

## PDF Download Free of Associate-Cloud-Engineer Valid Practice Test Questions [Q64-Q82]



## PDF Download Free of Associate-Cloud-Engineer Valid Practice Test Questions [Q64-Q82]

PDF Download Free of Associate-Cloud-Engineer Valid Practice Test Questions  
Associate-Cloud-Engineer Test Engine files, Associate-Cloud-Engineer Dumps PDF

Upon passing the Google Associate Cloud Engineer certification exam, you will receive a certificate that demonstrates your proficiency in operating and deploying applications, infrastructure, and services on the Google Cloud Platform. Google Associate Cloud Engineer Exam certification can help boost your career prospects, as it is recognized by employers and demonstrates that you have a strong understanding of cloud computing and the Google Cloud Platform.

**NO.64** A media company asked a Solutions Architect to design a highly available storage solution to serve as a centralized document store for their Amazon EC2 instances. The storage solution needs to be POSIX- compliant, scale dynamically, and be able to serve up to 100 concurrent EC2 instances.

Which solution meets these requirements?

- \* Create an Amazon S3 bucket and store all of the documents in this bucket.
- \* Create an Amazon EBS volume and allow multiple users to mount that volume to their EC2 instance(s).

- \* Use Amazon Glacier to store all of the documents.
- \* Create an Amazon Elastic File System (Amazon EFS) to store and share the documents.

Explanation/Reference:

Reference <https://aws.amazon.com/efs/enterprise-applications/>

**NO.65** 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.

**NO.66** A customer is running a critical payroll system in a production environment in one data center and a disaster recovery (DR) environment in another. The application includes load-balanced web servers and failover for the MySQL database. The customer's DR process is manual and error-prone. For this reason, management has asked IT to migrate the application to AWS and make it highly available so that IT no longer has to manually fail over the environment.

How should a Solutions Architect migrate the system to AWS?

- \* Migrate the production and DR environments to different Availability Zones within the same region. Let AWS manage failover between the environments.
- \* Migrate the production and DR environments to different regions. Let AWS manage failover between the environments.
- \* Migrate the production environment to a single Availability Zone, and set up instance recovery for Amazon EC2. Decommission the DR environment because it is no longer needed.
- \* Migrate the production environment to span multiple Availability Zones, using Elastic Load Balancing and Multi-AZ Amazon RDS. Decommission the DR environment because it is no longer needed.

**NO.67** You are building an application that stores relational data from users. Users across the globe will use this application. Your CTO is concerned about the scaling requirements because the size of the user base is unknown. You need to implement a database solution that can scale with your user growth with minimum configuration changes. Which storage solution should you use?

- \* Cloud SQL
- \* Cloud Spanner
- \* Cloud Firestore
- \* Cloud Datastore

CloudSpanner supports relational data model, globally. Option D is incorrect, Cloud Datastore is to store NonRelational data.

**NO.68** You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

- \* 1. Use nslookup to get the IP address for storage.googleapis.com.2. Negotiate with the security team to be able to give a public IP address to the servers.3. Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.
- \* 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).2. In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.
- \* 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.2. Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.3. Configure your new instances to use this ILB as proxy.
- \* 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network,

configure your DNS server to resolve \*.googleapis.com as a CNAME to restricted.googleapis.com.

**NO.69** You need to help a developer install the App Engine Go extensions. However, you've forgotten the exact name of the component. Which command could you run to show all of the available options?

- \* gcloud config list
- \* gcloud component list
- \* gcloud config components list
- \* gcloud components list

**NO.70** You have production and test workloads that you want to deploy on Compute Engine. Production VMs need to be in a different subnet than the test VMs. All the VMs must be able to reach each other over internal IP without creating additional routes. You need to set up VPC and the 2 subnets. Which configuration meets these requirements?

- \* Create a single custom VPC with 2 subnets.

Create each subnet in a different region and with a different CIDR range.

- \* Create a single custom VPC with 2 subnets.

Create each subnet in the same region and with the same CIDR range.

- \* Create 2 custom VPCs, each with a single subnet.

Create each subnet in a different region and with a different CIDR range.

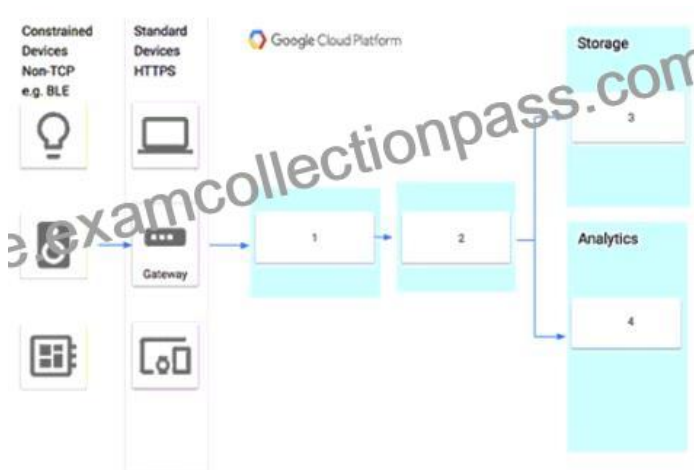
- \* Create 2 custom VPCs, each with a single subnet.

Create each subnet in the same region and with the same CIDR range.

Primary and secondary ranges for subnets cannot overlap with any allocated range, any primary or secondary range of another subnet in the same network, or any IP ranges of subnets in peered networks.

<https://cloud.google.com/vpc/docs/using-vpc#subnet-rules>

**NO.71** You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?



- \* Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
- \* Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery
- \* Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable

\* Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

**NO.72** You need to set up a policy so that videos stored in a specific Cloud Storage Regional bucket are moved to Coldline after 90 days, and then deleted after one year from their creation. How should you set up the policy?

- \* Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete actions. Set the SetStorageClass action to 90 days and the Delete action to 275 days (365 &#8211; 90)
- \* Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete actions. Set the SetStorageClass action to 90 days and the Delete action to 365 days.
- \* Use gsutil rewrite and set the Delete action to 275 days (365-90).
- \* Use gsutil rewrite and set the Delete action to 365 days.

<https://cloud.google.com/storage/docs/lifecycle#setstorageclass-cost>

# The object's time spent set at the original storage class counts towards any minimum storage duration that applies for the new storage class.

**NO.73** You have a Dockerfile that you need to deploy on Kubernetes Engine. What should you do?

- \* Use `kubectl app deploy <dockerfilename>`.
- \* Use `gcloud app deploy <dockerfilename>`.
- \* Create a docker image from the Dockerfile and upload it to Container Registry.

Create a Deployment YAML file to point to that image.

Use `kubectl` to create the deployment with that file.

- \* Create a docker image from the Dockerfile and upload it to Cloud Storage.

Create a Deployment YAML file to point to that image.

Use `kubectl` to create the deployment with that file.

**NO.74** You are using Container Registry to centrally store your company's container images in a separate project. In another project, you want to create a Google Kubernetes Engine (GKE) cluster. You want to ensure that Kubernetes can download images from Container Registry. What should you do?

- \* In the project where the images are stored, grant the Storage Object Viewer IAM role to the service account used by the Kubernetes nodes.
- \* When you create the GKE cluster, choose the Allow full access to all Cloud APIs option under

Access scopes.

- \* Create a service account, and give it access to Cloud Storage. Create a P12 key for this service account and use it as an `imagePullSecrets` in Kubernetes.
- \* Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account. If the cluster is in a different project or if the VMs in the cluster use a different service account, you must grant the service account the appropriate permissions to access the storage bucket used by Container Registry.

For the service account used by Compute Engine VMs, including VMs in Google Kubernetes Engine clusters, access is based on both Cloud IAM permissions and storage access scopes.

<https://cloud.google.com/container-registry/docs/access-control>

<https://cloud.google.com/container-registry/docs/using-with-google-cloud-platform>

**NO.75** You have a developer laptop with the Cloud SDK installed on Ubuntu. The Cloud SDK was installed from the Google Cloud Ubuntu package repository. You want to test your application locally on your laptop with Cloud Datastore. What should you do?

- \* Export Cloud Datastore data using `gcloud datastore export`.
- \* Create a Cloud Datastore index using `gcloud datastore indexes create`.
- \* Install the `google-cloud-sdk-datastore-emulator` component using the `apt get install` command.
- \* Install the `cloud-datastore-emulator` component using the `gcloud components install` command.

When you install SDK using `apt` Cloud SDK Component Manager is disabled and you need to install extra packages again using `apt`.

[https://cloud.google.com/sdk/docs/components#managing\\_cloud\\_sdk\\_components](https://cloud.google.com/sdk/docs/components#managing_cloud_sdk_components) Note: These instructions will not work if you have installed Cloud SDK using a package manager such as `APT` or `yum` because Cloud SDK Component Manager is disabled when using that method of installation.

**NO.76** You need to create an autoscaling managed instance group for an HTTPS web application. You want to make sure that unhealthy VMs are recreated. What should you do?

- \* Create a health check on port 443 and use that when creating the Managed Instance Group.
- \* Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.
- \* In the Instance Template, add the label `&#8216;health-check&#8217;`.
- \* In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Reference:

<https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances>

**NO.77** Your team has been working towards using desired state configuration for your entire infrastructure, which is why they're excited to store the Kubernetes Deployments in YAML. You created a Kubernetes Deployment with the `kubectl apply` command and passed on a YAML file. You need to edit the number of replicas. What steps should you take to update the Deployment?

- \* Edit the number of replicas in the YAML file and rerun the `kubectl apply`.
- \* Edit the YAML and push it to Github so that the git triggers deploy the change.
- \* Disregard the YAML file. Use the `kubectl scale` command.
- \* Edit the number of replicas in the YAML file and run the `kubectl set image` command.

**NO.78** You deployed a new application inside your Google Kubernetes Engine cluster using the YAML file specified below.

```
apiVersion: apps/v1          apiVersion: v1
kind: Deployment             kind: Service
metadata:                   metadata:
  name: myapp-deployment    name: myapp-service
spec:                       spec:
  selector:                 ports:
    matchLabels:            - port: 8000
      app: myapp            targetPort: 80
  replicas: 2               protocol: TCP
  template:                 selector:
    metadata:               app: myapp
      labels:
        app: myapp
    spec:
      containers:
        - name: myapp
          image: myapp:1.1
          ports:
            - containerPort: 80
```

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>

**NO.79** You need to create a copy of a custom Compute Engine virtual machine (VM) to facilitate an expected increase in application traffic due to a business acquisition. What should you do?

- \* Create a Compute Engine snapshot of your base VM. Create your images from that snapshot.
- \* Create a Compute Engine snapshot of your base VM. Create your instances from that snapshot.
- \* Create a custom Compute Engine image from a snapshot. Create your images from that image.
- \* Create a custom Compute Engine image from a snapshot. Create your instances from that image.

A custom image belongs only to your project. To create an instance with a custom image, you must first have a custom image.

Reference:

Preparing your instance for an image

You can create an image from a disk even while it is attached to a running VM instance. However, your image will be more reliable if you put the instance in a state that is easier for the image to capture. Use one of the following processes to prepare your boot disk for the image:

Stop the instance so that it can shut down and stop writing any data to the persistent disk.

If you can't stop your instance before you create the image, minimize the amount of writes to the disk and sync your file system.

Pause apps or operating system processes that write data to that persistent disk.

Run an app flush to disk if necessary. For example, MySQL has a FLUSH statement. Other apps might have similar processes.

Stop your apps from writing to your persistent disk.

Run `sudo sync`.

After you prepare the instance, create the image.

[https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images#prepare\\_instance\\_for\\_image](https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images#prepare_instance_for_image)

**NO.80** You're running an n-tier application on Compute Engine with an Apache web server serving up web requests. You want to consolidate all of your logging into Stackdriver. What's the best approach to get the Apache logs into Stackdriver?

- \* Create a log sink and export it to Stackdriver.
- \* Stackdriver logs application data from all instances by default.
- \* Enable Stackdriver monitoring when creating the instance.
- \* Install the Stackdriver monitoring and logging agents on the instance.

**NO.81** You are hosting an application from Compute Engine virtual machines (VMs) in us-central1-a.

You want to adjust your design to support the failure of a single Compute Engine zone, eliminate downtime, and minimize cost. What should you do?

- \* Create Compute Engine resources in us-central1-b.
- \* Balance the load across both us-central1-a and us-central1-b.
- \* Create a Managed Instance Group and specify us-central1-a as the zone.
- \* Configure the Health Check with a short Health Interval.
- \* Create an HTTP(S) Load Balancer.
- \* Create one or more global forwarding rules to direct traffic to your VMs.
- \* Perform regular backups of your application.

\* Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.

\* Restore from backups when notified.

<https://github.com/GoogleCloudPlatform/puppet-google-compute>

**NO.82** Your company has a large quantity of unstructured data in different file formats. You want to perform ETL transformations on the data. You need to make the data accessible on Google Cloud so it can be processed by a Dataflow job. What should you do?

- \* Upload the data to BigQuery using the bq command line tool.
- \* Upload the data to Cloud Storage using the gsutil command line tool.
- \* Upload the data into Cloud SQL using the import function in the console.
- \* Upload the data into Cloud Spanner using the import function in the console.

<https://cloud.google.com/solutions/performing-etl-from-relational-database-into-bigquery>

To become a Google Associate Cloud Engineer, candidates must pass the Associate-Cloud-Engineer exam, which is an online, proctored exam that can be taken from anywhere in the world. Associate-Cloud-Engineer exam fee is \$125, and candidates can register for the exam on the Google Cloud Platform website. Google Associate Cloud Engineer Exam certification is valid for two

years, after which candidates must renew their certification by passing the current version of the exam.

### Exam highlights

The Google Associate Cloud Engineer certification exam is 2 hours long and covers mostly two question formats, which are multiple choice and multiple select. The applicants can take the test as an online proctored option (remotely) or as an onsite proctored exam at one of the centers across the world. You can go through the official website for details of the test-taking process. To register for this exam, the individuals must pay the fee of \$125. This fee applies to a single try. If you do not pass your Google exam at the first attempt, you have to try again and pay another fee for registration. The test is available in the English, Spanish, Japanese, and Indonesian languages.

**Pass Your Google Cloud Certified Associate-Cloud-Engineer Exam on Sep 21, 2023 with 229 Questions:**

<https://www.examcollectionpass.com/Google/Associate-Cloud-Engineer-practice-exam-dumps.html>