high availability on the CloudSQL instances.

B. Move the application from App Engine to Compute Engine.

C. Modify the App Engine configuration to have additional idle instances.

D. Upgrade the GCS buckets to Multi-Regional.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 16

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: ([SHOW ANSWER](#))

<https://cloud.google.com/build/docs/deploying-builds/deploy-gke>

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

Valid Professional-Cloud-DevOps-Engineer Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the ExamDiscuss.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

<https://www.examdisscuss.com/Google/exam/Professional-Cloud-DevOps-Engineer/premium/> (200 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 18

You are deploying an application that needs to access sensitive information. You need to ensure that this information is encrypted and the risk of exposure is minimal if a breach occurs. What should you do?

- A.** Store the encryption keys in Cloud Key Management Service (KMS) and rotate the keys frequently
- B.** Leverage a continuous build pipeline that produces multiple versions of the secret for each instance of the application.
- C.** Inject the secret at the time of instance creation via an encrypted configuration management system.
- D.** Integrate the application with a Single sign-on (SSO) system and do not expose secrets to the application

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 19

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 ([LEAVE A REPLY](#))

"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: 20

- A.** Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- B.** Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.
- C.** Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.
- D.** Click "Share chart by URL" and provide the URL to the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

Answer: ([SHOW ANSWER](#))

<https://cloud.google.com/monitoring/access-control>

NEW QUESTION: 21

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 minutes. Commit to "next update" times.
- B.** Provide periodic updates to all stakeholders in a timely manner. Commit to a "next update" time in all communications.
- C.** Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing responses directly to customers.
- D.** Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

Answer: ([SHOW ANSWER](#))

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: 22

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: ([SHOW ANSWER](#))

<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION: 23

Your organization wants to implement Site Reliability Engineering (SRE) culture and principles. Recently, a service that you support had a limited outage. A manager on

another team asks you to provide a formal explanation of what happened so they can action remediations. What should you do?

- A.** Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it with the manager only.
- B.** Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it on the engineering organization's document portal.
- C.** Develop a postmortem that includes the root causes, resolution, lessons learned, the list of people responsible, and a list of action items for each person. Share it with the manager only.
- D.** Develop a postmortem that includes the root causes, resolution, lessons learned, and a prioritized list of action items. Share it on the engineering organization's document portal.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 24

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 25

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. aggregate logs per container, and then analyze CPU and memory consumption.

Answer: ([SHOW ANSWER](#))

<https://cloud.google.com/anthos/clusters/docs/on-prem/1.7/concepts/logging-and-monitoring>

NEW QUESTION: 26

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. Trace the origin of 500 errors and the root cause of increased latency.
- D. Record data for the postmortem document of the incident.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 27

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: ([SHOW ANSWER](#))

<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: 28

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 Stackdriver 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: ([SHOW ANSWER](#))

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/README.md> 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: 29

You use Cloud Build to build your application. You want to reduce the build time while minimizing cost and development effort. What should you do?

A. Use Cloud Storage to cache intermediate artifacts.

B. Use multiple smaller build steps to minimize execution time.

C. Run multiple Jenkins agents to parallelize the build.

D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Answer: ([SHOW ANSWER](#))

Valid Professional-Cloud-DevOps-Engineer Dumps shared by ExamDiscuss.com for Helping Passing Professional-Cloud-DevOps-Engineer Exam! ExamDiscuss.com now offer the **newest Professional-Cloud-DevOps-Engineer exam dumps**, the ExamDiscuss.com Professional-Cloud-DevOps-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com Professional-Cloud-DevOps-Engineer dumps with Test Engine here:

[https://www.examdiscuss.com/Google/exam/Professional-Cloud-DevOps-](https://www.examdiscuss.com/Google/exam/Professional-Cloud-DevOps-Engineer/premium/)

[Engineer/premium/](https://www.examdiscuss.com/Google/exam/Professional-Cloud-DevOps-Engineer/premium/) (**200 Q&As Dumps, 35%OFF Special Discount Code: freecram**)