



VCE & PDF

PassApply.com

<https://www.passapply.com/associate-cloud-engineer.html>

2024 Latest passapply ASSOCIATE-CLOUD-ENGINEER PDF and VCE dumps
Download

ASSOCIATE-CLOUD-ENGINEER^{Q&As}

Associate Cloud Engineer

**Pass Google ASSOCIATE-CLOUD-ENGINEER Exam
with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/associate-cloud-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

Your company completed the acquisition of a startup and is now merging the IT systems of both companies. The startup had a production Google Cloud project in their organization. You need to move this project into your organization and ensure that the project is billed to your organization. You want to accomplish this task with minimal effort. What should you do?

- A. Use the projects.move method to move the project to your organization. Update the billing account of the project to that of your organization.
- B. Ensure that you have an Organization Administrator Identity and Access Management (IAM) role assigned to you in both organizations. Navigate to the Resource Manager in the startup's Google Cloud organization, and drag the project to your company's organization.
- C. Create a Private Catalog for the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalog. Share the Catalog with your organization, and deploy the resources in your company's project.
- D. Create an infrastructure-as-code template for all resources in the project by using Terraform, and deploy that template to a new project in your organization. Delete the project from the startup's Google Cloud organization.

Correct Answer: A

<https://cloud.google.com/resource-manager/docs/project-migration-checklist>

QUESTION 2

Your projects incurred more costs than you expected last month. Your research reveals that a development GKE container emitted a huge number of logs, which resulted in higher costs. You want to disable the logs quickly using the minimum number of steps. What should you do?

- A. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE container resource.
- B. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE Cluster Operations resource.
- C. 1. Go to the GKE console, and delete existing clusters.
2.
Recreate a new cluster.
3.
Clear the option to enable legacy Stackdriver Logging.
- D. 1. Go to the GKE console, and delete existing clusters.
2.
Recreate a new cluster.
3.



Clear the option to enable legacy Stackdriver Monitoring.

Correct Answer: A

<https://cloud.google.com/logging/docs/api/v2/resource-list>

GKE Containers have more log than GKE Cluster Operations:

.-GKE Containe:

cluster_name: An immutable name for the cluster the container is running in.

namespace_id: Immutable ID of the cluster namespace the container is running in.

instance_id: Immutable ID of the GCE instance the container is running in.

pod_id: Immutable ID of the pod the container is running in.

container_name: Immutable name of the container.

zone: The GCE zone in which the instance is running.

VS

.-GKE Cluster Operations

project_id: The identifier of the GCP project associated with this resource, such as "my-project".

cluster_name: The name of the GKE Cluster.

location: The location in which the GKE Cluster is running.

QUESTION 3

You want to run a single caching HTTP reverse proxy on GCP for a latency-sensitive website. This specific reverse proxy consumes almost no CPU. You want to have a 30-GB in-memory cache, and need an additional 2 GB of memory for the rest of the processes. You want to minimize cost. How should you run this reverse proxy?

- A. Create a Cloud Memorystore for Redis instance with 32-GB capacity.
- B. Run it on Compute Engine, and choose a custom instance type with 6 vCPUs and 32 GB of memory.
- C. Package it in a container image, and run it on Kubernetes Engine, using n1-standard-32 instances as nodes.
- D. Run it on Compute Engine, choose the instance type n1-standard-1, and add an SSD persistent disk of 32 GB.

Correct Answer: A

Go to cloud console and create instance select Memorystore with Basic tier, select us-central1 and us-central1-a, and capacity 32GB, the cost estimate is \$0.023/GB/hr

select VM instance with custom machine type with 6 vCPUs and 32 GB memory, the same region and zone as Memorystore setting, the cost estimate is \$0.239/hr

Option B will definitely cost more as it adds on CPU usage cost even it uses little in this scenario, but still charge you.



So answer is A from real practice example.

QUESTION 4

You have a web application deployed as a managed instance group. You have a new version of the application to gradually deploy. Your web application is currently receiving live web traffic. You want to ensure that the available capacity does not decrease during the deployment. What should you do?

- A. Perform a rolling-action start-update with maxSurge set to 0 and maxUnavailable set to 1.
- B. Perform a rolling-action start-update with maxSurge set to 1 and maxUnavailable set to 0.
- C. Create a new managed instance group with an updated instance template. Add the group to the backend service for the load balancer. When all instances in the new managed instance group are healthy, delete the old managed instance group.
- D. Create a new instance template with the new application version. Update the existing managed instance group with the new instance template. Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

Correct Answer: B

We need to ensure the global capacity remains intact, for that reason we need to establish maxUnavailable to 0. On the other hand, we need to ensure new instances can be created. We do that by establishing the maxSurge to 1. Option C is more expensive and more difficult to set up and option D won't meet requirements since it won't keep global capacity intact.

QUESTION 5

You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

- A. 1. Create a consumer Gmail account.

2.

Write a script that monitors the CPU usage.

3.

When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.

- B. 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.

2.

Create an Alerting Policy in Stackdriver that uses the threshold as a trigger condition.

3.

Configure your email address in the notification channel.



C. 1. Create a Stackdriver Workspace, and associate your GCP project with it.

2.

Write a script that monitors the CPU usage and sends it as a custom metric to Stackdriver.

3.

Create an uptime check for the instance in Stackdriver.

D. 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: ([0-9] {1,3})%

2.

In Stackdriver Monitoring, create an Alerting Policy based on this metric.

3.

Configure your email address in the notification channel.

Correct Answer: B

[ASSOCIATE-CLOUD-ENGINEER VCE Dumps](#)

[ASSOCIATE-CLOUD-ENGINEER Practice Test](#)

[ASSOCIATE-CLOUD-ENGINEER Braindumps](#)