

Get Oct-2023 Download Latest & Valid Questions For Google Associate-Cloud-Engineer exam [Q61-Q76]



Get Oct-2023 Download Latest & Valid Questions For Google Associate-Cloud-Engineer exam [Q61-Q76]

Get Oct-2023 Download Latest & Valid Questions For Google Associate-Cloud-Engineer exam
Ensure Success With Updated Verified Associate-Cloud-Engineer Exam Dumps

QUESTION 61

You have a project for your App Engine application that serves a development environment. The required testing has succeeded and you want to create a new project to serve as your production environment. What should you do?

- * Use gcloud to create the new project, and then deploy your application to the new project.
- * Use gcloud to create the new project and to copy the deployed application to the new project.
- * Create a Deployment Manager configuration file that copies the current App Engine deployment into a new project.
- * Deploy your application again using gcloud and specify the project parameter with the new project name to create the new project.

You can deploy to a different project by using -project flag.

By default, the service is deployed the current project configured via:

```
$ gcloud config set core/project PROJECT
```

To override this value for a single deployment, use the -project flag:

```
$ gcloud app deploy ~/my_app/app.yaml -project=PROJECT
```

Ref: <https://cloud.google.com/sdk/gcloud/reference/app/deploy>

QUESTION 62

You need to update a deployment in Deployment Manager without any resource downtime in the deployment.

Which command should you use?

- * `gcloud deployment-manager deployments create –config <deployment-config-path>`
- * `gcloud deployment-manager deployments update –config <deployment-config-path>`
- * `gcloud deployment-manager resources create –config <deployment-config-path>`
- * `gcloud deployment-manager resources update –config <deployment-config-path>`

Explanation/Reference: <https://cloud.google.com/sdk/gcloud/reference/deployment-manager/deployments/update>

QUESTION 63

You deployed an LDAP server on Compute Engine that is reachable via TLS through port 636 using UDP. You want to make sure it is reachable by clients over that port. What should you do?

- * Add the network tag allow-udp-636 to the VM instance running the LDAP server.
- * Create a route called allow-udp-636 and set the next hop to be the VM instance running the LDAP server.
- * Add a network tag of your choice to the instance. Create a firewall rule to allow ingress on UDP port 636 for that network tag.
- * Add a network tag of your choice to the instance running the LDAP server. Create a firewall rule to allow egress on UDP port 636 for that network tag.

QUESTION 64

You manage three Google Cloud projects with the Cloud Monitoring API enabled. You want to follow Google-recommended practices to visualize CPU and network metrics for all three projects together. What should you do?

- * 1. Create a Cloud Monitoring Dashboard
- 2. Collect metrics and publish them into the Pub/Sub topics
- 3. Add CPU and network Charts (or each of the three projects)
 - * 1. Create a Cloud Monitoring Dashboard.
- 2. Select the CPU and Network metrics from the three projects.
- 3. Add CPU and network Charts for each of the three projects.
 - * 1 Create a Service Account and apply roles/viewer on the three projects
- 2. Collect metrics and publish them to the Cloud Monitoring API
- 3. Add CPU and network Charts for each of the three projects.
 - * 1. Create a fourth Google Cloud project

2 Create a Cloud Workspace from the fourth project and add the other three projects

QUESTION 65

You host a static website on Cloud Storage. Recently, you began to include links to PDF files on this site. Currently, when users click on the links to these PDF files, their browsers prompt them to save the file onto their local system. Instead, you want the clicked PDF files to be displayed within the browser window directly, without prompting the user to save the file locally. What should you do?

- * Enable Cloud CDN on the website frontend.
- * Enable `‘Share publicly’` on the PDF file objects.
- * Set Content-Type metadata to `application/pdf` on the PDF file objects.
- * Add a label to the storage bucket with a key of Content-Type and value of `application/pdf`.

QUESTION 66

You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements.

You want to select a cost-efficient option. Which storage option should you use?

- * Cold Storage
- * Nearline Storage
- * Regional Storage
- * Multi-Regional Storage

Explanation

Nearline, Coldline, and Archive offer ultra low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-effective storage option for long-term preservation of data.

Coldline is also ideal for cold storage-data your business expects to touch less than once a quarter. For warmer storage, choose Nearline: data you expect to access less than once a month, but possibly multiple times throughout the year. All storage classes are available across all GCP regions and provide unparalleled sub-second access speeds with a consistent API.

QUESTION 67

Users of your application are complaining of slowness when loading the application. You realize the slowness is because the App Engine deployment serving the application is deployed in `us-central` whereas all users of this application are closest to `europa-west3`. You want to change the region of the App Engine application to `europa-west3` to minimize latency. What's the best way to change the App Engine region?

- * Create a new project and create an App Engine instance in `europa-west3`
- * Use the `gcloud app region set` command and supply the name of the new region.
- * From the console, under the App Engine page, click edit, and change the region drop-down.
- * Contact Google Cloud Support and request the change.

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it cant be changed. The only way is to create a new project and create an App Engine instance in `europa-west3`, send all user traffic to this instance and delete the app engine instance in `us-central`.

Ref: <https://cloud.google.com/appengine/docs/locations>

QUESTION 68

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?

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

QUESTION 69

Your organization uses G Suite for communication and collaboration. All users in your organization have a G Suite account. You want to grant some G Suite users access to your Cloud Platform project. What should you do?

- * Enable Cloud Identity in the GCP Console for your domain.
- * Grant them the required IAM roles using their G Suite email address.
- * Create a CSV sheet with all users' email addresses. Use the gcloud command line tool to convert them into Google Cloud Platform accounts.
- * In the G Suite console, add the users to a special group called cloud-console-users@yourdomain.com.

Rely on the default behavior of the Cloud Platform to grant users access if they are members of this group.

QUESTION 70

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?

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

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 71

You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

- * Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
- * Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
- * Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
- * Use the Reports view in the Cloud Billing Console to view the desired cost information.

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

QUESTION 72

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?

- * Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
- * Select Cloud SQL (MySQL). Select the create failover replicas option.
- * Select Cloud Spanner. Set up your instance with 2 nodes.

* Select Cloud Spanner. Set up your instance as multi-regional.

Reference:

<https://cloud.google.com/sql/docs/mysql/backup-recovery/restore>

QUESTION 73

A photo-sharing website running on AWS allows users to generate thumbnail images of photos stored in Amazon S3. An Amazon DynamoDB table maintains the locations of photos, and thumbnails are easily re-created from the originals if they are accidentally deleted.

How should the thumbnail images be stored to ensure the LOWEST cost?

- * Amazon S3 Standard-Infrequent Access (S3 Standard-IA) with cross-region replication
- * Amazon S3
- * Amazon Glacier
- * Amazon S3 with cross-region replication

QUESTION 74

You're about to deploy your team's App Engine application. They're using the Go runtime with a Standard Environment. Which command should you use to deploy the application?

- * `gcloud app deploy app.yaml`
- * `gcloud app-engine apply app.yaml`
- * `gcloud app apply app.yaml`
- * `gcloud app-engine deploy app.yaml`

QUESTION 75

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?

- * Create a Service of type NodePort for Deployment A and an Ingress Resource for that Service.

Have Deployment B use the Ingress IP address.

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

A is not correct because this exposes Deployment A over the public internet.

B is not correct because LoadBalancer will expose the service publicly.

C is not correct because this exposes the service externally using a cloud provider's load balancer, and Ingress can work only with nodeport, not LoadBalancer.

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.

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

QUESTION 76

You are configuring service accounts for an application that spans multiple projects. Virtual machines (VMs) running in the web-applications project need access to BigQuery datasets in crm-databases-proj. You want to follow Google-recommended practices to give access to the service account in the web-applications project.

What should you do?

- * Give `project owner` for web-applications appropriate roles to crm-databases- proj
- * Give `project owner` role to crm-databases-proj and the web-applications project.
- * Give `project owner` role to crm-databases-proj and bigquery.dataViewer role to web-applications.
- * Give bigquery.dataViewer role to crm-databases-proj and appropriate roles to web-applications.

Exam Materials for You to Prepare & Pass Associate-Cloud-Engineer Exam:

<https://www.examlabs.com/Google/Google-Cloud-Certified/best-Associate-Cloud-Engineer-exam-dumps.html>