

2024 Updated Verified Associate-Cloud-Engineer dumps Q&As - Pass Guarantee or Full Refund [Q152-Q173]



2024 Updated Verified Associate-Cloud-Engineer dumps Q&As - Pass Guarantee or Full Refund
Associate-Cloud-Engineer PDF Questions and Testing Engine With 268 Questions

Google Associate-Cloud-Engineer Exam is a certification exam offered by Google that tests an individual's knowledge and skills in cloud computing. Associate-Cloud-Engineer exam is designed to assess the examinee's ability to deploy and manage Google Cloud Platform (GCP) services, including computing, storage, networking, and security. Associate-Cloud-Engineer exam is intended for individuals who have experience in cloud computing and are looking to validate their skills and knowledge.

Google Associate-Cloud-Engineer certification exam is one of the most popular cloud certification exams in the market today. This is because GCP is quickly becoming one of the most widely used cloud platforms, and there is a growing demand for cloud engineers who are skilled in deploying and managing GCP projects. Individuals who earn this certification can expect to have an advantage in the job market and may be able to command higher salaries.

NEW QUESTION 152

You created an instance of SQL Server 2017 on Compute Engine to test features in the new version. You want to connect to this instance using the fewest number of steps. What should you do?

- * Install a RDP client on your desktop. Verify that a firewall rule for port 3389 exists.
- * Install a RDP client in your desktop. Set a Windows username and password in the GCP Console. Use the credentials to log in to the instance.
- * Set a Windows password in the GCP Console. Verify that a firewall rule for port 22 exists. Click the RDP button in the GCP Console and supply the credentials to log in.
- * Set a Windows username and password in the GCP Console. Verify that a firewall rule for port 3389 exists. Click the RDP button in the GCP Console, and supply the credentials to log in.

<https://cloud.google.com/compute/docs/instances/connecting-to-windows#remote-desktop-connection-app>

<https://cloud.google.com/compute/docs/instances/windows/generating-credentials>

<https://cloud.google.com/compute/docs/instances/connecting-to-windows#before-you-begin>

NEW QUESTION 153

Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

- * Multi-Regional Storage
- * Regional Storage
- * Nearline Storage
- * Coldline Storage

Explanation/Reference: <https://cloud.google.com/storage/docs/storage-classes#nearline>

NEW QUESTION 154

You are managing a project for the Business Intelligence (BI) department in your company. A data pipeline ingests data into BigQuery via streaming. You want the users in the BI department to be able to run the custom SQL queries against the latest data in BigQuery. What should you do?

- * Create a Data Studio dashboard that uses the related BigQuery tables as a source and give the BI team view access to the Data Studio dashboard.
- * Create a Service Account for the BI team and distribute a new private key to each member of the BI team.
- * Use Cloud Scheduler to schedule a batch Dataflow job to copy the data from BigQuery to the BI team's internal data warehouse.
- * Assign the IAM role of BigQuery User to a Google Group that contains the members of the BI team.

Explanation

When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A member with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/bigquery.dataOwner) on these new datasets.

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

NEW QUESTION 155

You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes.

You want to use Google services. What should you do?

- * 1. Create a consumer Gmail account.

- 2. Write a script that monitors the CPU usage.

- 3. When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.
 - * 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.

 - 2. Create an Alerting Policy in Stackdriver that uses the threshold as a trigger condition.

 - 3. Configure your email address in the notification channel.
 - * 1. Create a Stackdriver Workspace, and associate your GCP project with it.

 - 2. Write a script that monitors the CPU usage and sends it as a custom metric to Stackdriver.

 - 3. Create an uptime check for the instance in Stackdriver.
 - * 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: `[(0-9] {1,3})%`

 - 2. In Stackdriver Monitoring, create an Alerting Policy based on this metric.

 - 3. Configure your email address in the notification channel.

NEW QUESTION 156

You have successfully created a development environment in a project for an application. This application uses Compute Engine and Cloud SQL. Now, you need to create a production environment for this application.

The security team has forbidden the existence of network routes between these 2 environments, and asks you to follow Google-recommended practices. What should you do?

- * Create a new project, enable the Compute Engine and Cloud SQL APIs in that project, and replicate the setup you have created in the development environment.
- * Create a new production subnet in the existing VPC and a new production Cloud SQL instance in your existing project, and deploy your application using those resources.
- * Create a new project, modify your existing VPC to be a Shared VPC, share that VPC with your new project, and replicate the setup you have in the development environment in that new project, in the Shared VPC.
- * Ask the security team to grant you the Project Editor role in an existing production project used by another division of your company. Once they grant you that role, replicate the setup you have in the development environment in that project.

Explanation

This aligns with Googles recommended practices. By creating a new project, we achieve complete isolation between development and production environments; as well as isolate this production application from production applications of other departments.

Ref: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#define-hierarchy>

NEW QUESTION 157

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

NEW QUESTION 158

You have several users who need access to some very specific Google Cloud functionality. You'd like to follow the principle of least privilege. What's the best way to ensure these users can list Cloud Storage buckets, list BigQuery jobs, and list compute disks?

- * Add the users to the viewer role.
- * Use the Cloud Storage Bucket Viewer, BigQuery Job User, and Compute User predefined roles.
- * Create a custom role for this job role, add the required permissions, and add the users to the role.
- * Add the users to a group, apply the Cloud Storage Bucket Viewer, BigQuery Job User, and Compute User predefined roles.

NEW QUESTION 159

Your company has a single sign-on (SSO) identity provider that supports Security Assertion Markup Language (SAML) integration with service providers. Your company has users in Cloud Identity. You would like users to authenticate using your company's SSO provider.

What should you do?

- * In Cloud Identity, set up SSO with Google as an identity provider to access custom SAML apps.
- * In Cloud Identity, set up SSO with a third-party identity provider with Google as a service provider.
- * Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Mobile

& Desktop Apps.

- * Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Web Server Applications.

https://support.google.com/cloudidentity/answer/6262987?hl=en&ref_topic=7558767 Google offers a SAML-based single sign-on (SSO) service that provides partner companies with full control over the authorization and authentication of hosted user accounts that can access web-based applications like Gmail or Google Calendar. Using the SAML model, Google acts as the service provider and provides services such as Gmail and Start Pages.

NEW QUESTION 160

You have a batch workload that runs every night and uses a large number of virtual machines (VMs). It is fault-tolerant and can tolerate some of the VMs being terminated. The current cost of VMs is too high. What should you do?

- * Run a test using simulated maintenance events. If the test is successful, use preemptible N1 Standard VMs when running future jobs.
- * Run a test using simulated maintenance events. If the test is successful, use N1 Standard VMs when running future jobs.
- * Run a test using a managed instance group. If the test is successful, use N1 Standard VMs in the managed instance group when running future jobs.
- * Run a test using N1 standard VMs instead of N2. If the test is successful, use N1 Standard VMs when running future jobs.

Creating and starting a preemptible VM instance This page explains how to create and use a preemptible virtual machine (VM) instance. A preemptible instance is an instance you can create and run at a much lower price than normal instances. However, Compute Engine might terminate (preempt) these instances if it requires access to those resources for other tasks. Preemptible instances will always terminate after 24 hours.

To learn more about preemptible instances, read the preemptible instances documentation. Preemptible instances are recommended only for fault-tolerant applications that can withstand instance preemptions. Make sure your application can handle preemptions before you decide to create a preemptible instance. To understand the risks and value of preemptible instances, read the preemptible instances documentation.<https://cloud.google.com/compute/docs/instances/create-start-preemptible-instance>

NEW QUESTION 161

Your company requires all developers to have the same permissions, regardless of the Google Cloud project they are working on. Your company's security policy also restricts developer permissions to Compute Engine.

Cloud Functions, and Cloud SQL. You want to implement the security policy with minimal effort. What should you do?

- * * Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions in one project within the Google Cloud organization.
 - * Copy the role across all projects created within the organization with the `gcloud iam roles copy` command.
 - * Assign the role to developers in those projects.
 - * * Add all developers to a Google group in Google Groups for Workspace.
 - * Assign the predefined role of Compute Admin to the Google group at the Google Cloud organization level.
 - * * Add all developers to a Google group in Cloud Identity.
 - * Assign predefined roles for Compute Engine, Cloud Functions, and Cloud SQL permissions to the Google group for each project in the Google Cloud organization.
 - * * Add all developers to a Google group in Cloud Identity.
 - * Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions at the Google Cloud organization level.
 - * Assign the custom role to the Google group.
- <https://www.cloudskillsboost.google/focuses/1035?parent=catalog#:~:text=custom%20role%20at%20the%20org>

NEW QUESTION 162

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.

For unstructured data use cloud storage. Use Big Query for analytics, data warehouse with structured data.

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

NEW QUESTION 163

Your customer has implemented a solution that uses Cloud Spanner and notices some read latency-related performance issues on one table. This table is accessed only by their users using a primary key. The table schema is shown below.

```
CREATE TABLE Persons (  
    person_id INT64 NOT NULL, // sequential number based on number of registration  
    account_creation_date DATE, // system date  
    birthdate DATE, // customer birthdate  
    firstname STRING (255), // first name  
    lastname STRING (255), // last name  
    profile_picture BYTES (255) // profile picture  
    ) PRIMARY KEY (person_id)
```

You want to resolve the issue. What should you do?

- A. Remove the profile_picture field from the table.
- B. Add a secondary index on the person_id column.
- C. Change the primary key to not have monotonically increasing values.
- D. Create a secondary index using the following Data Definition Language (DDL):

```
CREATE INDEX person_id_ix  
ON Persons (  
    person_id,  
    firstname,  
    lastname  
    ) STORING (  
    profile_picture  
    )
```

- * Option A
- * Option B
- * Option C
- * Option D

NEW QUESTION 164

Your company is moving from an on-premises environment to Google Cloud Platform (GCP). You have multiple development teams that use Cassandra environments as backend databases. They all need a development environment that is isolated from other Cassandra instances. You want to move to GCP quickly and with minimal support effort. What should you do?

- * 1. Build an instruction guide to install Cassandra on GCP.
- 2. Make the instruction guide accessible to your developers.
 - * 1. Advise your developers to go to Cloud Marketplace.
- 2. Ask the developers to launch a Cassandra image for their development work.
 - * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.
- 2. Use the snapshot to create instances for your developers.
 - * 1. Build a Cassandra Compute Engine instance and take a snapshot of it.
- 2. Upload the snapshot to Cloud Storage and make it accessible to your developers.
- 3. Build instructions to create a Compute Engine instance from the snapshot so that developers can do it themselves.

Explanation:

<https://medium.com/google-cloud/how-to-deploy-cassandra-and-connect-on-google-cloud-platform-with-a-few-clicks-11ee3d7001d1>

<https://cloud.google.com/blog/products/databases/open-source-cassandra-now-managed-on-google-cloud>

<https://cloud.google.com/marketplace>

You can deploy Cassandra as a Service, called Astra, on the Google Cloud Marketplace. Not only do you get a unified bill for all GCP services, your Developers can now create Cassandra clusters on Google Cloud in minutes and build applications with Cassandra as a database as a service without the operational overhead of managing Cassandra

NEW QUESTION 165

You want to configure 10 Compute Engine instances for availability when maintenance occurs. Your requirements state that these instances should attempt to automatically restart if they crash. Also, the instances should be highly available including during system maintenance. What should you do?

- * Create an instance template for the instances. Set the `Automatic Restart` to on. Set the `On-host maintenance` to Migrate VM instance. Add the instance template to an instance group.
- * Create an instance template for the instances. Set `Automatic Restart` to off. Set `On-host maintenance` to Terminate VM instances. Add the instance template to an instance group.
- * Create an instance group for the instances. Set the `Autohealing` health check to healthy (HTTP).
- * Create an instance group for the instance. Verify that the `Advanced creation options` setting for `do not retry machine creation` is set to off.

Create an instance template for the instances so VMs have same specs. Set the `~Automatic Restart` to on to VM automatically restarts upon crash. Set the `~On-host maintenance` to Migrate VM instance. This will take care of VM during maintenance window. It will migrate VM instance making it highly available Add the instance template to an instance group so instances can be managed.

- * `onHostMaintenance`: Determines the behavior when a maintenance event occurs that might cause your instance to reboot.
- * [Default] `MIGRATE`, which causes Compute Engine to live migrate an instance when there is a maintenance event.
- * `TERMINATE`, which stops an instance instead of migrating it.
- * `automaticRestart`: Determines the behavior when an instance crashes or is stopped by the system.
- * [Default] `true`, so Compute Engine restarts an instance if the instance crashes or is stopped.
- * `false`, so Compute Engine does not restart an instance if the instance crashes or is stopped.

Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash.

And Enabling Migrate VM Instance enables live migrates i.e. compute instances are migrated during system maintenance and remain running during the migration.

Automatic Restart If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the `automaticRestart` field to true. This setting does not apply if the instance is taken offline through a user action, such as calling `sudo shutdown`, or during a

zone outage. Ref: <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart> Enabling the Migrate VM Instance option migrates your instance away from an infrastructure maintenance event, and your instance remains running during the migration. Your instance might experience a short period of decreased performance, although generally, most instances should not notice any difference. This is ideal for instances that require constant uptime and can tolerate a short period of decreased performance. Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#live_mi

NEW QUESTION 166

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 list` to get the project ID, and then run `gcloud services list --project <project ID>`.
 - * Run `gcloud init` to set the current project to my-project, and then run `gcloud services list --available`.
 - * Run `gcloud info` to 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`.
- `gcloud services list --available` returns not only the enabled services in the project but also services that CAN be enabled.`

<https://cloud.google.com/sdk/gcloud/reference/services/list--available> Run the following command to list the enabled APIs and services in your current project:

```
gcloud services list
```

whereas, Run the following command to list the APIs and services available to you in your current project:

```
gcloud services list --available
```

<https://cloud.google.com/sdk/gcloud/reference/services/list--available>

```
--available
```

Return the services available to the project to enable. This list will include any services that the project has already enabled.

To list the services the current project has enabled for consumption, run:

```
gcloud services list --enabled
```

To list the services the current project can enable for consumption, run:

```
gcloud services list --available
```

NEW QUESTION 167

You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their application on Google Cloud. They manage the resources in their own project, but they need access to the BigQuery dataset in your project. You want to provide the partner company with access to the dataset. What should you do?

- * Create a Service Account in your own project, and grant this Service Account access to BigQuery in your project.
- * Create a Service Account in your own project, and ask the partner to grant this Service Account access to BigQuery in their project.
- * Ask the partner to create a Service Account in their project, and have them give the Service Account access to BigQuery in their project.

* Ask the partner to create a Service Account in their project, and grant their Service Account access to the BigQuery dataset in your project.

NEW QUESTION 168

You have a virtual machine that is currently configured with 2 vCPUs and 4 GB of memory. It is running out of memory. You want to upgrade the virtual machine to have 8 GB of memory.

What should you do?

- * Rely on live migration to move the workload to a machine with more memory.
- * Use `gcloud` to add metadata to the VM. Set the key to `required-memory-size` and the value to 8 GB.
- * Stop the VM, change the machine type to `n1-standard-8`, and start the VM.
- * Stop the VM, increase the memory to 8 GB, and start the VM.

NEW QUESTION 169

Several employees at your company have been creating projects with Cloud Platform and paying for it with their personal credit cards, which the company reimburses. The company wants to centralize all these projects under a single, new billing account. What should you do?

- * Contact `cloud-billing@google.com` with your bank account details and request a corporate billing account for your company.
- * Create a ticket with Google Support and wait for their call to share your credit card details over the phone.
- * In the Google Platform Console, go to the Resource Manager and move all projects to the root Organization.
- * In the Google Cloud Platform Console, create a new billing account and set up a payment method.

NEW QUESTION 170

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.

Reference:

<https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances>

NEW QUESTION 171

You are building a new version of an application hosted in an App Engine environment. You want to test the new version with 1% of users before you completely switch your application over to the new version.

What should you do?

- * Deploy a new version of your application in Google Kubernetes Engine instead of App Engine and then use GCP Console to split traffic.
- * Deploy a new version of your application in a Compute Engine instance instead of App Engine and then use GCP Console to split traffic.

- * Deploy a new version as a separate app in App Engine. Then configure App Engine using GCP Console to split traffic between the two apps.
 - * Deploy a new version of your application in App Engine. Then go to App Engine settings in GCP Console and split traffic between the current version and newly deployed versions accordingly.
- <https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

NEW QUESTION 172

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-58d4bbb995-lp86m  0/1    Pending  0         9m
myapp-deployment-58d4bbb995-qjpkq  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-58d4bbb995-lp86m Pod and check for warning messages.
- * View logs of the container in myapp-deployment-58d4bbb995-lp86m pod and check for warning messages.

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-application/#debugging-pods> You can't view logs of a pod that isn't deployed, so D is incorrect.

C allows you to check the pod deployment messages and look for errors.

NEW QUESTION 173

You need to run an important query in BigQuery but expect it to return a lot of records. You want to find out how much it will cost to run the query. You are using on-demand pricing. What should you do?

- * Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.
- * Use the command line to run a dry run query to estimate the number of bytes read.

Then convert that bytes estimate to dollars using the Pricing Calculator.

- * Use the command line to run a dry run query to estimate the number of bytes returned.

Then convert that bytes estimate to dollars using the Pricing Calculator.

- * Run a select count (*) to get an idea of how many records your query will look through.

Then convert that number of rows to dollars using the Pricing Calculator.

On-demand pricing

Under on-demand pricing, BigQuery charges for queries by using one metric: the number of bytes processed (also referred to as bytes read). You are charged for the number of bytes processed whether the data is stored in BigQuery or in an external data source such as Cloud Storage, Drive, or Cloud Bigtable. On-demand pricing is based solely on usage.

https://cloud.google.com/bigquery/pricing#on_demand_pricing

Exam Engine for Associate-Cloud-Engineer Exam Free Demo & 365 Day Updates:

<https://www.vceprep.com/Associate-Cloud-Engineer-latest-vce-prep.html>