Pass Google Cloud Certified Associate-Cloud-Engineer exam [Jan 25, 2025 Updated 290 Questions [Q71-Q94

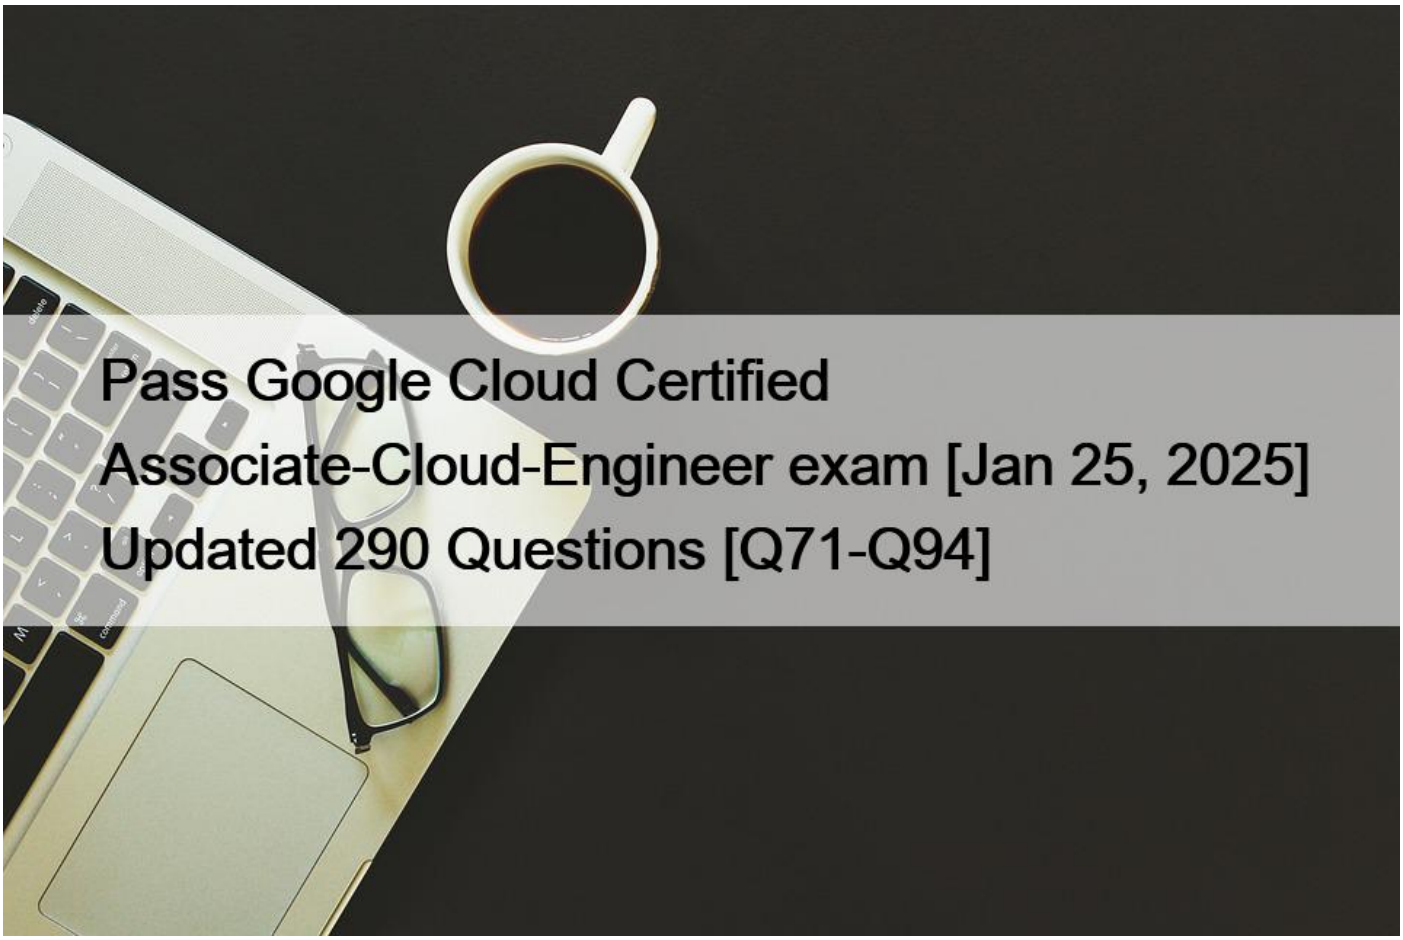

Pass Google Cloud Certified Associate-Cloud-Engineer exam [Jan 25, 2025] Updated 290 Questions
Google Associate-Cloud-Engineer Actual Questions and 100% Cover Real Exam Questions

Google Associate-Cloud-Engineer certification is designed to validate the skills of individuals in the field of cloud computing. Google Associate Cloud Engineer Exam certification is awarded to those who have successfully passed the Google Associate Cloud Engineer exam that focuses on the Google Cloud Platform (GCP). Google Associate Cloud Engineer Exam certification is an industry-recognized credential that can help IT professionals advance in their careers and succeed in the competitive job market.

Q71. You need to monitor resources that are distributed over different projects in Google Cloud Platform. You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?
* Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
* For each project, create a Stackdriver account. In each project, create a service account for that project and grant it the role of Stackdriver Account Editor in all other projects.
* Configure a single Stackdriver account, and link all projects to the same account.
* Configure a single Stackdriver account for one of the projects. In Stackdriver, create a Group and add the other project names as

criteria for that Group.
Explanation

When you intially click on Monitoring(Stackdriver Monitoring) it creates a workspac(a stackdriver account) linked to the ACTIVE(CURRENT) Project from which it was clicked.

Now if you change the project and again click onto Monitoring it would create an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don’t want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).

If you have accidently created two diff workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.

If we have only one workspace and two projects we can simply add other GCP Project under Monitoring > Settings > GCP Projects > Add GCP Projects.

https://cloud.google.com/monitoring/settings/multiple-projects

Nothing about groups https://cloud.google.com/monitoring/settings?hl=en

Q72. A Solutions Architect is building an online shopping application where users will be able to browse items, add items to a cart, and purchase the items. Images of items will be stored in Amazon S3 buckets organized by item category. When an item is no longer available for purchase, the item image will be deleted from the S3 bucket.

Occasionally, during testing, item images deleted from the S3 bucket are still visible to some users.

What is a flaw in this design approach?
* Defining S3 buckets by item may cause partition distribution errors, which will impact performance.
* Amazon S3 DELETE requests are eventually consistent, which may cause other users to view items that have already been purchased
* Amazon S3 DELETE requests apply a lock to the S3 bucket during the operation, causing other users to be blocked
* Using Amazon S3 for persistence exposes the application to a single point of failure
Explanation

Q73. You deployed a new application inside your Google Kubernetes Engine cluster using the YAML file specified below.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp-deployment
spec:
  selector:
    matchLabels:
      app: myapp
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp
    spec:
      containers:
      - name: myapp
        image: myapp:1.1
        ports:
        - containerPort: 80
```

```
apiVersion: v1
kind: Service
metadata:
  name: myapp-service
spec:
  ports:
  - port: 8000
    targetPort: 80
    protocol: TCP
  selector:
    app: myapp
```

You check the status of the deployed pods and notice that one of them is still in PENDING status:

```
kubectl get pods -l app=myapp
NAME                                  READY     STATUS    RESTART   AGE
myapp-deployment-58ddbbb995-lp86m     0/1       Pending   0         9m
myapp-deployment-58ddbbb995-qjpkg     1/1       Running   0         9m
```


You want to find out why the pod is stuck in pending status. What should you do?
* Review details of the myapp-service Service object and check for error messages.
* Review details of the myapp-deployment Deployment object and check for error messages.
* Review details of myapp-deployment-58ddbbb995-lp86m Pod and check for warning messages.
* View logs of the container in myapp-deployment-58ddbbb995-lp86m pod and check for warning messages.
https://kubernetes.io/docs/tasks/debug-application-cluster/debug-application/#debugging-pods

Q74. 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.
https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP

/MIME_Types#importance_of_setting_the_correct_mime_type

Q75. You have an application that uses Cloud Spanner as a database backend to keep current state information about users. Cloud Bigtable logs all events triggered by users. You export Cloud Spanner data to Cloud Storage during daily backups. One of your analysts asks you to join data from Cloud Spanner and Cloud Bigtable for specific users. You want to complete this ad hoc request as efficiently as possible. What should you do?
* Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
* Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
* Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
* Create two separate BigQuery external tables on Cloud Storage and Cloud Bigtable. Use the BigQuery console to join these tables through user fields, and apply appropriate filters.

Q76. You significantly changed a complex Deployment Manager template and want to confirm that the dependencies of all defined resources are properly met before committing it to the project. You want the most rapid feedback on your changes. What should you do?
* Use granular logging statements within a Deployment Manager template authored in Python.
* Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.
* Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.
* Execute the Deployment Manager template using the ?preview option in the same project, and observe the state of interdependent resources.

Deployment Manager provides the preview feature to check on what resources would be created.

https://cloud.google.com/deployment-manager/docs/deployments/updating-deployments

Q77. 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.
https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_th

Q78. You need to manage a Cloud Spanner Instance for best query performance. Your instance in production runs in a single Google Cloud region. You need to improve performance in the shortest amount of time. You want to follow Google best practices for service configuration. What should you do?
* Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45% If you exceed this threshold, add nodes lo your instance.
* Create an alert in Cloud Monitoring to alert when the percentage to high priority CPU utilization reaches 45% Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage
* Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65% If you exceed this threshold, add nodes to your instance
* Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%. Use database query statistics to identity queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.
https://cloud.google.com/spanner/docs/cpu-utilization#recommended-max

Q79. You have deployed an application on a single Compute Engine instance. The application writes logs to disk. Users start reporting errors with the application. You want to diagnose the problem. What should you do?
* Navigate to Cloud Logging and view the application logs.
* Connect to the instance’s serial console and read the application logs.
* Configure a Health Check on the instance and set a Low Healthy Threshold value.
* Install and configure the Cloud Logging Agent and view the logs from Cloud Logging.
Reference:

Cloud Loging knows nothing about applications installed on the system without an agent collecting logs. Using the serial console is not a best-practice and is impractical on a large scale.

The VM images for Compute Engine and Amazon Elastic Compute Cloud (EC2) don’t include the Logging agent, so you must complete these steps to install it on those instances. The agent runs under both Linux and Windows. Source:
https://cloud.google.com/logging/docs/agent/logging/installation

Q80. You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?
* Deploy the new version in the same application and use the –migrate option.
* Deploy the new version in the same application and use the –splits option to give a weight of 99 to the current version and a weight of 1 to the new version.
* Create a new App Engine application in the same project. Deploy the new version in that application.

Use the App Engine library to proxy 1% of the requests to the new version.
* Create a new App Engine application in the same project. Deploy the new version in that application.Configure your network load balancer to send 1% of the traffic to that new application.
Explanation

https://cloud.google.com/appengine/docs/standard/python/splitting-traffic#gcloud

Q81. You have just created a new project which will be used to deploy a globally distributed application. You will use Cloud Spanner for data storage. You want to create a Cloud Spanner instance. You want to perform the first step in preparation of creating the instance.

What should you do?
* Grant yourself the IAM role of Cloud Spanner Admin
* Create a new VPC network with subnetworks in all desired regions
* Configure your Cloud Spanner instance to be multi-regional
* Enable the Cloud Spanner API

Q82. You need to produce a list of the enabled Google Cloud Platform APIs for a GCP project using the gcloud command line in the Cloud Shell. The project name is my-project. What should you do?
* Run gcloud projects listto get the project ID, and then run gcloud services list –

-project <project ID>.
* Run gcloud initto set the current project to my-project, and then run gcloud services list –available.
* Run gcloud infoto view the account value, and then run gcloud services list — account <Account>.
* Run gcloud projects describe <project ID>to verify the project value, and then run gcloud services list –available.

Q83. You have a Compute Engine instance hosting an application used between 9 AM and 6 PM on weekdays.

You want to back up this instance daily for disaster recovery purposes. You want to keep the backups for 30 days. You want the Google-recommended solution with the least management overhead and the least number of services. What should you do?
* 1. Update your instances’ metadata to add the following value: snapshot-schedule: 0 1 * * *

2. Update your instances’ metadata to add the following value: snapshot-retention: 30
* 1. In the Cloud Console, go to the Compute Engine Disks page and select your instance’s disk.

2. In the Snapshot Schedule section, select Create Schedule and configure the following parameters:

-Schedule frequency: Daily

-Start time: 1:00 AM – 2:00 AM

-Autodelete snapshots after 30 days
* 1. Create a Cloud Function that creates a snapshot of your instance’s disk.

2.Create a Cloud Function that deletes snapshots that are older than 30 days.

3.Use Cloud Scheduler to trigger both Cloud Functions daily at 1:00 AM.
* 1. Create a bash script in the instance that copies the content of the disk to Cloud Storage.

2.Create a bash script in the instance that deletes data older than 30 days in the backup Cloud Storage bucket.

3.Configure the instance’s crontab to execute these scripts daily at 1:00 AM.
Creating scheduled snapshots for persistent disk This document describes how to create a snapshot schedule to regularly and automatically back up your zonal and regional persistent disks. Use snapshot schedules as a best practice to back up your Compute Engine workloads. After creating a snapshot schedule, you can apply it to one or more persistent disks.
https://cloud.google.com/compute/docs/disks/scheduled-snapshots

Q84. You want to add a new auditor to a Google Cloud Platform project. The auditor should be allowed to read, but not modify, all project items.

How should you configure the auditor’s permissions?
* Create a custom role with view-only project permissions. Add the user’s account to the custom role.
* Create a custom role with view-only service permissions. Add the user’s account to the custom role.
* Select the built-in IAM project Viewer role. Add the user’s account to this role.
* Select the built-in IAM service Viewer role. Add the user’s account to this role.
https://cloud.google.com/resource-manager/docs/access-control-proj

Q85. Your preview application, deployed on a single-zone Google Kubernetes Engine (GKE) cluster in us-centrall, has gained popularity. You are now ready to make the application generally available. You need to deploy the application to production while ensuring high availability and resilience. You also want to follow Google- recommended practices. What should you do?
* Use the gcloud container clusters create command with the options–enable-multi-networking and– enable-autoscaling to create an autoscaling zonal cluster and deploy the application to it.
* Use the gcloud container clusters create-auto command to create an autopilot cluster and deploy the application to it.
* Use the gcloud container clusters update command with the option-region us-centrall to update the cluster and deploy the application to it.
* Use the gcloud container clusters update command with the option-node-locations us-centrall-a,us- centrall-b to update the cluster and deploy the application to the nodes.

Q86. Your organization has user identities in Active Directory. Your organization wants to use Active Directory as their source of truth for identities. Your organization wants to have full control over the Google accounts used by employees for all Google services, including your Google Cloud Platform (GCP) organization.

What should you do?
* Use Google Cloud Directory Sync (GCDS) to synchronize users into Cloud Identity.
* Use the cloud Identity APIs and write a script to synchronize users to Cloud Identity.
* Export users from Active Directory as a CSV and import them to Cloud Identity via the Admin Console.
* Ask each employee to create a Google account using self signup. Require that each employee use their company email address and password.
Google Cloud Directory Sync enables administrators to synchronize users, groups and other data from an Active Directory/LDAP service to their Google Cloud domain directory

https://tools.google.com/dlpage/dirsync/

Q87. 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?
* 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.

* 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.
* 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.
* 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.
We recommend collecting users with the same responsibilities into groups and assigning IAM roles to the groups rather than to individual users. For example, you can create a “data scientist” group and assign appropriate roles to enable interaction with BigQuery and Cloud Storage. When a new data scientist joins your team, you can simply add them to the group and they will inherit the defined permissions. You can create and manage groups through the Admin Console.

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

Q88. 30. You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images. The microservice responsible for the image rendering requires a large amount of CPU time compared to the memory it requires. The other microservices are workloads that are optimized for n1-standard machine types. You need to optimize your cluster so that all workloads are using resources as efficiently as possible. What should you do?
* Assign the pods of the image rendering microservice a higher pod priority than the older microservices
* Create a node pool with compute-optimized machine type nodes for the image rendering microservice Use the node pool with general-purpose machine type nodes for the other microservices
* Use the node pool with general-purpose machine type nodes for lite mage rendering microservice Create a nodepool with compute-optimized machine type nodes for the other microservices
* Configure the required amount of CPU and memory in the resource requests specification of the image rendering microservice deployment Keep the resource requests for the other microservices at the default

Q89. You built an application on your development laptop that uses Google Cloud services. Your application uses Application Default Credentials for authentication and works fine on your development laptop. You want to migrate this application to a Compute Engine virtual machine (VM) and set up authentication using Google- recommended practices and minimal changes. What should you do?
* Assign appropriate access for Google services to the service account used by the Compute Engine VM.
* Create a service account with appropriate access for Google services, and configure the application to use this account.
* Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
* Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.
In general, Google recommends that each instance that needs to call a Google API should run as a service account with the minimum permissions necessary for that instance to do its job. In practice, this means you should configure service accounts for your instances with the following process: Create a new service account rather than using the Compute Engine default service account. Grant IAM roles to that service account for only the resources that it needs. Configure the instance to run as that service account. Grant the instance the https://www.googleapis.com/auth/cloud-platform scope to allow full access to all Google Cloud APIs, so that the IAM permissions of the instance are completely determined by the IAM roles of the service account. Avoid granting more access than necessary and regularly check your service account permissions to make sure they are up-to-date.
https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances#best_practices

Q90. You are planning to migrate the following on-premises data management solutions to Google Cloud:

* One MySQL cluster for your main database

* Apache Kafka for your event streaming platform

* One Cloud SOL for PostgreSOL database for your analytical and reporting needs You want to implement Google-recommended solutions for the migration. You need to ensure that the new solutions provide global scalability and require minimal operational and infrastructure management. What should you do?
* Migrate from MySQL to Cloud SQL, from Kafka to Memorystore, and from Cloud SQL for PostgreSQL to Cloud SQL
* Migrate from MySQL to Cloud Spanner, from Kafka to Memorystore, and from Cloud SOL for PostgreSQL to Cloud SQL
* Migrate from MySQL to Cloud SOL, from Kafka to Pub/Sub, and from Cloud SOL for PostgreSQL to BigQuery.
* Migrate from MySQL to Cloud Spanner, from Kafka to Pub/Sub. and from Cloud SQL for PostgreSQL to BigQuery

Q91. You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your organizational policies require you to use virtual machines directly. You need to ensure that the application scaling is operationally efficient and completed as quickly as possible. What should you do?
* Create a Google Kubernetes Engine cluster, and use horizontal pod autoscaling to scale the application.
* Create an instance template, and use the template in a managed instance group with autoscaling configured.
* Create an instance template, and use the template in a managed instance group that scales up and down based on the time of day.
* Use a set of third-party tools to build automation around scaling the application up and down, based on Stackdriver CPU usage monitoring.

Managed instance groups offer autoscaling capabilities that let you automatically add or delete instances from a managed instance group based on increases or decreases in load (CPU Utilization in this case). Autoscaling helps your apps gracefully handle increases in traffic and reduce costs when the need for resources is lower. You define the autoscaling policy and the autoscaler performs automatic scaling based on the measured load (CPU Utilization in this case). Autoscaling works by adding more instances to your instance group when there is more load (upscaling), and deleting instances when the need for instances is lowered (downscaling).
Ref: https://cloud.google.com/compute/docs/autoscaler

Q92. You want to configure a solution for archiving data in a Cloud Storage bucket. The solution must be cost-effective. Data with multiple versions should be archived after 30 days. Previous versions are accessed once a month for reporting. This archive data is also occasionally updated at month-end. What should you do?
* Add a bucket lifecycle rule that archives data with newer versions after 30 days to Coldline Storage.
* Add a bucket lifecycle rule that archives data with newer versions after 30 days to Nearline Storage.
* Add a bucket lifecycle rule that archives data from regional storage after 30 days to Coldline Storage.
* Add a bucket lifecycle rule that archives data from regional storage after 30 days to Nearline Storage.

Q93. You are building a backend service for an ecommerce platform that will persist transaction data from mobile and web clients. After the platform is launched, you expect a large volume of global transactions. Your business team wants to run SQL queries to analyze the dat a. You need to build a highly available and scalable data store for the platform. What should you do?
* Create a multi-region Cloud Spanner instance with an optimized schema.
* Create a multi-region Firestore database with aggregation query enabled.
* Create a multi-region Cloud SQL for PostgreSQL database with optimized indexes.
* Create a multi-region BigQuery dataset with optimized tables.

Q94. You have an object in a Cloud Storage bucket that you want to share with an external company.

The object contains sensitive data. You want access to the content to be removed after four hours. The external company does not have a Google account to which you can grant specific user-based access privileges. You want to use the most secure method that requires the fewest steps. What should you do?
* Create a signed URL with a four-hour expiration and share the URL with the company.
* Set object access to `public’ and use object lifecycle management to remove the object after four hours.

* Configure the storage bucket as a static website and furnish the object’s URL to the company.

Delete the object from the storage bucket after four hours.
* Create a new Cloud Storage bucket specifically for the external company to access.

Copy the object to that bucket. Delete the bucket after four hours have passed.
Signed URLs are used to give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account.

https://cloud.google.com/storage/docs/access-control/signed-urls

Google Associate-Cloud-Engineer certification is ideal for individuals who are looking to start a career in cloud computing or enhance their existing skills. Google Associate Cloud Engineer Exam certification provides a strong foundation in GCP and is recognized by employers as a valuable asset. It is also a prerequisite for more advanced certifications such as the Professional Cloud Architect and Professional Cloud Developer.

Google Associate-Cloud-Engineer Real 2025 Braindumps Mock Exam Dumps:
https://www.actualtests4sure.com/Associate-Cloud-Engineer-test-questions.html]