

Google

Exam Questions Professional-Cloud-DevOps-Engineer

Google Cloud Certified - Professional Cloud DevOps Engineer Exam



NEW QUESTION 1

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?

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

Answer: C

NEW QUESTION 2

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 3

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?

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

Answer: D

Explanation:

<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 4

You need to define Service Level Objectives (SLOs) for a high-traffic multi-region web application. Customers expect the application to always be available and have fast response times. Customers are currently happy with the application performance and availability. Based on current measurement, you observe that the 90th percentile of latency is 120ms and the 95th percentile of latency is 275ms over a 28-day window. What latency SLO would you recommend to the team to publish?

- A. 90th percentile – 100ms 95th percentile – 250ms
- B. 90th percentile – 120ms 95th percentile – 275ms
- C. 90th percentile – 150ms 95th percentile – 300ms
- D. 90th percentile – 250ms 95th percentile – 400ms

Answer: C

Explanation:

<https://sre.google/sre-book/service-level-objectives/>

NEW QUESTION 5

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?

- A. Use the filter-record-transformer Fluentd filter plugin to remove the fields from the log entries in flight.
- B. Use the fluent-plugin-record-reformer Fluentd output plugin to remove the fields from the log entries in flight.
- C. Wait for the application developers to patch the application, and then verify that the log entries are no longer exposing PII.
- D. 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.

Answer: A

NEW QUESTION 6

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 Slackdriver 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: A

Explanation:

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/REA> 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 7

You support a large service with a well-defined Service Level Objective (SLO). The development team deploys new releases of the service multiple times a week. If a major incident causes the service to miss its SLO, you want the development team to shift its focus from working on features to improving service reliability. What should you do before a major incident occurs?

- A. Develop an appropriate error budget policy in cooperation with all service stakeholders.
- B. Negotiate with the product team to always prioritize service reliability over releasing new features.
- C. Negotiate with the development team to reduce the release frequency to no more than once a week.
- D. Add a plugin to your Jenkins pipeline that prevents new releases whenever your service is out of SLO.

Answer: A

Explanation:

Reason : Incident has not occurred yet, even when development team is already pushing new features multiple times a week. The option A says, to define an error budget "policy", not to define error budget (It is already present). Just simple means to bring in all stakeholders, and decide how to consume the error budget effectively that could bring balance between feature deployment and reliability.

The goals of this policy are to: -- Protect customers from repeated SLO misses -- Provide an incentive to balance reliability with other features

<https://sre.google/workbook/error-budget-policy/>

NEW QUESTION 8

You created a Stackdriver chart for CPU utilization in a dashboard within your workspace project. You want to share the chart with your Site Reliability Engineering (SRE) team only. You want to ensure you follow the principle of least privilege. What should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 9

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
- 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 items
- 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 person
- 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 person
- H. Share it on the engineering organization's document portal.

Answer: B

NEW QUESTION 10

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 10

Your team of Infrastructure DevOps Engineers is growing, and you are starting to use Terraform to manage infrastructure. You need a way to implement code versioning and to share code with other team members. What should you do?

- A. Store the Terraform code in a version-control system
- B. Establish procedures for pushing new versions and merging with the master.
- C. Store the Terraform code in a network shared folder with child folders for each version release
- D. Ensure that everyone works on different files.
- E. Store the Terraform code in a Cloud Storage bucket using object versioning
- F. Give access to the bucket to every team member so they can download the files.
- G. Store the Terraform code in a shared Google Drive folder so it syncs automatically to every team member's computer
- H. Organize files with a naming convention that identifies each new version.

Answer: A

Explanation:

<https://www.terraform.io/docs/cloud/guides/recommended-practices/part3.3.html>

NEW QUESTION 15

Your development team has created a new version of their service's API. You need to deploy the new versions of the API with the least disruption to third-party developers and end users of third-party installed applications. What should you do?

- A. Introduce the new version of the API. Announce deprecation of the old version of the API
- B. Deprecate the old version of the API. Contact remaining users of the old API. Provide best effort support to users of the old API
- C. Turn down the old version of the API.
- D. Announce deprecation of the old version of the API
- E. Introduce the new version of the API. Contact remaining users on the old API
- F. Deprecate the old version of the API
- G. Turn down the old version of the API. Provide best effort support to users of the old API.
- H. Announce deprecation of the old version of the API
- I. Contact remaining users on the old API. Introduce the new version of the API
- J. Deprecate the old version of the API. Provide best effort support to users of the old API
- K. Turn down the old version of the API.
- L. Introduce the new version of the API
- M. Contact remaining users of the old API. Announce deprecation of the old version of the API
- N. Deprecate the old version of the API. Turn down the old version of the API. Provide best effort support to users of the old API.

Answer: A

NEW QUESTION 20

You have migrated an e-commerce application to Google Cloud Platform (GCP). You want to prepare the application for the upcoming busy season. What should you do first to prepare for the busy season?

- A. Load test the application to profile its performance for scaling.
- B. Enable AutoScaling on the production clusters, in case there is growth.
- C. Pre-provision double the compute power used last season, expecting growth.
- D. Create a runbook on inflating the disaster recovery (DR) environment if there is growth.

Answer: A

Explanation:

<https://cloud.google.com/blog/topics/retail/preparing-for-peak-holiday-season-while-wfh>

NEW QUESTION 24

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: A

Explanation:

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

NEW QUESTION 27

You are managing an application that exposes an HTTP endpoint without using a load balancer. The latency of the HTTP responses is important for the user experience. You want to understand what HTTP latencies all of your users are experiencing. You use Stackdriver Monitoring. What should you do?

- A. • In your application, create a metric with a metricKind set to DELTA and a valueType set to DOUBLE. • In Stackdriver's Metrics Explorer, use a Slacked Bar graph to visualize the metric.

B. • In your application, create a metric with a metricKind set to CUMULATIVE and a valueType set to DOUBLE. • In Stackdriver's Metrics Explorer, use a Line graph to visualize the metric.

C. • In your application, create a metric with a metricKind set to gauge and a valueType set to distribution. • In Stackdriver's Metrics Explorer, use a Heatmap graph to visualize the metric.

D. • In your application, create a metric with a metricKind set to METRIC_KIND_UNSPECIFIED and a valueType set to INT64. • In Stackdriver's Metrics Explorer, use a Stacked Area graph to visualize the metric.

Answer: C

Explanation:

<https://sre.google/workbook/implementing-slos/> <https://cloud.google.com/architecture/adopting-slos/>
Latency is commonly measured as a distribution. Given a distribution, you can measure various percentiles.
For example, you might measure the number of requests that are slower than the historical 99th percentile.

NEW QUESTION 30

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.)

- A. An explanation of the root cause of the incident
- B. A list of employees responsible for causing the incident
- C. A list of action items to prevent a recurrence of the incident
- D. Your opinion of the incident's severity compared to past incidents
- E. Copies of the design documents for all the services impacted by the incident

Answer: AC

Explanation:

For a postmortem to be truly blameless, it must focus on identifying the contributing causes of the incident without indicting any individual or team for bad or inappropriate behavior.

NEW QUESTION 32

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 37

.....

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](#)