



**Free Questions for Associate-Cloud-Engineer by
go4braindumps**

Shared by Preston on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You need to migrate invoice documents stored on-premises to Cloud Storage. The documents have the following storage requirements:

- * Documents must be kept for five years.
- * Up to five revisions of the same invoice document must be stored, to allow for corrections.
- * Documents older than 365 days should be moved to lower cost storage tiers.

You want to follow Google-recommended practices to minimize your operational and development costs. What should you do?

Options:

- A-** Enable retention policies on the bucket, and use Cloud Scheduler to invoke a Cloud Function to move or delete your documents based on their metadata.
- B-** Enable retention policies on the bucket, use lifecycle rules to change the storage classes of the objects, set the number of versions, and delete old files.
- C-** Enable object versioning on the bucket, and use Cloud Scheduler to invoke a Cloud Functions instance to move or delete your documents based on their metadata.
- D-** Enable object versioning on the bucket, use lifecycle conditions to change the storage class of the objects, set the number of versions, and delete old files.

Answer:

B

Question 2

Question Type: MultipleChoice

You need to extract text from audio files by using the Speech-to-Text API. The audio files are pushed to a Cloud Storage bucket. You need to implement a fully managed, serverless compute solution that requires authentication and aligns with Google-recommended practices. You want to automate the call to the API by submitting each file to the API as the audio file arrives in the bucket. What should you do?

Options:

- A-** Run a Kubernetes job to scan the bucket regularly for incoming files, and call the Speech-to-Text API for each unprocessed file.
- B-** Create an App Engine standard environment triggered by Cloud Storage bucket events to submit the file URI to the Google Speech-to-Text API.
- C-** Run a Python script by using a Linux cron job in Compute Engine to scan the bucket regularly for incoming files, and call the Speech-to-Text API for each unprocessed file.

D- Create a Cloud Function triggered by Cloud Storage bucket events to submit the file URI to the Google Speech-to-Text API.

Answer:

B

Question 3

Question Type: MultipleChoice

Your manager asks you to deploy a workload to a Kubernetes cluster. You are not sure of the workloads resource requirements or how the requirements might vary depending on usage patterns, external dependencies, or other factors. You need a solution that makes cost-effective recommendations regarding CPU and memory requirements, and allows the workload to function consistently in any situation. You want to follow Google-recommended practices. What should you do?

Options:

- A-** Configure the Horizontal Pod Autoscaler for availability, and configure the cluster autoscaler for suggestions.
- B-** Configure the Horizontal Pod Autoscaler for availability, and configure the Vertical Pod Autoscaler recommendations for suggestions.
- C-** Configure the Vertical Pod Autoscaler recommendations for availability, and configure the Cluster autoscaler for suggestions.

D- Configure the Vertical Pod Autoscaler recommendations for availability, and configure the Horizontal Pod Autoscaler for suggestions.

Answer:

D

Question 4

Question Type: MultipleChoice

Your company has multiple projects linked to a single billing account in Google Cloud. You need to visualize the costs with specific metrics that should be dynamically calculated based on company-specific criteria.

a. You want to automate the process. What should you do?

Options:

A- In the Google Cloud console, visualize the costs related to the projects in the Reports section.

B- In the Google Cloud console, visualize the costs related to the projects in the Cost breakdown section.

C- In the Google Cloud console, use the export functionality of the Cost table. Create a Looker Studio dashboard on top of the CSV export.

D- Configure Cloud Billing data export to BigQuery for the billing account. Create a Looker Studio dashboard on top of the BigQuery export.

Answer:

D

Explanation:

Cloud Billing export to BigQuery enables you to export detailed Google Cloud billing data (such as usage, cost estimates, and pricing data) automatically throughout the day to a BigQuery dataset that you specify. Then you can access your Cloud Billing data from BigQuery for detailed analysis, or use a tool like Looker Studio to visualize your data.

Question 5

Question Type: MultipleChoice

You are building a multi-player gaming application that will store game information in a database. As the popularity of the application increases, you are concerned about delivering consistent performance. You need to ensure an optimal gaming performance for global users, without increasing the management complexity. What should you do?

Options:

- A- Use Cloud SQL database with cross-region replication to store game statistics in the EU, US, and APAC regions.
- B- Use Cloud Spanner to store user data mapped to the game statistics.
- C- Use BigQuery to store game statistics with a Redis on Memorystore instance in the front to provide global consistency.
- D- Store game statistics in a Bigtable database partitioned by username.

Answer:

B

Question 6

Question Type: MultipleChoice

You deployed an application on a managed instance group in Compute Engine. The application accepts Transmission Control Protocol (TCP) traffic on port 389 and requires you to preserve the IP address of the client who is making a request. You want to expose the application to the internet by using a load balancer. What should you do?

Options:

- A-** Expose the application by using an external TCP Network Load Balancer.
- B-** Expose the application by using a TCP Proxy Load Balancer.
- C-** Expose the application by using an SSL Proxy Load Balancer.
- D-** Expose the application by using an internal TCP Network Load Balancer.

Answer:

B

Question 7

Question Type: MultipleChoice

You have a managed instance group comprised of preemptible VM's. All of the VM's keep deleting and recreating themselves every minute. What is a possible cause of this behavior?

Options:

A- Your zonal capacity is limited, causing all preemptible VM's to be shutdown to recover capacity. Try deploying your group to another zone.

- B-** You have hit your instance quota for the region.
- C-** Your managed instance group's VM's are toggled to only last 1 minute inpreemptible settings.
- D-** Your managed instance group's health check is repeatedly failing, either to amisconfigured health check or misconfigured firewall rules not allowing the healthcheck to access the instance

Answer:

D

Explanation:

as the instances (normal or preemptible) would be terminated and relaunched if the health check fails either due to application not configured properly or the instances firewall do not allow health check to happen.

GCP provides health check systems that connect to virtual machine (VM) instances on a configurable, periodic basis. Each connection attempt is called a probe. GCP records the success or failure of each probe.

Health checks and load balancers work together. Based on a configurable number of sequential successful or failed probes, GCP computes an overall health state for each VM in the load balancer. VMs that respond successfully for the configured number of times are considered healthy. VMs that fail to respond successfully for a separate number of times are unhealthy.

GCP uses the overall health state of each VM to determine its eligibility for receiving new requests. In addition to being able to configure probe frequency and health state thresholds, you can configure the criteria that define a successful probe.

Question 8

Question Type: MultipleChoice

You have files in a Cloud Storage bucket that you need to share with your suppliers. You want to restrict the time that the files are available to your suppliers to 1 hour. You want to follow Google recommended practices. What should you do?

Options:

- A-** Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command `gsutil signurl -m 1h gs:///*`.
- B-** Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command `gsutil signurl -d 1h gs:///**`.
- C-** Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command `gsutil signurl -p 60m gs:///`.
- D-** Create a JSON key for the Default Compute Engine Service Account. Execute the command `gsutil signurl -t 60m gs:///***`

Answer:

B

Explanation:

This command correctly specifies the duration that the signed url should be valid for by using the -d flag. The default is 1 hour so omitting the -d flag would have also resulted in the same outcome. Times may be specified with no suffix (default hours), or with s = seconds, m = minutes, h = hours, d = days. The max duration allowed is 7d.

Ref:<https://cloud.google.com/storage/docs/gsutil/commands/signurl>

Question 9

Question Type: MultipleChoice

You have been asked to migrate a docker application from datacenter to cloud. Your solution architect has suggested uploading docker images to GCR in one project and running an application in a GKE cluster in a separate project. You want to store images in the project img-278322 and run the application in the project prod-278986. You want to tag the image as acme_track_n_trace:v1. You want to follow Google-recommended practices. What should you do?

Options:

- A- Run `gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace`
- B- Run `gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace:v1`
- C- Run `gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace`

D- Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace:v1

Answer:

B

Explanation:

Run gcloud builds submit tag gcr.io/img-278322/acme_track_n_trace:v1. is the right answer.

This command correctly tags the image as acme_track_n_trace:v1 and uploads the image to the img-278322 project.

Ref:<https://cloud.google.com/sdk/gcloud/reference/builds/submit>

Question 10

Question Type: MultipleChoice

You have a number of applications that have bursty workloads and are heavily dependent on topics to decouple publishing systems from consuming systems. Your company would like to go serverless to enable developers to focus on writing code without worrying about infrastructure. Your solution architect has already identified Cloud Pub/Sub as a suitable alternative for decoupling systems. You have been asked to identify a suitable GCP Serverless service that is easy to use with Cloud Pub/Sub. You want the ability to scale down to

zero when there is no traffic in order to minimize costs. You want to follow Google recommended practices. What should you suggest?

Options:

- A- Cloud Run for Anthos
- B- Cloud Run
- C- App Engine Standard
- D- Cloud Functions.

Answer:

D

Explanation:

Cloud Functions is Google Cloud's event-driven serverless compute platform that lets you run your code locally or in the cloud without having to provision servers. Cloud Functions scales up or down, so you pay only for compute resources you use. Cloud Functions have excellent integration with Cloud Pub/Sub, lets you scale down to zero and is recommended by Google as the ideal serverless platform to use when dependent on Cloud Pub/Sub. 'If you're building a simple API (a small set of functions to be accessed via HTTP or Cloud Pub/Sub), we recommend using Cloud Functions.' Ref:<https://cloud.google.com/serverless-options>

Question 11

Question Type: MultipleChoice

You created a Kubernetes deployment by running `kubectl run nginx image=nginx replicas=1`. After a few days, you decided you no longer want this deployment. You identified the pod and deleted it by running `kubectl delete pod`. You noticed the pod got recreated.

```
$ kubectl get pods
```

```
NAME READY STATUS RESTARTS AGE
```

```
nginx-84748895c4-nqqmt 1/1 Running 0 9m41s
```

```
$ kubectl delete pod nginx-84748895c4-nqqmt
```

```
pod nginx-84748895c4-nqqmt deleted
```

```
$ kubectl get pods
```

```
NAME READY STATUS RESTARTS AGE
```

```
nginx-84748895c4-k6bzl 1/1 Running 0 25s
```

What should you do to delete the deployment and avoid pod getting recreated?

Options:

- A- `kubectl delete deployment nginx`
- B- `kubectl delete --deployment=nginx`
- C- `kubectl delete pod nginx-84748895c4-k6bz1 --no-restart 2`
- D- `kubectl delete inginx`

Answer:

A

Explanation:

This command correctly deletes the deployment. Pods are managed by kubernetes workloads (deployments). When a pod is deleted, the deployment detects the pod is unavailable and brings up another pod to maintain the replica count. The only way to delete the workload is by deleting the deployment itself using the `kubectl delete deployment` command.

```
$ kubectl delete deployment nginx
```

```
deployment.apps/nginx deleted
```

Ref:<https://kubernetes.io/docs/reference/kubectl/cheatsheet/#deleting-resources>

To Get Premium Files for Associate-Cloud-Engineer Visit

<https://www.p2pexams.com/products/associate-cloud-engineer>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/associate-cloud-engineer>

