

Google Associate Cloud Engineer Actual Exam

1. We highly recommend you should take Google Associate Cloud Engineer Actual Exam Version because it include actual exam questions and highlighted answers are collected and verified in our exam. It will help you pass exam in easier way.

A. 1. You are a project owner and need your co-worker to deploy a new version of your application to App Engine. You want to follow Google's recommended practices. Which IAM roles should you grant your co-worker?

Answer(s): C

2. <https://cloud.google.com/iam/docs/understanding-roles>

2. Your company has reserved a monthly budget for your project. You want to be informed automatically of your project spend so that you can take action when you approach the limit. What should you do?

A. A. Link a credit card with a monthly limit equal to your budget.

B. B. Create a budget alert for 50%, 90%, and 100% of your total monthly budget.

C. C. In App Engine Settings, set a daily budget at the rate of 1/30 of your monthly budget.

D. D. In the GCP Console, configure billing export to BigQuery. Create a saved view that queries your total spend.

Answer(s): B

3. https://cloud.google.com/appengine/pricing#spending_limit

<https://cloud.google.com/billing/docs/how-to/budgets>

3. You have a project using BigQuery. You want to list all BigQuery jobs for that project. You want to set this project as the default for the bq command-line tool. What should you do?

A. A. Use “gcloud config set project” to set the default project.

B. B. Use “bq config set project” to set the default project.

C. C. Use “gcloud generate config-url” to generate a URL to the Google Cloud Platform Console to set the default project.

D. D. Use “bq generate config-url” to generate a URL to the Google Cloud Platform Console to set the default project.

Answer(s): A

4. <https://cloud.google.com/bigquery/docs/reference/bq-cli-reference>

<https://cloud.google.com/sdk/gcloud/reference/config/set>

4. Your project has all its Compute Engine resources in the europe-west1 region. You want to set europe-west1 as the default region for gcloud commands. What should you do?

A. A. Use Cloud Shell instead of the command line interface of your device. Launch Cloud Shell after you navigate to a resource in the europe-west1 region. The europe-west1 region will automatically become the default region.

B. B. Use “gcloud config set compute/region europe-west1” to set the default region for future gcloud commands.

C. C. Use “gcloud config set compute/zone europe-west1” to set the default region for future gcloud commands.

D. D. Create a VPN from on-premises to a subnet in europe-west1, and use that connection when executing gcloud commands.

Answer(s): B

5. <https://cloud.google.com/compute/docs/regions-zones/changing-default-zone-region>

5. You developed a new application for App Engine and are ready to deploy it to production. You need to estimate the costs of running your application on Google Cloud Platform as accurately as possible. What should you do?

A. A. Create a YAML file with the expected usage. Pass this file to the “gcloud app estimate” command to get an accurate estimation.

B. B. Multiply the costs of your application when it was in development by the number of expected users to get an accurate estimation.

C. C. Use the pricing calculator for App Engine to get an accurate estimation of the expected charges.

D. D. Create a ticket with Google Cloud Billing Support to get an accurate estimation.

Answer(s): C is correct because this is the proper way to estimate charges

6. Your company processes high volumes of IoT data that are time-stamped. The total data volume can be several petabytes. The data needs to be written and changed at a high speed. You want to use the most performant storage option for your data. Which product should you use?

A. A. Cloud Datastore

B. B. Cloud Storage

C. C. Cloud Bigtable

D. D. BigQuery

Answer(s): C is correct because Cloud Bigtable is the most performant storage option to work with IoT and time series data.

7. <https://cloud.google.com/bigtable/docs/schema-design-time-series>

7. Your application has a large international audience and runs stateless virtual machines within a managed instance group across multiple locations. One feature of the application lets users upload files and share them with other users. Files must be available for 30 days; after that, they are removed from the system entirely. Which storage solution should you choose?

A. A. A Cloud Datastore database.

B. B. A multi-regional Cloud Storage bucket.

C. C. Persistent SSD on virtual machine instances.

D. D. A managed instance group of Filestore servers.

Answer(s): B is correct because buckets can be multi-regional and have lifecycle management.

8. You have a definition for an instance template that contains a web application. You are asked to deploy the application so that it can scale based on the HTTP traffic it receives. What should you do?

A. A. Create a VM from the instance template. Create a custom image from the VM's disk. Export the image to Cloud Storage. Create an HTTP load balancer and add the Cloud Storage bucket as its backend service.

B. B. Create a VM from the instance template. Create an App Engine application in Automatic Scaling mode that forwards all traffic to the VM.

C. C. Create a managed instance group based on the instance template. Configure autoscaling based on HTTP traffic and configure the instance group as the backend service of an HTTP load balancer.

D. D. Create the necessary amount of instances required for peak user traffic based on the instance template. Create an unmanaged instance group and add the instances to that instance group. Configure the instance group as the Backend Service of an HTTP load balancer.

Answer(s): C is correct because a managed instance group can use an instance template to scale based on HTTP traffic.

9. [https://cloud.google.com/compute/docs/instance-](https://cloud.google.com/compute/docs/instance-groups/#managed_instance_groups_and_autoscaling)

[groups/#managed_instance_groups_and_autoscaling](https://cloud.google.com/compute/docs/instance-groups/#managed_instance_groups_and_autoscaling)

<https://cloud.google.com/compute/docs/images/export-image>

<https://cloud.google.com/compute/docs/load-balancing/http/adding-a-backend-bucket-to-content-based-load-balancing>

9. You are creating a Kubernetes Engine cluster to deploy multiple pods inside the cluster. All container logs must be stored in BigQuery for later analysis. You want to follow Google-recommended practices. Which two approaches can you take?

A. A. Turn on Stackdriver Logging during the Kubernetes Engine cluster creation.

B. B. Turn on Stackdriver Monitoring during the Kubernetes Engine cluster creation.

C. C. Develop a custom add-on that uses Cloud Logging API and BigQuery API. Deploy the add-on to your Kubernetes Engine cluster.

D. D. Use the Stackdriver Logging export feature to create a sink to Cloud Storage. Create a Cloud Dataflow job that imports log files from Cloud Storage to BigQuery.

E. E. Use the Stackdriver Logging export feature to create a sink to BigQuery. Specify a filter expression to export log records related to your Kubernetes Engine cluster only.

Answer(s): A is correct because creating a cluster with Stackdriver Logging option will enable all the container logs to be stored in Stackdriver Logging.

10. E is correct because Stackdriver Logging support exporting logs to BigQuery by creating sinks

<https://cloud.google.com/kubernetes-engine/docs/how-to/logging>

https://cloud.google.com/logging/docs/export/configure_export_v2

<https://kubernetes.io/docs/reference/labels-annotations-taints/>

10. You need to create a new Kubernetes Cluster on Google Cloud Platform that can autoscale the number of worker nodes. What should you do?

A. A. Create a cluster on Kubernetes Engine and enable autoscaling on Kubernetes Engine.

B. B. Create a cluster on Kubernetes Engine and enable autoscaling on the instance group of the cluster.

C. C. Configure a Compute Engine instance as a worker and add it to an unmanaged instance group. Add a load balancer to the instance group and rely on the load balancer to create additional Compute Engine instances when needed.

D. D. Create Compute Engine instances for the workers and the master, and install Kubernetes. Rely on Kubernetes to create additional Compute Engine instances when needed.

Answer(s): A is correct because this is the way to set up an autoscaling Kubernetes cluster.

11. <https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler>

11. You have an application server running on Compute Engine in the europe-west1-d zone. You

need to ensure high availability and replicate the server to the europe-west2-c zone using the fewest steps possible. What should you do?

A. A. Create a snapshot from the disk. Create a disk from the snapshot in the europe-west2-c zone. Create a new VM with that disk.

B. B. Create a snapshot from the disk. Create a disk from the snapshot in the europe-west1-d zone and then move the disk to europe-west2-c. Create a new VM with that disk.

C. C. Use "gcloud" to copy the disk to the europe-west2-c zone. Create a new VM with that disk.

D. D. Use "gcloud compute instances move" with parameter "--destination-zone europe-west2-c" to move the instance to the new zone.

Answer(s): A is correct because this makes sure the VM gets replicated in the new zone.

12. Your company has a mission-critical application that serves users globally. You need to select a transactional, relational data storage system for this application. Which two products should you consider

A. A. BigQuery

B. B. Cloud SQL

C. C. Cloud Spanner

D. D. Cloud Bigtable

E. E. Cloud Datastore

Answer(s): B is correct because Cloud SQL is a relational and transactional database in the list.

13. C is correct because Spanner is a relational and transactional database in the list.

13. You have a Kubernetes cluster with 1 node-pool. The cluster receives a lot of traffic and needs to grow. You decide to add a node. What should you do?

A. A. Use "gcloud container clusters resize" with the desired number of nodes.

B. B. Use “kubectl container clusters resize” with the desired number of nodes.

C. C. Edit the managed instance group of the cluster and increase the number of VMs by 1.

D. D. Edit the managed instance group of the cluster and enable autoscaling.

Answer(s): A is correct because this resizes the cluster to the desired number of nodes.

14. You created an update for your application on App Engine. You want to deploy the update without impacting your users. You want to be able to roll back as quickly as possible if it fails. What should you do?

A. A. Delete the current version of your application. Deploy the update using the same version identifier as the deleted version.

B. B. Notify your users of an upcoming maintenance window. Deploy the update in that maintenance window.

C. C. Deploy the update as the same version that is currently running.

D. D. Deploy the update as a new version. Migrate traffic from the current version to the new version.

Answer(s): D is correct because this makes sure there is no downtime and you can roll back the fastest.

15. <https://cloud.google.com/appengine/docs/admin-api/migrating-splitting-traffic>

15. You have created a Kubernetes deployment, called Deployment-A, with 3 replicas on your cluster. Another deployment, called Deployment-B, needs access to Deployment-A. You cannot expose Deployment-A outside of the cluster. What should you do?

A. A. Create a Service of type NodePort for Deployment A and an Ingress Resource for that Service. Have Deployment B use the Ingress IP address.

B. B. Create a Service of type LoadBalancer for Deployment A. Have Deployment B use the Service IP address.

C. C. Create a Service of type LoadBalancer for Deployment A and an Ingress Resource for that Service. Have Deployment B use the Ingress IP address.

D. D. Create a Service of type ClusterIP for Deployment A. Have Deployment B use the Service IP address.

Answer(s): D is correct because this exposes the service on a cluster-internal IP address. Choosing this method makes the service reachable only from within the cluster.

16. <https://kubernetes.io/docs/concepts/services-networking/service/>

16. You need to estimate the annual cost of running a Bigquery query that is scheduled to run nightly. What should you do?

A. A. Use "gcloud query --dry_run" to determine the number of bytes read by the query. Use this number in the Pricing Calculator.

B. B. Use "bq query --dry_run" to determine the number of bytes read by the query. Use this number in the Pricing Calculator.

C. C. Use "gcloud estimate" to determine the amount billed for a single query. Multiply this amount by 365.

D. D. Use "bq estimate" to determine the amount billed for a single query. Multiply this amount by 365.

Answer(s): B is correct because this is the correct way to estimate the yearly BigQuery querying costs.

17. You want to find out who in your organization has Owner access to a project called "my-project". What should you do?

A. A. In the Google Cloud Platform Console, go to the IAM page for your organization and apply the filter "Role:Owner".

B. B. In the Google Cloud Platform Console, go to the IAM page for your project and apply the filter "Role:Owner".

C. C. Use "gcloud iam list-grantable-role --project my-project" from your Terminal.

D. D. Use "gcloud iam list-grantable-role" from Cloud Shell on the project page.

Answer(s): B is correct because this shows you the Owners of the project.

18. You want to create a new role for your colleagues that will apply to all current and future projects created in your organization. The role should have the permissions of the BigQuery Job User and Cloud Bigtable User roles. You want to follow Google's recommended practices. How should you create the new role?

A. A. Use "gcloud iam combine-roles --global" to combine the 2 roles into a new custom role.

B. B. For one of your projects, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role. Use "gcloud iam promote-role" to promote the role from a project role to an organization role.

C. C. For all projects, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role.

D. D. For your organization, in the Google Cloud Platform Console under Roles, select both roles and combine them into a new custom role.

Answer(s): D is correct because this creates a new role with the combined permissions on the organization level.

19. You work in a small company where everyone should be able to view all resources of a specific project. You want to grant them access following Google's recommended practices. What should you do?

A. A. Create a script that uses "gcloud projects add-iam-policy-binding" for all users' email addresses and the Project Viewer role.

B. B. A. Create a script that uses "gcloud iam roles create" for all users' email addresses and the Project Viewer role.

C. C. Create a new Google Group and add all users to the group. Use "gcloud projects add-iam-policy-binding" with the Project Viewer role and Group email address.

D. D. Create a new Google Group and add all members to the group. Use "gcloud iam roles create" with the Project Viewer role and Group email address.

Answer(s): C is correct because Google recommends to use groups where possible.

20. <https://cloud.google.com/sdk/gcloud/reference/iam/>

20. You need to verify the assigned permissions in a custom IAM role. What should you do?

A. A. Use the GCP Console, IAM section to view the information.

B. B. Use the "gcloud init" command to view the information.

C. C. Use the GCP Console, Security section to view the information.

D. D. Use the GCP Console, API section to view the information.

Answer(s): A is correct because this is the correct console area to view permission assigned to a custom role in a particular project.
