



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 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. Run multiple Jenkins agents to parallelize the build.
- C. Use multiple smaller build steps to minimize execution time.
- D. Use larger Cloud Build virtual machines (VMs) by using the machine-type option.

Answer: C

Explanation:

<https://cloud.google.com/storage/docs/best-practices>

https://cloud.google.com/build/docs/speeding-up-builds#caching_directories_with_google_cloud_storage Caching directories with Google Cloud Storage To

increase the speed of a build, reuse the results from a

previous build. You can copy the results of a previous build to a Google Cloud Storage bucket, use the results for faster calculation, and then copy the new results back to the bucket. Use this method when your build takes a long time and produces a small number of files that does not take time to copy to and from Google Cloud Storage.

upvoted 2 times

NEW QUESTION 3

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: B

Explanation:

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

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. Bucketize the 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 encountered a major service outage that affected all users of the service for multiple hours. After several hours of incident management, the service returned to normal, and user access was restored. You need to provide an incident summary to relevant stakeholders following the Site Reliability Engineering recommended practices. What should you do first?

- A. Call individual stakeholders to explain what happened.
- B. Develop a post-mortem to be distributed to stakeholders.
- C. Send the Incident State Document to all the stakeholders.
- D. Require the engineer responsible to write an apology email to all stakeholders.

Answer: B

NEW QUESTION 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 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 8

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 9

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 need to run a business-critical workload on a fixed set of Compute Engine instances for several months. The workload is stable with the exact amount of resources allocated to it. You want to lower the costs for this workload without any performance implications. What should you do?

- A. Purchase Committed Use Discounts.
- B. Migrate the instances to a Managed Instance Group.
- C. Convert the instances to preemptible virtual machines.
- D. Create an Unmanaged Instance Group for the instances used to run the workload.

Answer: A

NEW QUESTION 10

You 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 12

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

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

Answer: B

NEW QUESTION 15

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 17

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 Profile
- 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 policy
- G. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received
- 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 21

You support a stateless web-based API that is deployed on a single Compute Engine instance in the europe-west2-a zone. The Service Level Indicator (SLI) for service availability is below the specified Service Level Objective (SLO). A postmortem has revealed that requests to the API regularly time out. The time outs are due to the API having a high number of requests and running out of memory. You want to improve service availability. What should you do?

- A. Change the specified SLO to match the measured SLI.
- B. Move the service to higher-specification compute instances with more memory.
- C. Set up additional service instances in other zones and load balance the traffic between all instances.
- D. Set up additional service instances in other zones and use them as a failover in case the primary instance is unavailable.

Answer: C

NEW QUESTION 26

Your application images are built and pushed to Google Container Registry (GCR). You want to build an automated pipeline that deploys the application when the image is updated while minimizing the development effort. What should you do?

- A. Use Cloud Build to trigger a Spinnaker pipeline.
- B. Use Cloud Pub/Sub to trigger a Spinnaker pipeline.
- C. Use a custom builder in Cloud Build to trigger a Jenkins pipeline.
- D. Use Cloud Pub/Sub to trigger a custom deployment service running in Google Kubernetes Engine(GKE).

Answer: B

Explanation:

<https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud> <https://spinnaker.io/guides/user/pipeline/triggers/pubsub/>

NEW QUESTION 29

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

Answer: B

NEW QUESTION 32

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 33

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

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

Answer: B

Explanation:

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

NEW QUESTION 36

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 NEG's 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 40

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

Answer: BE

NEW QUESTION 41

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

Answer: C

NEW QUESTION 46

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 50

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 53

.....

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](#)