

[Dec 16, 2022 Google Associate-Cloud-Engineer Real Exam Questions and Answers FREE [Q75-Q95]



[Dec 16, 2022] Google Associate-Cloud-Engineer Real Exam Questions and Answers FREE
Pass Google Associate-Cloud-Engineer Exam Info and Free Practice Test

Planning & Configuring Cloud Solutions

The focus of the second topic is on the competency of the individuals in planning and evaluating the GCP product use with the help of Pricing Calculator. The test takers should also have the skills in planning & configuring compute resources, network resources, as well as data storage options.

Q75. You need to set a budget alert for use of Compute Engine services on one of the three Google Cloud Platform projects that you manage. All three projects are linked to a single billing account.

What should you do?

- * Verify that you are the project billing administrator. Select the associated billing account and create a budget and alert for the appropriate project.
- * Verify that you are the project billing administrator. Select the associated billing account and create a budget and a custom alert.
- * Verify that you are the project administrator. Select the associated billing account and create a budget for the appropriate project.
- * Verify that you are project administrator. Select the associated billing account and create a budget and a custom alert.

You can rely on default alert. No need for custom alert.

Q76. Your company publishes large files on an Apache web server that runs on a Compute Engine instance. The Apache web server is not the only application running in the project. You want to receive an email when the egress network costs for the server exceed 100 dollars for the current month as measured by Google Cloud Platform (GCP). What should you do?

- * Set up a budget alert on the project with an amount of 100 dollars, a threshold of 100%, and notification type of email;
- * Set up a budget alert on the billing account with an amount of 100 dollars, a threshold of 100%, and notification type of email;
- * Export the billing data to BigQuery.

Create a Cloud Function that uses BigQuery to sum the egress network costs of the exported billing data for the Apache web server for the current month and sends an email if it is over 100 dollars.

Schedule the Cloud Function using Cloud Scheduler to run hourly.

- * Use the Stackdriver Logging Agent to export the Apache web server logs to Stackdriver Logging.

Create a Cloud Function that uses BigQuery to parse the HTTP response log data in Stackdriver for the current month and sends an email if the size of all HTTP responses, multiplied by current GCP egress prices, totals over 100 dollars. Schedule the Cloud Function using Cloud Scheduler to run hourly.

The Apache web server is not the only application running in the project, so the budget alert is not suitable because it is related to all applications and not only the network traffic as needed.

Q77. The core business of your company is to rent out construction equipment at a large scale. All the equipment that is being rented out has been equipped with multiple sensors that send event information every few seconds. These signals can vary from engine status, distance traveled, fuel level, and more. Customers are billed based on the consumption monitored by these sensors. You expect high throughput; up to thousands of events per hour per device; and need to retrieve consistent data based on the time of the event. Storing and retrieving individual signals should be atomic. What should you do?

- * Create a file in Cloud Storage per device and append new data to that file.
- * Create a file in Cloud Filestore per device and append new data to that file.
- * Ingest the data into Datastore. Store data in an entity group based on the device.
- * Ingest the data into Cloud Bigtable. Create a row key based on the event timestamp.

Q78. 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

Q79. You have a web application deployed as a managed instance group. You have a new version of the application to gradually deploy. Your web application is currently receiving live web traffic.

You want to ensure that the available capacity does not decrease during the deployment. What should you do?

- * Perform a rolling-action start-update with maxSurge set to 0 and maxUnavailable set to 1.
- * Perform a rolling-action start-update with maxSurge set to 1 and maxUnavailable set to 0.
- * Create a new managed instance group with an updated instance template.

Add the group to the backend service for the load balancer.

When all instances in the new managed instance group are healthy, delete the old managed instance group.

* Create a new instance template with the new application version.

Update the existing managed instance group with the new instance template.

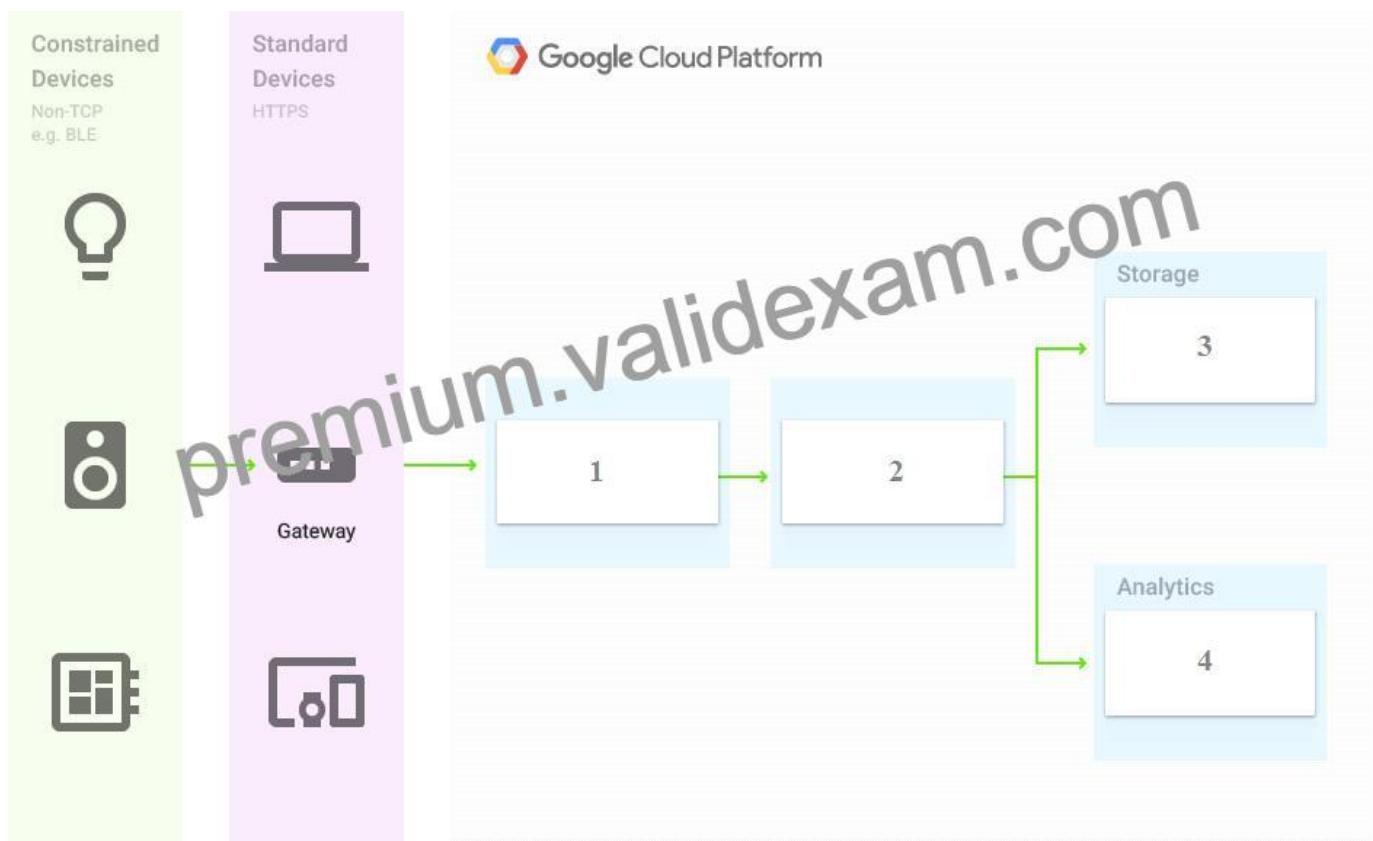
Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

We need to ensure the global capacity remains intact, for that reason we need to establish maxUnavailable to 0. On the other hand, we need to ensure new instances can be created. We do that by establishing the maxSurge to 1.

Option C is more expensive and more difficult to set up and option D won't meet requirements since it won't keep global capacity intact.

<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#options>

Q80. 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

Reference:

<https://cloud.google.com/solutions/correlating-time-series-dataflow>

Q81. 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>`

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

Q82. 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.

Q83. You have been asked to create robust Virtual Private Network (VPN) connectivity between a new Virtual Private Cloud (VPC) and a remote site. Key requirements include dynamic routing, a shared address space of 10.19.0.1/22, and no overprovisioning of tunnels during a failover event. You want to follow Google-recommended practices to set up a high availability Cloud VPN. What should you do?

- * Use a custom mode VPC network, configure static routes, and use active/passive routing
- * Use an automatic mode VPC network, configure static routes, and use active/active routing
- * Use a custom mode VPC network use Cloud Router border gateway protocol (86P) routes, and use active/passive routing
- * Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Q84. You are asked to set up application performance monitoring on Google Cloud projects A, B, and C as a single pane of glass. You want to monitor CPU, memory, and disk. What should you do?

- * Enable API and then share charts from project A, B, and C.
- * Enable API and then give the metrics.reader role to projects A, B, and C.
- * Enable API and then use default dashboards to view all projects in sequence.
- * Enable API, create a workspace under project A, and then add project B and C.

Q85. 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.

Reference:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-app>

Q86. You deployed an App Engine application using `gcloud app deploy`, but it did not deploy to the intended project. You want to find out why this happened and where the application deployed. What should you do?

- * Check the `app.yaml` file for your application and check project settings.

- * Check the web-application.xml file for your application and check project settings.
- * Go to Deployment Manager and review settings for deployment of applications.
- * Go to the Cloud Shell and run `gcloud config list` to review the Google Cloud configuration used for deployment.

Q87. You want to deploy an application on Cloud Run that processes messages from a Cloud Pub/Sub topic. You want to follow Google-recommended practices. What should you do?

- * 1. Create a Cloud Function that uses a Cloud Pub/Sub trigger on that topic. 2. Call your application on Cloud Run from the Cloud Function for every message.
- * 1. Grant the Pub/Sub Subscriber role to the service account used by Cloud Run. 2. Create a Cloud Pub/Sub subscription for that topic. 3. Make your application pull messages from that subscription.
- * 1. Create a service account. 2. Give the Cloud Run Invoker role to that service account for your Cloud Run application. 3. Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint.
- * 1. Deploy your application on Cloud Run on GKE with the connectivity set to Internal. 2. Create a Cloud Pub/Sub subscription for that topic. 3. In the same Google Kubernetes Engine cluster as your application, deploy a container that takes the messages and sends them to your application.

Q88. 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.

Q89. You have one GCP account running in your default region and zone and another account running in a non-default region and zone. You want to start a new Compute Engine instance in these two Google Cloud Platform accounts using the command line interface. What should you do?

- * Create two configurations using `gcloud config configurations create [NAME]`. Run `gcloud config configurations activate [NAME]` to switch between accounts when running the commands to start the Compute Engine instances.
- * Create two configurations using `gcloud config configurations create [NAME]`. Run `gcloud config configurations list` to start the Compute Engine instances.
- * Activate two configurations using `gcloud config configurations activate [NAME]`. Run `gcloud config list` to start the Compute Engine instances.
- * Activate two configurations using `gcloud config configurations activate [NAME]`. Run `gcloud config configurations list` to start the Compute Engine instances.

Q90. 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.

Reference:

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

Q91. You and your team have been working on a new application over the past couple weeks. While it's still in development, it's becoming a bit costly for your limited budget. The entire team had a meeting on Friday to talk about how to save money until you're able to launch. One of your team members suggested shutting down some services overnight and during the weekend. Though no official decision was made before leaving for the weekend, a junior team member sent out an email saying he found a solution to the problem.

When arriving at the office on Monday, you find that your project is no longer in the drop-down inside the Console. What's the most likely reason for the missing project?

- * The project was moved to another account.
- * Your trial credits expired, and the project was removed.
- * The engineer removed the project and attached all of the resources directly to the Organization.
- * The engineer clicked the "shut down" link thinking it was like a light switch that he could toggle on and off.

Q92. 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.

Q93. Your company uses BigQuery for data warehousing. Over time, many different business units in your company have created 1000+ datasets across hundreds of projects. Your CIO wants you to examine all datasets to find tables that contain an employee_ssn column. You want to minimize effort in performing this task. What should you do?

- * Go to Data Catalog and search for employee_ssn in the search box.
- * Write a shell script that uses the bq command line tool to loop through all the projects in your organization.
- * Write a script that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find the employee_ssn column.
- * Write a Cloud Dataflow job that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find employee_ssn column.

Data Catalog allows you to discover, manage, and understand data assets across Google Cloud Platform. Data Catalog API natively indexes Cloud BigQuery, Cloud Storage, and Cloud Pub/Sub data assets.

The Data Catalog API can be used to:

Search for data assets across different projects and GCP resources

Create tags and tag templates to add structured business metadata unique to your organization Automate the tagging process of data assets as part of a data ingestion pipeline

<https://cloud.google.com/data-catalog/docs/how-to/search>

Q94. 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.

Q95. You need to create a new billing account and then link it with an existing Google Cloud Platform project. What should you do?

- * Verify that you are Project Billing Manager for the GCP project. Update the existing project to link it to the existing billing account.
- * Verify that you are Project Billing Manager for the GCP project. Create a new billing account and link the new billing account to the existing project.
- * Verify that you are Billing Administrator for the billing account. Create a new project and link the new project to the existing billing account.
- * Verify that you are Billing Administrator for the billing account. Update the existing project to link it to the existing billing account.

Latest Associate-Cloud-Engineer Exam Dumps Google Exam:

<https://www.validexam.com/Associate-Cloud-Engineer-latest-dumps.html>]