

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.pass2lead.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 create a custom IAM role for use with a GCP service. All permissions in the role must be suitable for production use. You also want to clearly share with your organization the status of the custom role. This will be the first version of the custom role. What should you do?

- A. Use permissions in your role that use the `supported` support level for role permissions. Set the role stage to ALPHA while testing the role permissions.
- B. Use permissions in your role that use the `supported` support level for role permissions. Set the role stage to BETA while testing the role permissions.
- C. Use permissions in your role that use the `testing` support level for role permissions. Set the role stage to ALPHA while testing the role permissions.
- D. Use permissions in your role that use the `testing` support level for role permissions. Set the role stage to BETA while testing the role permissions.

Correct Answer: A

Testing and deploying

Custom roles include a launch stage, which is stored in the stage property for the role. The launch stage is informational; it helps you keep track of whether each role is ready for widespread use.

Each custom role can have one of the following launch stages:

Launch stages

ALPHA The role is still being developed or tested, or it includes permissions for Google Cloud services or features that are not yet public. It is not ready for widespread use. BETA The role has been tested on a limited basis, or it includes permissions for Google Cloud services or features that are not generally available. GA The role has been widely tested, and all of its permissions are for Google Cloud services or features that are generally available.

QUESTION 2

You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

- A. Open the Cloud Spanner console to review configurations.
- B. Open the IAM and admin console to review IAM policies for Cloud Spanner roles.
- C. Go to the Stackdriver Monitoring console and review information for Cloud Spanner.
- D. Go to the Stackdriver Logging console, review admin activity logs, and filter them for Cloud Spanner IAM roles.

Correct Answer: D

Activity logs captures the time when the users were given the IAM roles for Cloud Spanner

QUESTION 3

You want to verify the IAM users and roles assigned within a GCP project named my-project. What should you do?

- A. Run `gcloud iam roles list`. Review the output section.
- B. Run `gcloud iam service-accounts list`. Review the output section.
- C. Navigate to the project and then to the IAM section in the GCP Console. Review the members and roles.
- D. Navigate to the project and then to the Roles section in the GCP Console. Review the roles and status.

Correct Answer: C

C is the only logical answers. If you go IAM and Admin > IAM: You can see Principals and Roles. Users, groups, service accounts are Principals

QUESTION 4

Your engineers have asked you to set up a subnet with the largest IP address range possible. Which of the following ranges would work best?

- A. 0.0.0.0/0
- B. 10.0.0.0/32
- C. 192.168.0.0/16
- D. 10.0.0.0/8

Correct Answer: B

QUESTION 5

An external member of your team needs list access to compute images and disks in one of your projects. You want to follow Google-recommended practices when you grant the required permissions to this user. What should you do?

- A. Create a custom role, and add all the required `compute.disks.list` and `compute.images.list` permissions as `includedPermissions`. Grant the custom role to the user at the project level.
- B. Create a custom role based on the Compute Image User role. Add the `compute.disks.list` to the `includedPermissions` field. Grant the custom role to the user at the project level.
- C. Create a custom role based on the Compute Storage Admin role. Exclude unnecessary permissions from the custom role. Grant the custom role to the user at the project level.
- D. Grant the Compute Storage Admin role at the project level.

Correct Answer: A

<https://cloud.google.com/iam/docs/custom-roles-permissions-support>

Both `compute.disks.list` and `compute.images.list` are available as permissions for custom roles. Makes more sense to

make a new custom role than going off an admin one then adjusting it.

QUESTION 6

You are operating a Google Kubernetes Engine (GKE) cluster for your company where different teams can run non-production workloads. Your Machine Learning (ML) team needs access to Nvidia Tesla P100 GPUs to train their models. You want to minimize effort and cost. What should you do?

- A. Ask your ML team to add the "accelerator: gpu" annotation to their pod specification.
- B. Recreate all the nodes of the GKE cluster to enable GPUs on all of them.
- C. Create your own Kubernetes cluster on top of Compute Engine with nodes that have GPUs. Dedicate this cluster to your ML team.
- D. Add a new, GPU-enabled, node pool to the GKE cluster. Ask your ML team to add the cloud.google.com/gke-accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Correct Answer: D

Add a new, GPU-enabled, node pool to the GKE cluster. Ask your ML team to add the cloud.google.com/gke-accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

QUESTION 7

You want to permanently delete a Pub/Sub topic managed by Config Connector in your Google Cloud project. What should you do?

- A. Use kubectl to create the label deleted-by-cnrm and to change its value to true for the topic resource.
- B. Use kubectl to delete the topic resource.
- C. Use gcloud CLI to delete the topic.
- D. Use gcloud CLI to update the topic label managed-by-cnrm to false.

Correct Answer: B

https://cloud.google.com/config-connector/docs/how-to/getting-started#deleting_a_resource

QUESTION 8

You have one GCP account running in your default region and zone and another account running in a non-default region and zone. You want to start a new Compute Engine instance in these two Google Cloud Platform accounts using the command line interface. What should you do?

- A. Create two configurations using gcloud config configurations create [NAME]. Run gcloud config configurations activate [NAME] to switch between accounts when running the commands to start the Compute Engine instances.
- B. Create two configurations using gcloud config configurations create [NAME]. Run gcloud configurations list to start the Compute Engine instances.

C. Activate two configurations using gcloud configurations activate [NAME]. Run gcloud config list to start the Compute Engine instances.

D. Activate two configurations using gcloud configurations activate [NAME]. Run gcloud configurations list to start the Compute Engine instances.

Correct Answer: A

Reference: <https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate>

All the other options don't make any sense when day say "Run gcloud configurations list to start the Compute Engine instances". How the heck are you expecting to "start" GCE instances doing "configuration list". Obviously B,C,D don't make any sense.

QUESTION 9

What's the easiest way to ensure that the nodes in your Kubernetes cluster are always up-to-date with the latest stable version of Kubernetes?

A. Opt into the Kubernetes Node Update program from the quotas page.

B. Run the kubectl nodes update command.

C. Run the kubectl nodes upgrade command.

D. Enable the automatic node upgrades setting.

Correct Answer: D

QUESTION 10

You keep attempting to execute a command, and while there's no error being thrown, you suspect something is going wrong. You want to check out the gcloud logs. However, you can't recall in which directory they're located. Your script skills

are not what they could be, so you can't rely on using those fancy command line skills.

Which command could you run to show you where the log directory is located?

A. google logs

B. gcloud info

C. gcloud logging

D. gcloud

Correct Answer: B

QUESTION 11

The core business of your company is to rent out construction equipment at a large scale. All the equipment that is being rented out has been equipped with multiple sensors that send event information every few seconds. These signals can vary from engine status, distance traveled, fuel level, and more. Customers are billed based on the consumption monitored by these sensors. You expect high throughput - up to thousands of events per hour per device - and need to retrieve consistent data based on the time of the event. Storing and retrieving individual signals should be atomic. What should you do?

- A. Create a file in Cloud Storage per device and append new data to that file.
- B. Create a file in Cloud Filestore per device and append new data to that file.
- C. Ingest the data into Datastore. Store data in an entity group based on the device.
- D. Ingest the data into Cloud Bigtable. Create a row key based on the event timestamp.

Correct Answer: D

Keyword need to look for

- "High Throughput",

- "Consistent",

-

"Property based data insert/fetch like engine status, distance traveled, fuel level, and more." which can be designed in column,

-

"Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds" This will go for per byte situation,

- Export data based on the time of the event.

- Atomic

BigTable will fit all requirements.

DataStore is not fully Atomic

CloudStorage is not an option where we can export data based on time of event. We need another solution to do that

Firestore can be used with MobileSDK.

So go with Option D: Big Table

QUESTION 12

You have an application running in Google Kubernetes Engine (GKE) with cluster autoscaling enabled. The application exposes a TCP endpoint. There are several replicas of this application. You have a Compute Engine instance in the same region, but in another Virtual Private Cloud (VPC), called gce-network, that has no overlapping IP ranges with the first VPC. This instance needs to connect to the application on GKE. You want to minimize effort. What should you do?

- A. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.

2.
Set the service's externalTrafficPolicy to Cluster.
3.
Configure the Compute Engine instance to use the address of the load balancer that has been created.
- B. 1. In GKE, create a Service of type NodePort that uses the application's Pods as backend.
 2.
Create a Compute Engine instance called proxy with 2 network interfaces, one in each VPC.
 3.
Use iptables on this instance to forward traffic from gce-network to the GKE nodes.
 4.
Configure the Compute Engine instance to use the address of proxy in gce-network as endpoint.
- C. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.
 2.
Add an annotation to this service: cloud.google.com/load-balancer-type: Internal
 3.
Peer the two VPCs together.
 4.
Configure the Compute Engine instance to use the address of the load balancer that has been created.
- D. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.
 2.
Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG's instances.
 3.
Configure the Compute Engine instance to use the address of the load balancer that has been created.

Correct Answer: A

It's never mentioned in the question that traffic cannot go through the Internet but it's mentioned that effort should be minimized. A requires a lot less effort than C to accomplish the same (no VPC peering, per example).

QUESTION 13

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

- A. Set up an export job for the first of the month. Write the export file to an Archive class Cloud Storage bucket.
- B. Save the automatic first-of-the-month backup for three years. Store the backup file in an Archive class Cloud Storage bucket.
- C. Set up an on-demand backup for the first of the month. Write the backup to an Archive class Cloud Storage bucket.
- D. Convert the automatic first-of-the-month backup to an export file. Write the export file to a Coldline class Cloud Storage bucket.

Correct Answer: A

https://cloud.google.com/sql/docs/mysql/backup-recovery/backups#can_i_export_a_backup
https://cloud.google.com/sql/docs/mysql/import-export#automating_export_operations

QUESTION 14

You are building a product on top of Google Kubernetes Engine (GKE). You have a single GKE cluster. For each of your customers, a Pod is running in that cluster, and your customers can run arbitrary code inside their Pod. You want to maximize the isolation between your customers' Pods. What should you do?

- A. Use Binary Authorization and whitelist only the container images used by your customers' Pods.
- B. Use the Container Analysis API to detect vulnerabilities in the containers used by your customers' Pods.
- C. Create a GKE node pool with a sandbox type configured to gvisor. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.
- D. Use the cos_containerd image for your GKE nodes. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Correct Answer: C

Reference: <https://cloud.google.com/kubernetes-engine/sandbox/>

QUESTION 15

You have deployed multiple Linux instances on Compute Engine. You plan on adding more instances in the coming weeks. You want to be able to access all of these instances through your SSH client over the internet without having to configure specific access on the existing and new instances. You do not want the Compute Engine instances to have a public IP. What should you do?

- A. Configure Cloud Identity-Aware Proxy for HTTPS resources.
- B. Configure Cloud Identity-Aware Proxy for SSH and TCP resources
- C. Create an SSH keypair and store the public key as a project-wide SSH Key.
- D. Create an SSH keypair and store the private key as a project-wide SSH Key.

Correct Answer: B

<https://cloud.google.com/iap/docs/using-tcp-forwarding>

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

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

[ASSOCIATE-CLOUD-ENGINEER Braindumps](#)