



Free Questions for **Professional-Cloud-DevOps-Engineer** by **go4braindumps**

Shared by **Carson** on **06-06-2022**

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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.)

Options:

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

Answer:

D, E

Question 2

Question Type: MultipleChoice

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.)

Options:

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

Answer:

A, D

Question 3

Question Type: MultipleChoice

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?

Options:

- 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. Determine who approved and pushed the new release to production.
- D-** Use the Git history to find the related code commit. Prevent the engineer who made that commit from working on production services.

Answer:

B

Question 4

Question Type: MultipleChoice

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?

Options:

- 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:

B

Question 5

Question Type: MultipleChoice

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?

Options:

A- Introduce the new version of the API.

Announce deprecation of the old version of the API.

Deprecate the old version of the API.

Contact remaining users of the old API.

Provide best effort support to users of the old API.

Turn down the old version of the API.

B- Announce deprecation of the old version of the API.

Introduce the new version of the API.

Contact remaining users on the old API.

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.

C- Announce deprecation of the old version of the API.

Contact remaining users on the old API.

Introduce the new version of the API.

Deprecate the old version of the API.

Provide best effort support to users of the old API.

Turn down the old version of the API.

- D-** Introduce the new version of the API.
- Contact remaining users of the old API.
- Announce deprecation of the old version of the API.
- 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:

B

Question 6

Question Type: MultipleChoice

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?

Options:

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

Answer:

A

Question 7

Question Type: MultipleChoice

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?

Options:

- 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

Question 8

Question Type: MultipleChoice

You are responsible for the reliability of a high-volume enterprise application. A large number of users report that an important subset of the application's functionality -- a data intensive reporting feature -- is consistently failing with an HTTP 500 error. When you investigate your application's dashboards, you notice a strong correlation between the failures and a metric that represents the size of an internal queue used for generating reports. You trace the failures to a reporting backend that is experiencing high I/O wait times. You quickly fix the issue by resizing the backend's persistent disk (PD). How you need to create an availability Service Level Indicator (SLI) for the report generation feature. How would you define it?

Options:

A- As the I/O wait times aggregated across all report generation backends

- B-** As the proportion of report generation requests that result in a successful response
- C-** As the application's report generation queue size compared to a known-good threshold
- D-** As the reporting backend PD throughout capacity compared to a known-good threshold

Answer:

C

Question 9

Question Type: MultipleChoice

You are ready to deploy a new feature of a web-based application to production. You want to use Google Kubernetes Engine (GKE) to perform a phased rollout to half of the web server pods.

What should you do?

Options:

- A-** Use a partitioned rolling update.
- B-** Use Node taints with NoExecute.

- C- Use a replica set in the deployment specification.
- D- Use a stateful set with parallel pod management policy.

Answer:

A

Question 10

Question Type: MultipleChoice

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.)

Options:

- 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:

A, B

**To Get Premium Files for Professional-Cloud-DevOps-Engineer
Visit**

<https://www.p2pexams.com/products/professional-cloud-devops-engineer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-devops-engineer>

