



VCE & PDF

GeekCert.com

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

2024 Latest geekcert 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.geekcert.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 have an application that runs on Compute Engine VM instances in a custom Virtual Private Cloud (VPC). Your company's security policies only allow the use to internal IP addresses on VM instances and do not let VM instances connect to the internet. You need to ensure that the application can access a file hosted in a Cloud Storage bucket within your project. What should you do?

- A. Enable Private Service Access on the Cloud Storage Bucket.
- B. Add storage.googleapis.com to the list of restricted services in a VPC Service Controls perimeter and add your project to the list to protected projects.
- C. Enable Private Google Access on the subnet within the custom VPC.
- D. Deploy a Cloud NAT instance and route the traffic to the dedicated IP address of the Cloud Storage bucket.

Correct Answer: C

Private Google Access lets you connect VM instances to GCP services without external IP addresses and only internal. A is wrong because even though Private Services Access lets you also access GCP and other services through internal IPs, it also allows the VMs to have external IPs. <https://cloud.google.com/vpc/docs/private-google-access>

QUESTION 2

You need to monitor resources that are distributed over different projects in Google Cloud Platform. You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?

- A. Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
- B. For each project, create a Stackdriver account. In each project, create a service account for that project and grant it the role of Stackdriver Account Editor in all other projects.
- C. Configure a single Stackdriver account, and link all projects to the same account.
- D. Configure a single Stackdriver account for one of the projects. In Stackdriver, create a Group and add the other project names as criteria for that Group.

Correct Answer: C

First of all D is incorrect, Groups are used to define alerts on set of resources(such as VM instances, databases, and load balancers). FYI tried adding Two projects into a group it did not allowed me as the "AND"/"OR" criteria for the group

failed with this combination of resources.

C is correct because,

When you initially click on Monitoring(Stackdriver Monitoring) it creates a workspace(a stackdriver account) linked to the ACTIVE(CURRENT) Project from which it was clicked.

Now if you change the project and again click onto Monitoring it would create an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don't want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).



If you have accidentally created two diff workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.

If we have only one workspace and two projects we can simply add other GCP Project under

Monitoring > Settings > GCP Projects > Add GCP Projects.

In both of these cases we did not create a GROUP, we just linked GCP Project to the workspace(stackdriver account).

QUESTION 3

You've created a Pod using the kubectl run command. Now you're attempting to remove the Pod, and it keeps being recreated. Which command might help you as you attempt to remove the pod?

- A. gcloud container describe pods
- B. kubectl get pods
- C. kubectl get secrets
- D. kubectl get deployments

Correct Answer: D

QUESTION 4

You've been asked to help onboard a new member of the big-data team. They need full access to BigQuery. Which type of role would be the most efficient to set up while following the principle of least privilege?

- A. Primitive Role
- B. Custom Role
- C. Managed Role
- D. Predefined Role

Correct Answer: D

QUESTION 5

Your company uses BigQuery for data warehousing. Over time, many different business units in your company have created 1000+ datasets across hundreds of projects. Your CIO wants you to examine all datasets to find tables that contain an employee_ssn column. You want to minimize effort in performing this task. What should you do?

- A. Go to Data Catalog and search for employee_ssn in the search box.
- B. Write a shell script that uses the bq command line tool to loop through all the projects in your organization.
- C. Write a script that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find the employee_ssn column.



D. Write a Cloud Dataflow job that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find employee_ssn column.

Correct Answer: A

A is the correct answer, Data Catalog can be used to search the column with keyword:value pair,

Filter your search by adding a keyword:value to your search terms in the search box:

Keyword Description name: Match data asset name ***column: Match column name or nested column name
description: Match table description

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

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

[ASSOCIATE-CLOUD-ENGINEER Braindumps](#)