

[Dec 26, 2023 Passing Key To Getting Professional-Cloud-DevOps-Engineer Certified Exam Engine PDF [Q25-Q42]



[Dec 26, 2023] Passing Key To Getting Professional-Cloud-DevOps-Engineer Certified Exam Engine PDF
Professional-Cloud-DevOps-Engineer Exam Dumps Pass with Updated Dec-2023 Tests Dumps

Google Professional-Cloud-DevOps-Engineer certification exam is a highly respected certification in the field of cloud computing. It is designed to validate the skills and knowledge of professionals who are responsible for building and managing cloud-based solutions using Google Cloud Platform. Google Cloud Certified - Professional Cloud DevOps Engineer Exam certification exam requires a deep understanding of DevOps practices, as well as a comprehensive knowledge of Google Cloud Platform tools and services.

NEW QUESTION 25

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?

- * Bucketize the request latencies into ranges, and then compute the percentile at 100 ms.
- * Bucketize the request latencies into ranges, and then compute the median and 90th percentiles.

- * Count the number of home page requests that load in under 100 ms, and then divide by the total number of home page requests.
- * Count the number of home page requests that load in under 100 ms. and then divide by the total number of all web application requests.

NEW QUESTION 26

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?

- * Push the images to Google Container Registry (GCR) using the gcr.io hostname.
- * Push the images to Google Container Registry (GCR) using the us.gcr.io hostname.
- * Push the images to Google Container Registry (GCR) using the eu.gcr.io hostname.
- * Push the images to a private image registry running on a Compute Engine instance in the eu-west-1 region.

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 27

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.)

- * Make the service's SLO more strict.
- * Increase the service's deployment velocity and/or risk.
- * Shift engineering time to other services that need more reliability.
- * Get the product team to prioritize reliability work over new features.
- * Change the implementation of your Service Level Indicators (SLIs) to increase coverage.

(<https://sre.google/workbook/implementing-slos/#slo-decision-matrix>)

NEW QUESTION 28

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?

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

NEW QUESTION 29

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?

- * Focus on responding to internal stakeholders at least every 30 minutes. Commit to a next update times.
- * Provide periodic updates to all stakeholders in a timely manner. Commit to a next update time in all communications.
- * Delegate the responding to internal stakeholder emails to another member of the Incident Response Team. Focus on providing

responses directly to customers.

* Provide all internal stakeholder emails to the Incident Commander, and allow them to manage internal communications. Focus on providing responses directly to customers.

NEW QUESTION 30

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 31

You are writing a postmortem for an incident that severely affected users. You want to prevent similar incidents in the future. Which two of the following sections should you include in the postmortem? (Choose two.)

* An explanation of the root cause of the incident

* A list of employees responsible for causing the incident

* A list of action items to prevent a recurrence of the incident

* Your opinion of the incident's severity compared to past incidents

* Copies of the design documents for all the services impacted by the incident

NEW QUESTION 32

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?

* Share the workspace Project ID with the SRE team. Assign the SRE team the Monitoring Viewer IAM role in the workspace project.

* Share the workspace Project ID with the SRE team. Assign the SRE team the Dashboard Viewer IAM role in the workspace project.

* 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.

* 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.

Explanation

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

NEW QUESTION 33

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.

NEW QUESTION 34

You support a user-facing web application. When analyzing the application's error budget over the previous six months, you notice that the application never consumed more than 5% of its error budget. You hold a SLO review with business stakeholders and confirm that the SLO is set appropriately. You want your application's reliability to more closely reflect its SLO. What steps can you take to further that goal while balancing velocity, reliability, and business needs?

Choose 2 answers

- * Add more serving capacity to all of your application's zones
- * Implement and measure all other available SLIs for the application
- * Announce planned downtime to consume more error budget and ensure that users are not depending on a tighter SLO
- * Have more frequent or potentially risky application releases
- * Tighten the SLO to match the application's observed reliability

Explanation

The best options for furthering your application's reliability goal while balancing velocity, reliability, and business needs are to have more frequent or potentially risky application releases and to tighten the SLO to match the application's observed reliability. Having more frequent or potentially risky application releases can help you increase the change velocity and deliver new features faster. However, this also increases the likelihood of consuming more error budget and reducing the reliability of your service. Therefore, you should monitor your error budget consumption and adjust your release policies accordingly. For example, you can freeze or slow down releases when the error budget is low, or accelerate releases when the error budget is high. Tightening the SLO to match the application's observed reliability can help you align your service quality with your users' expectations and business needs. However, this also means that you have less room for error and need to maintain a higher level of reliability. Therefore, you should ensure that your SLO is realistic and achievable, and that you have sufficient engineering resources and processes to meet it.

NEW QUESTION 35

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

- * Before merging new code, require 2 different peers to review the code changes.
- * Adopt the blue/green deployment strategy when releasing new code via a CD server.
- * Integrate a code linting tool to validate coding standards before any code is accepted into the repository.
- * Require developers to run automated integration tests on their local development environments before release.
- * 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.

NEW QUESTION 36

You are configuring the frontend tier of an application deployed in Google Cloud. The frontend tier is hosted in gcp and deployed using a managed instance group with an Envoy-based external HTTP(S) load balancer in front. The application is deployed entirely within the europe-west2 region and only serves users based in the United Kingdom. You need to choose the most cost-effective

network tier and load balancing configuration What should you use?

- * Premium Tier with a global load balancer
- * Premium Tier with a regional load balancer
- * Standard Tier with a global load balancer
- * Standard Tier with a regional load balancer

Explanation

The most cost-effective network tier and load balancing configuration for your frontend tier is to use Premium Tier with a regional load balancer. Premium Tier is a network tier that provides high-performance and low-latency network connectivity across Google's global network. A regional load balancer is a load balancer that distributes traffic within a single region. Since your application is deployed entirely within the europe-west2 region and only serves users based in the United Kingdom, you can use Premium Tier with a regional load balancer to optimize the network performance and cost.

NEW QUESTION 37

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 38

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?

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

NEW QUESTION 39

You built a serverless application by using Cloud Run and deployed the application to your production environment. You want to identify the resource utilization of the application for cost optimization. What should you do?

- * Use Cloud Trace with distributed tracing to monitor the resource utilization of the application
- * Use Cloud Profiler with Ops Agent to monitor the CPU and memory utilization of the application
- * Use Cloud Monitoring to monitor the container CPU and memory utilization of the application
- * Use Cloud Ops to create logs-based metrics to monitor the resource utilization of the application

Explanation

The best option for giving developers the ability to test the latest revisions of the service before the service is exposed to customers is to run the `gcloud run deploy booking-engine --no-traffic --tag dev` command and use the `https://dev-booking-engine-abcdef.a.run.app` URL for testing. The `gcloud run deploy` command is a command that deploys a new revision of your service or updates an existing service. By using the `--no-traffic` flag, you can prevent any traffic from being sent to the new revision. By using the `--tag` flag, you can assign a tag to the new revision, such as `dev`. This way, you can create a new revision of your service without affecting your customers. You can also use the tag-based URL (e.g.,

<https://dev—booking-engine-abcdef.a.run.app>) to access and test the new revision.

NEW QUESTION 40

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. You want to prevent these fields from being written in new log entries as quickly as possible. What should you do?

- * Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- * Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- * Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- * Stage log entries to Cloud Storage, and then trigger a Cloud Function to remove the fields and write the entries to Stackdriver via the Stackdriver Logging API.

NEW QUESTION 41

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?

- * Roll back the experimental canary release.
- * Start monitoring latency, traffic, errors, and saturation.
- * Record data for the postmortem document of the incident.
- * Trace the origin of 500 errors and the root cause of increased latency.

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?

- * 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.

Explanation

<https://cloud.google.com/build/docs/securing-builds/use-encrypted-credentials>

Earning the Google Professional-Cloud-DevOps-Engineer certification can demonstrate to employers that an individual has the

skills and knowledge needed to effectively manage cloud technologies with DevOps practices. Google Cloud Certified - Professional Cloud DevOps Engineer Exam certification can also lead to better job opportunities and higher salaries for cloud professionals.

Professional-Cloud-DevOps-Engineer exam questions for practice in 2023 Updated 164 Questions:

<https://www.examlabs.com/Google/Cloud-DevOps-Engineer/best-Professional-Cloud-DevOps-Engineer-exam-dumps.html>