

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.leads4pass.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

You need to set a budget alert for use of Compute Engine services on one of the three Google Cloud Platform projects that you manage. All three projects are linked to a single billing account. What should you do?

- A. Verify that you are the project billing administrator. Select the associated billing account and create a budget and alert for the appropriate project.
- B. Verify that you are the project billing administrator. Select the associated billing account and create a budget and a custom alert.
- C. Verify that you are the project administrator. Select the associated billing account and create a budget for the appropriate project.
- D. Verify that you are project administrator. Select the associated billing account and create a budget and a custom alert.

Correct Answer: A

Billing Administrator is in between organization and projects, So he can chose and select appropriate project.

QUESTION 2

You recently deployed a new version of an application to App Engine and then discovered a bug in the release. You need to immediately revert to the prior version of the application. What should you do?

- A. Run `gcloud app restore`.
- B. On the App Engine page of the GCP Console, select the application that needs to be reverted and click Revert.
- C. On the App Engine Versions page of the GCP Console, route 100% of the traffic to the previous version.
- D. Deploy the original version as a separate application. Then go to App Engine settings and split traffic between applications so that the original version serves 100% of the requests.

Correct Answer: C

Reference: <https://medium.com/google-cloud/app-engine-project-cleanup-9647296e796a>

Option A is wrong as `gcloud app restore` was used for backup and restore and has been deprecated. Option B is wrong as there is no application revert functionality available. Option D is wrong as App Engine maintains version and need not be redeployed.

QUESTION 3

You are responsible for a web application on Compute Engine. You want your support team to be notified automatically if users experience high latency for at least 5 minutes. You need a Google-recommended solution with no development cost. What should you do?

- A. Export Cloud Monitoring metrics to BigQuery and use a Looker Studio dashboard to monitor your web application's latency.

- B. Create an alert policy to send a notification when the HTTP response latency exceeds the specified threshold.
- C. Implement an App Engine service which invokes the Cloud Monitoring API and sends a notification in case of anomalies.
- D. Use the Cloud Monitoring dashboard to observe latency and take the necessary actions when the response latency exceeds the specified threshold.

Correct Answer: B

<https://cloud.google.com/monitoring/alerts#alerting-example>

QUESTION 4

You have a large 5-TB AVRO file stored in a Cloud Storage bucket. Your analysts are proficient only in SQL and need access to the data stored in this file. You want to find a cost-effective way to complete their request as soon as possible. What should you do?

- A. Load data in Cloud Datastore and run a SQL query against it.
- B. Create a BigQuery table and load data in BigQuery. Run a SQL query on this table and drop this table after you complete your request.
- C. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.
- D. Create a Hadoop cluster and copy the AVRO file to NFS by compressing it. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Correct Answer: C

Breaking down the question into key points

1.

5-TB AVRO file stored in a Cloud Storage bucket.

2.

Analysts are proficient only in SQL

3.

cost-effective way to complete their request as soon as possible

- A.Load data in Cloud Datastore... (Not Correct because Cloud Datastore is not a good option to run SQL Queries)
- B. ...Load data in BigQuery.... (Not Cost Effective because loading the data which is already present in the bucket into BigQuery again is expensive)
- C. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request. (This is the right answer as it meets all the requirements from the question)
- D. Create a Hadoop cluster and copy the AVRO file to NFS by compressing it. Load the file in a hive table and provide access to your analysts so that they can run SQL queries. (Too roundabout and indirect. Not the right option)

QUESTION 5

Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1-standard-2 nodes. You need to deploy additional pods requiring n2-highmem-16 nodes without any downtime. What should you do?

- A. Use gcloud container clusters upgrade. Deploy the new services.
- B. Create a new Node Pool and specify machine type n2-highmem-16. Deploy the new pods.
- C. Create a new cluster with n2-highmem-16 nodes. Redeploy the pods and delete the old cluster.
- D. Create a new cluster with both n1-standard-2 and n2-highmem-16 nodes. Redeploy the pods and delete the old cluster.

Correct Answer: B

B is correct answer, read below form google docs;

This tutorial demonstrates how to migrate workloads running on a Google Kubernetes Engine (GKE) cluster to a new set of nodes within the same cluster without incurring downtime for your application. Such a migration can be useful if you

want to migrate your workloads to nodes with a different machine type.

Background

A node pool is a subset of machines that all have the same configuration, including machine type (CPU and memory) authorization scopes. Node pools represent a subset of nodes within a cluster; a container cluster can contain one or more

node pools.

When you need to change the machine profile of your Compute Engine cluster, you can create a new node pool and then migrate your workloads over to the new node pool.

To migrate your workloads without incurring downtime, you need to:

Mark the existing node pool as unschedulable.

Drain the workloads running on the existing node pool.

Delete the existing node pool.

https://cloud.google.com/kubernetes-engine/docs/tutorials/migrating-node-pool#creating_a_node_pool_with_large_machine_type

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

[ASSOCIATE-CLOUD-ENGINEER Exam Questions](#)

[ASSOCIATE-CLOUD-ENGINEER Braindumps](#)