

**100%** Money Back  
**Guarantee**

**Vendor:**Google

**Exam**

**Code:**PROFESSIONAL-CLOUD-DEVOPS-ENGINEER

**Exam Name:**Professional Cloud DevOps Engineer

**Version:**Demo

### QUESTION 1

You are deploying an application to Cloud Run. The application requires a password to start. Your organization requires that all passwords are rotated every 24 hours, and your application must have the latest password. You need to deploy the application with no downtime. What should you do?

- A. Store the password in Secret Manager and send the secret to the application by using environment variables.
- B. Store the password in Secret Manager and mount the secret as a volume within the application.
- C. Use Cloud Build to add your password into the application container at build time. Ensure that Artifact Registry is secured from public access.
- D. Store the password directly in the code. Use Cloud Build to rebuild and deploy the application each time the password changes.

Correct Answer: A

The best solution is to store the password in Secret Manager and send the secret to the application by using environment variables. This will allow you to rotate the password without having to rebuild and deploy the application each time.

---

### QUESTION 2

As a Site Reliability Engineer, you support an application written in Go that runs on Google Kubernetes Engine (GKE) in production. After releasing a new version of the application, you notice the application runs for about 15 minutes and then restarts. You decide to add Cloud Profiler to your application and now notice that the heap usage grows constantly until the application restarts. What should you do?

- A. Increase the CPU limit in the application deployment.
- B. Add high memory compute nodes to the cluster.
- C. Increase the memory limit in the application deployment.
- D. Add Cloud Trace to the application, and redeploy.

Correct Answer: C

---

### QUESTION 3

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

Correct Answer: A

A: Metric for App engine and for version

B: NO, Metrics for Cloud Load Balancing.

C: NO, Metrics for Cloud Load Balancing. D Metric for App engine and for instance

An App Engine app is made up of a single application resource that consists of one or more services. Each service can be configured to use different runtimes and to operate with different performance settings. Within each service, you deploy versions of that service. Each version then runs within one or more instances, depending on how much traffic you configured it to handle.

if the version runs within one or more instances we need for the version.

---

#### QUESTION 4

You recently noticed that one of your services has exceeded the error budget for the current rolling window period. Your company's product team is about to launch a new feature. You want to follow Site Reliability Engineering (SRE) practices. What should you do?

A. Notify the team about the lack of error budget and ensure that all their tests are successful so the launch will not further risk the error budget

B. Notify the team that their error budget is used up. Negotiate with the team for a launch freeze or tolerate a slightly worse user experience.

C. Escalate the situation and request additional error budget.

D. Look through other metrics related to the product and find SLOs with remaining error budget. Reallocate the error budgets and allow the feature launch.

Correct Answer: B

Negotiate with the team for a launch freeze or tolerate a slightly worse user experience.

---

#### QUESTION 5

Your company runs services by using multiple globally distributed Google Kubernetes Engine (GKE) clusters. Your operations team has set up workload monitoring that uses Prometheus-based tooling for metrics, alerts, and generating dashboards. This setup does not provide a method to view metrics globally across all clusters. You need to implement a scalable solution to support global Prometheus querying and minimize management overhead. What should you do?

A. Configure Prometheus cross-service federation for centralized data access.

B. Configure workload metrics within Cloud Operations for GKE.

C. Configure Prometheus hierarchical federation for centralized data access.

D. Configure Google Cloud Managed Service for Prometheus.

Correct Answer: D

<https://cloud.google.com/stackdriver/docs/managed-prometheus>

---

### QUESTION 6

You are developing the deployment and testing strategies for your CI/CD pipeline in Google Cloud. You must be able to:

? Reduce the complexity of release deployments and minimize the duration of deployment rollbacks.

? Test real production traffic with a gradual increase in the number of affected users.

You want to select a deployment and testing strategy that meets your requirements. What should you do?

- A. Recreate deployment and canary testing
- B. Blue/green deployment and canary testing
- C. Rolling update deployment and A/B testing
- D. Rolling update deployment and shadow testing

Correct Answer: B

---

### QUESTION 7

You are working with a government agency that requires you to archive application logs for seven years. You need to configure Stackdriver to export and store the logs while minimizing costs of storage. What should you do?

- A. Create a Cloud Storage bucket and develop your application to send logs directly to the bucket.
- B. Develop an App Engine application that pulls the logs from Stackdriver and saves them in BigQuery.
- C. Create an export in Stackdriver and configure Cloud Pub/Sub to store logs in permanent storage for seven years.
- D. Create a sink in Stackdriver, name it, create a bucket on Cloud Storage for storing archived logs, and then select the bucket as the log export destination.

Correct Answer: D

<https://cloud.google.com/logging/docs/routing/overview>

---

### QUESTION 8

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. 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.

Correct Answer: A

---

## QUESTION 9

You are analyzing Java applications in production. All applications have Cloud Profiler and Cloud Trace installed and configured by default. You want to determine which applications need performance tuning. What should you do? (Choose two.)

A. Examine the wall-clock time and the CPU time of the application. If the difference is substantial increase the CPU resource allocation.

B. Examine the wall-clock time and the CPU time of the application. If the difference is substantial, increase the memory resource allocation.

C. Examine the wall-clock time and the CPU time of the application. If the difference is substantial, increase the local disk storage allocation.

D. Examine the latency time the wall-clock time and the CPU time of the application. If the latency time is slowly burning down the error budget, and the difference between wall-clock time and CPU time is minimal mark the application for optimization.

E. Examine the heap usage of the application. If the usage is low, mark the application for optimization.

Correct Answer: AD

<https://cloud.google.com/profiler/docs/concepts-profiling>

---

## QUESTION 10

You work for a global organization and are running a monolithic application on Compute Engine. You need to select the machine type for the application to use that optimizes CPU utilization by using the fewest number of steps. You want to use historical system metrics to identify the machine type for the application to use. You want to follow Google-recommended practices. What should you do?

- A. Use the Recommender API and apply the suggested recommendations.
- B. Create an Agent Policy to automatically install Ops Agent in all VMs.
- C. Install the Ops Agent in a fleet of VMs by using the gcloud CLI.
- D. Review the Cloud Monitoring dashboard for the VM and choose the machine type with the lowest CPU utilization.

Correct Answer: A

---

### QUESTION 11

You are performing a semi-annual capacity planning exercise for your flagship service. You expect a service user growth rate of 10% month-over-month over the next six months. Your service is fully containerized and runs on Google Cloud Platform (GCP), using a Google Kubernetes Engine (GKE) Standard regional cluster on three zones with cluster autoscaler enabled. You currently consume about 30% of your total deployed CPU capacity, and you require resilience against the failure of a zone. You want to ensure that your users experience minimal negative impact as a result of this growth or as a result of zone failure, while avoiding unnecessary costs. How should you prepare to handle the predicted growth?

- A. Verify the maximum node pool size, enable a horizontal pod autoscaler, and then perform a load test to verify your expected resource needs.
- B. Because you are deployed on GKE and are using a cluster autoscaler, your GKE cluster will scale automatically regardless of growth rate.
- C. Because you are at only 30% utilization, you have significant headroom and you won't need to add any additional capacity for this rate of growth.
- D. Proactively add 60% more node capacity to account for six months of 10% growth rate, and then perform a load test to make sure you have enough.

Correct Answer: A

<https://cloud.google.com/kubernetes-engine/docs/concepts/horizontalpodautoscaler>

---

### QUESTION 12

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 tool. Ensure that your team members add their SMS/phone numbers to the external tool.
- C. Ensure that your team members set their SMS/phone numbers in their Stackdriver Profile. Select the SMS notification option for each alerting policy and then select the appropriate SMS/phone numbers from the list.
- D. Configure a Slack notification for each alerting policy. Set up a Slack-to-SMS integration to send SMS messages when Slack messages are received. Ensure that your team members add their SMS/phone numbers to the external

integration.

Correct Answer: C

[https://cloud.google.com/monitoring/support/notification-options#creating\\_channels](https://cloud.google.com/monitoring/support/notification-options#creating_channels)