



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 create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectl to examine the status of your Pod and observe that one of them is still in Pending status:

```
$ kubectl get pods -l app=myapp
```

NAME	READY	STATUS	RESTART	AGE
myapp-deployment-58ddb995-lp86m	0/1	Pending	0	9m
myapp-deployment-58ddb995-qjpkg	1/1	Running	0	9m

What is the most likely cause?

- A. The pending Pod's resource requests are too large to fit on a single node of the cluster.
- B. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.
- C. The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.
- D. The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' status. It is currently being rescheduled on a new node.

Correct Answer: B

Reasons for a Pod Status Pending: Troubleshooting Reason #1: Not enough CPU Troubleshooting Reason #2: Not enough memory Troubleshooting Reason #3: Not enough CPU and memory

<https://managedk8s.com/kubernetes/k8sbot/troubleshooting/pending/pod/2019/02/22/pending-pod.html>

QUESTION 2

You are creating an application that will run on Google Kubernetes Engine. You have identified MongoDB as the most suitable database system for your application and want to deploy a managed MongoDB environment that provides a support SLA. What should you do?

- A. Create a Cloud Bigtable cluster, and use the HBase API.
- B. Deploy MongoDB Atlas from the Google Cloud Marketplace.
- C. Download a MongoDB installation package, and run it on Compute Engine instances.
- D. Download a MongoDB installation package, and run it on a Managed Instance Group.

Correct Answer: C

QUESTION 3



You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

A. 1. Use nslookup to get the IP address for storage.googleapis.com.

2.

Negotiate with the security team to be able to give a public IP address to the servers.

3.

Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.

B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).

2.

In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.

3.

Configure your servers to use that instance as a proxy to access Cloud Storage.

C. 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.

2.

Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.

3.

Configure your new instances to use this ILB as proxy.

D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.

2.

Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.

3.

In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

Correct Answer: D

VPN/interconnect without exposing to public

QUESTION 4

You have designed a solution on Google Cloud Platform (GCP) that uses multiple GCP products. Your company has asked you to estimate the costs of the solution. You need to provide estimates for the monthly total cost. What should



you do?

- A. For each GCP product in the solution, review the pricing details on the products pricing page. Use the pricing calculator to total the monthly costs for each GCP product.
- B. For each GCP product in the solution, review the pricing details on the products pricing page. Create a Google Sheet that summarizes the expected monthly costs for each product.
- C. Provision the solution on GCP. Leave the solution provisioned for 1 week. Navigate to the Billing Report page in the Google Cloud Platform Console. Multiply the 1 week cost to determine the monthly costs.
- D. Provision the solution on GCP. Leave the solution provisioned for 1 week. Use Stackdriver to determine the provisioned and used resource amounts. Multiply the 1 week cost to determine the monthly costs.

Correct Answer: A

For each Google Cloud product in the solution, review the pricing details on the products pricing page. Use the pricing calculator to total the monthly costs for each Google Cloud product.

QUESTION 5

You are creating a Google Kubernetes Engine (GKE) cluster with a cluster autoscaler feature enabled. You need to make sure that each node of the cluster will run a monitoring pod that sends container metrics to a third-party monitoring solution. What should you do?

- A. Deploy the monitoring pod in a StatefulSet object.
- B. Deploy the monitoring pod in a DaemonSet object.
- C. Reference the monitoring pod in a Deployment object.
- D. Reference the monitoring pod in a cluster initializer at the GKE cluster creation time.

Correct Answer: B

B is right: <https://kubernetes.io/docs/concepts/workloads/controllers/daemonset/> Some typical uses of a DaemonSet are:

running a cluster storage daemon on every node
running a logs collection daemon on every node
running a node monitoring daemon on every node

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

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

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