



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

You are the organization and billing administrator for your company. The engineering team has the Project Creator role on the organization. You do not want the engineering team to be able to link projects to the billing account. Only the finance team should be able to link a project to a billing account, but they should not be able to make any other changes to projects. What should you do?

- A. Assign the finance team only the Billing Account User role on the billing account.
- B. Assign the engineering team only the Billing Account User role on the billing account.
- C. Assign the finance team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.
- D. Assign the engineering team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.

Correct Answer: C

When granted in combination with the `_Billing Account User` role, the `_Project Billing Manager` role allows a user to attach the project to the billing account, but does not grant any rights over resources - YOU NEED BOTH

QUESTION 2

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

- A. 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.
- 2. Update your instances' metadata to add the following value: logs-destination: bq://platform-logs.

- B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.
- 2.

Create a Cloud Function that is triggered by messages in the logs topic.

- 3.

Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.

- C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.
- 2.

Click Create Export.

- 3.

Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.



D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.

2.

Configure this Cloud Function to create a BigQuery Job that executes this query:

```
INSERT INTO dataset.platform-logs (timestamp, log)
```

```
SELECT timestamp, log FROM compute.logs
```

```
WHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)
```

3.

Use Cloud Scheduler to trigger this Cloud Function once a day.

Correct Answer: C

C. is correct, Sinks control how Cloud Logging routes logs. Using sinks, you can route some or all of your logs to supported destinations.

QUESTION 3

You are using multiple configurations for gcloud. You want to review the configured Kubernetes Engine cluster of an inactive configuration using the fewest possible steps. What should you do?

- A. Use gcloud config configurations describe to review the output.
- B. Use gcloud config configurations activate and gcloud config list to review the output.
- C. Use kubectl config get-contexts to review the output.
- D. Use kubectl config use-context and kubectl config view to review the output.

Correct Answer: D

Reference: <https://medium.com/google-cloud/kubernetes-engine-kubectl-config-b6270d2b656c>

```
kubectl config view -o jsonpath='{.users[0].name}' # display the first user kubectl config view -o  
jsonpath='{.users[*].name}' # get a list of users kubectl config get-contexts # display list of contexts kubectl config  
current-context # display the current-context kubectl config use-context my-cluster-name # set the default context to my-  
cluster-name
```

QUESTION 4

You have been asked to create robust Virtual Private Network (VPN) connectivity between a new Virtual Private Cloud (VPC) and a remote site. Key requirements include dynamic routing, a shared address space of 10.19.0.1/22, and no overprovisioning of tunnels during a failover event. You want to follow Google-recommended practices to set up a high availability Cloud VPN. What should you do?

- A. Use a custom mode VPC network, configure static routes, and use active/passive routing
- B. Use an automatic mode VPC network, configure static routes, and use active/active routing



- C. Use a custom mode VPC network use Cloud Router border gateway protocol (86P) routes, and use active/passive routing
- D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Correct Answer: C

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/best-practices>

QUESTION 5

You need to create a Compute Engine instance in a new project that doesn't exist yet. What should you do?

- A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.
- B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the --project flag to specify a new project.
- C. Using the Cloud SDK, create the new instance, and use the --project flag to specify the new project. Answer yes when prompted by Cloud SDK to enable the Compute Engine API.
- D. Enable the Compute Engine API in the Cloud Console. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Correct Answer: A

Quickstart: Creating a New Instance Using the Command Line Before you begin

1.

In the Cloud Console, on the project selector page, select or create a Cloud project.

2.

Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project.

To use the gcloud command-line tool for this quickstart, you must first install and initialize the Cloud SDK:

1.

Download and install the Cloud SDK using the instructions given on Installing Google Cloud SDK.

2.

Initialize the SDK using the instructions given on Initializing Cloud SDK.

To use gcloud in Cloud Shell for this quickstart, first activate Cloud Shell using the instructions given on Starting Cloud Shell.

<https://cloud.google.com/ai-platform/deep-learning-vm/docs/quickstart-cli#before-you-begin>



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 VCE Dumps](#)

[ASSOCIATE-CLOUD-ENGINEER Study Guide](#)

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