



Google

*Associate-Cloud-Engineer
Google Associate Cloud Engineer Exam*

Questions & Answers PDF

For More Information:

<https://www.certswarrior.com/>

Features:

- 90 Days Free Updates
- 30 Days Money Back Guarantee
- Instant Download Once Purchased
- 24/7 Online Chat Support
- Its Latest Version

Latest Version: 25.3

Question: 1

Every employee of your company has a Google account. Your operational team needs to manage a large number of instances on Compute Engine. Each member of this team needs only administrative access to the servers. Your security team wants to ensure that the deployment of credentials is operationally efficient and must be able to determine who accessed a given instance. What should you do?

- A. Generate a new SSH key pair. Give the private key to each member of your team. Configure the public key in the metadata of each instance.
- B. Ask each member of the team to generate a new SSH key pair and to send you their public key. Use a configuration management tool to deploy those keys on each instance.
- C. Ask each member of the team to generate a new SSH key pair and to add the public key to their Google account. Grant the “compute.osAdminLogin” role to the Google group corresponding to this team.
- D. Generate a new SSH key pair. Give the private key to each member of your team. Configure the public key as a project-wide public SSH key in your Cloud Platform project and allow project-wide public SSH keys on each instance.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instances/managing-instance-access>

Question: 2

You need to create a custom VPC with a single subnet. The subnet’s range must be as large as possible. Which range should you use?

- A. .00.0.0/0
- B. 10.0.0.0/8
- C. 172.16.0.0/12
- D. 192.168.0.0/16

Answer: B

Explanation:

https://cloud.google.com/vpc/docs/vpc#manually_created_subnet_ip_ranges

Question: 3

You want to select and configure a cost-effective solution for relational data on Google Cloud Platform. You are working with a small set of operational data in one geographic location. You need to support point-in-time recovery. What should you do?

- A. Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
- B. Select Cloud SQL (MySQL). Select the create failover replicas option.
- C. Select Cloud Spanner. Set up your instance with 2 nodes.
- D. Select Cloud Spanner. Set up your instance as multi-regional.

Answer: A

Explanation:

Reference: <https://cloud.google.com/sql/docs/mysql/backup-recovery/restore>
<https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#disk-usage>

Question: 4

You want to configure autohealing for network load balancing for a group of Compute Engine instances that run in multiple zones, using the fewest possible steps. You need to configure re-creation of VMs if they are unresponsive after 3 attempts of 10 seconds each. What should you do?

- A. Create an HTTP load balancer with a backend configuration that references an existing instance group.
Set the health check to healthy (HTTP).
- B. Create an HTTP load balancer with a backend configuration that references an existing instance group.
Define a balancing mode and set the maximum RPS to 10.
- C. Create a managed instance group. Set the Autohealing health check to healthy (HTTP).
- D. Create a managed instance group. Verify that the autoscaling setting is on.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instance-groups>

<https://cloud.google.com/load-balancing/docs/network/transition-to-backend-services#console>

In order to enable auto-healing, you need to group the instances into a managed instance group. Managed instance groups (MIGs) maintain the high availability of your applications by proactively keeping your virtual machine (VM) instances available. An auto-healing policy on the MIG relies on an application-based health check to verify that an application is responding as expected. If the auto-healer determines that an application isn't responding, the managed instance group automatically recreates that instance.

It is important to use separate health checks for load balancing and for auto-healing. Health checks for load balancing can and should be more aggressive because these health checks determine whether an instance receives user traffic. You want to catch non-responsive instances quickly, so you can redirect

traffic if necessary. In contrast, health checking for auto-healing causes Compute Engine to proactively replace failing instances, so this health check should be more conservative than a load balancing health check.

Question: 5

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.

Answer: D

Explanation:

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

<https://kubernetes.io/docs/reference/kubectl/cheatsheet/>



CERTSWARRIOR

FULL PRODUCT INCLUDES:

Money Back Guarantee



Instant Download after Purchase



90 Days Free Updates



PDF Format Digital Download



24/7 Live Chat Support



Latest Syllabus Updates



For More Information – Visit link below:

<https://www.certswarrior.com>

16 USD Discount Coupon Code: U89DY2AQ