

Latest [Apr 15, 2023] Real Google Professional-Cloud-DevOps-Engineer Exam Dumps Questions [Q47-Q68]



Latest [Apr 15, 2023] Real Google Professional-Cloud-DevOps-Engineer Exam Dumps Questions [Q47-Q68]

Latest [Apr 15, 2023] Real Google Professional-Cloud-DevOps-Engineer Exam Dumps Questions
Professional-Cloud-DevOps-Engineer Dumps To Pass Cloud DevOps Engineer Exam in One Day (Updated 82 Questions)

Google Professional Cloud DevOps Engineer Certification Path

The Google Professional Cloud DevOps Engineer Certification is the highest level of certification mainly focusing on the Solution Architect Professional. There is no prerequisite for this exam but still, it would be best to follow some sequence to prove immense knowledge as a Google Professional Cloud DevOps Engineer.

Domains Tested in Final Evaluation The candidates who succeed in passing the Google Professional Cloud DevOps Engineer validation will add an international designation under their belt that helps them consolidate their position in the company and receive generous offers from their employers. So, if applicants want to differentiate themselves from other colleagues, they should demonstrate the following skills: - Identifying issues related to service performance and implementing solutions to optimize it;- Identifying the most effective principles related to site reliability engineering and applying them to a definite service;- Checking and implementing strategies related to service monitoring;- Coordinating teams to identify and manage incidents related to service development.- Implementing and building CI/CD service pipelines;

Google Professional-Cloud-DevOps-Engineer Exam Syllabus Topics:

TopicDetailsTopic 1- Implementing service monitoring strategies- Develop plan to optimize areas of greatest cost or lowest utilizationTopic 2- Create environments dynamically per feature branch with GKE, Cloud Deployment Manager- Collecting third-party and structured logs with Stackdriver Logging, FluentdTopic 3- Manage different development environments- Vulnerability analysis with Container RegistryTopic 4- Set ACL to allow metric writing for custom metrics with IAM, Stackdriver Monitoring- Building and implementing CI- CD pipelines for a serviceTopic 5- Testing a new version with Spinnaker- Decide on the number of environments and their purposeTopic 6- Perform investigation to isolate most likely actual cause- Applying site reliability engineering principles to a serviceTopic 7- Artifact repositories with Container Registry- Balance change, velocity, and reliability of the service

NEW QUESTION 47

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.)

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

NEW QUESTION 48

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?

- * File a bug with the development team so they can find the root cause of the crashing instance.
- * Create a Managed Instance Group with a single instance and use health checks to determine the system status.
- * Add a Load Balancer in front of the Compute Engine instance and use health checks to determine the system status.
- * Create a Stackdriver Monitoring dashboard with SMS alerts to be able to start recreating the crashed instance promptly after it has crashed.

NEW QUESTION 49

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?

- * Publish various metrics from the application directly to the Stackdriver Monitoring API, and then observe these custom metrics in Stackdriver.
- * Install the Cloud Pub/Sub client libraries, push various metrics from the application to various topics, and then observe the aggregated metrics in Stackdriver.
- * 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.
- * 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.

NEW QUESTION 50

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?

- * Install the Stackdriver Error Reporting library for Python, and then run your code on a Compute Engine VM.
- * Install the Stackdriver Error Reporting library for Python, and then run your code on Google Kubernetes Engine.
- * Install the Stackdriver Error Reporting library for Python, and then run your code on App Engine flexible environment.
- * 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.

<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 51

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?

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

<https://cloud.google.com/architecture/continuous-delivery-toolchain-spinnaker-cloud>

<https://spinnaker.io/guides/user/pipeline/triggers/pubsub/>

NEW QUESTION 52

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?

- * flex/connections/current
- * tcp_ssl_proxy/new_connections
- * tcp_ssl_proxy/open_connections
- * flex/instance/connections/current

Explanation/Reference: https://cloud.google.com/monitoring/api/metrics_gcp

NEW QUESTION 53

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?

- * Upgrade the GCS buckets to Multi-Regional.
- * Enable high availability on the CloudSQL instances.
- * Move the application from App Engine to Compute Engine.
- * Modify the App Engine configuration to have additional idle instances.

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 54

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?

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

NEW QUESTION 55

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?

- * 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.
- * 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.
- * 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.
- * 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.

https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#roles_and_groups

NEW QUESTION 56

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?

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

NEW QUESTION 57

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?

- * Create a Cloud Storage bucket and use the built-in encryption at rest. Store the secrets in the bucket and grant Cloud Build access to the bucket.
- * Encrypt the secrets and store them in the application repository. Store a decryption key in a separate repository and grant Cloud Build access to the repository.
- * Use client-side encryption to encrypt the secrets and store them in a Cloud Storage bucket. Store a decryption key in the bucket and grant Cloud Build access to the bucket.
- * Use Cloud Key Management Service (Cloud KMS) to encrypt the secrets and include them in your Cloud Build deployment configuration. Grant Cloud Build access to the KeyRing.

NEW QUESTION 58

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?

- * 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 to a versioned Cloud Storage bucket as the latest master version.
- * 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.
- * 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.
- * Rename the folder structure with a predefined naming convention that increments the version.
- * 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.

NEW QUESTION 59

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?

- * Add logic to each Cloud Build step to HTTP POST the build information to a webhook.
- * Add a new step at the end of the pipeline in Cloud Build to HTTP POST the build information to a webhook.
- * Use Stackdriver Logging to create a logs-based metric from the Cloud Build logs. Create an Alert with a Webhook notification type.
- * Create a Cloud Pub/Sub push subscription to the Cloud Build cloud-builds PubSub topic to HTTP POST the build information to a webhook.

NEW QUESTION 60

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?

- * Prompt developers for secrets at build time. Instruct developers to not store secrets at rest.
- * Store secrets in a separate configuration file on Git. Provide select developers with access to the configuration file.
- * Store secrets in Cloud Storage encrypted with a key from Cloud KMS. Provide the CI/CD pipeline with access to Cloud KMS via IAM.
- * 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

NEW QUESTION 61

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?

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

NEW QUESTION 62

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?

- * Assign the Container Developer role to the Cloud Build service account.
- * Specify the Container Developer role for Cloud Build in the cloudbuild.yaml file.
- * Create a new service account with the Container Developer role and use it to run Cloud Build.
- * Create a separate step in Cloud Build to retrieve service account credentials and pass these to kubectl.

NEW QUESTION 63

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?

- * * Deploy the Stackdriver logging agent to the application servers.

- * Give the developers the IAM Logs Viewer role to access Stackdriver and view logs.
- * * 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.
- * * Deploy the Stackdriver monitoring agent to the application servers.

- * Give the developers the IAM Monitoring Viewer role to access Stackdriver and view metrics.
- * * 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.

NEW QUESTION 64

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?

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

NEW QUESTION 65

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

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

NEW QUESTION 66

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?

- * Assign one owner for each action item and any necessary collaborators.

- * Assign multiple owners for each item to guarantee that the team addresses items quickly
- * Assign collaborators but no individual owners to the items to keep the postmortem blameless.
- * Assign the team lead as the owner for all action items because they are in charge of the SRE team.

NEW QUESTION 67

You are part of an organization that follows SRE practices and principles. You are taking over the management of a new service from the Development Team, and you conduct a Production Readiness Review (PRR). After the PRR analysis phase, you determine that the service cannot currently meet its Service Level Objectives (SLOs). You want to ensure that the service can meet its SLOs in production. What should you do next?

- * Adjust the SLO targets to be achievable by the service so you can bring it into production.
- * Notify the development team that they will have to provide production support for the service.
- * Identify recommended reliability improvements to the service to be completed before handover.
- * Bring the service into production with no SLOs and build them when you have collected operational data.

NEW QUESTION 68

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?

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

Professional-Cloud-DevOps-Engineer Exam Brain Dumps - Study Notes and Theory:
<https://www.validbraindumps.com/Professional-Cloud-DevOps-Engineer-exam-prep.html>