

Associate Cloud Engineer

Passing Score: 800
Time Limit: 120 min
File Version: 1

Associate Cloud Engineer



<https://www.gratisexam.com/>

<https://www.gratisexam.com/>

Exam A

QUESTION 1

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?



<https://www.gratisexam.com/>

- A. 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.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys>

QUESTION 2

You are using multiple configurations for gcloud. You want to review the configured Kubernetes Engine cluster of an inactive configuration using the fewest possible steps. What should you do?

- A. Use `gcloud config configurations describe` to review the output.
- B. Use `gcloud config configurations activate` and `gcloud config list` to review the output.
- C. Use `kubectl config get-contexts` to review the output.
- D. Use `kubectl config use-context` and `kubectl config view` to review the output.

<https://www.gratisexam.com/>

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://medium.com/google-cloud/kubernetes-engine-kubectl-config-b6270d2b656c>

QUESTION 3

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?

- A. Multi-Regional Storage
- B. Regional Storage
- C. Nearline Storage
- D. Coldline Storage

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

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

QUESTION 4

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?

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

Correct Answer: D
Section: (none)
Explanation

Explanation/Reference:

Reference: <https://www.whizlabs.com/blog/google-cloud-interview-questions/>

QUESTION 5

You have an application that looks for its licensing server on the IP 10.0.3.21. You need to deploy the licensing server on Compute Engine. You do not want to change the configuration of the application and want the application to be able to reach the licensing server. What should you do?

- A. Reserve the IP 10.0.3.21 as a static internal IP address using `gcloud` and assign it to the licensing server.
- B. Reserve the IP 10.0.3.21 as a static public IP address using `gcloud` and assign it to the licensing server.
- C. Use the IP 10.0.3.21 as a custom ephemeral IP address and assign it to the licensing server.
- D. Start the licensing server with an automatic ephemeral IP address, and then promote it to a static internal IP address.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 6

You have a Dockerfile that you need to deploy on Kubernetes Engine. What should you do?

- A. Use `kubectl app deploy <dockerfilename>`.
- B. Use `gcloud app deploy <dockerfilename>`.
- C. 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.
- D. 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 7

Your development team needs a new Jenkins server for their project. You need to deploy the server using the fewest steps possible. What should you do?

- A. Download and deploy the Jenkins Java WAR to App Engine Standard.
- B. Create a new Compute Engine instance and install Jenkins through the command line interface.
- C. Create a Kubernetes cluster on Compute Engine and create a deployment with the Jenkins Docker image.
- D. Use GCP Marketplace to launch the Jenkins solution.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine>

QUESTION 8

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?

- A. Create a Google Kubernetes Engine cluster, and use horizontal pod autoscaling to scale the application.
- B. Create an instance template, and use the template in a managed instance group with autoscaling configured.
- C. Create an instance template, and use the template in a managed instance group that scales up and down based on the time of day.
- D. Use a set of third-party tools to build automation around scaling the application up and down, based on Stackdriver CPU usage monitoring.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 9

You are analyzing Google Cloud Platform service costs from three separate projects. You want to use this information to create service cost estimates by service type, daily and monthly, for the next six months using standard query syntax. What should you do?

- A. Export your bill to a Cloud Storage bucket, and then import into Cloud Bigtable for analysis.
- B. Export your bill to a Cloud Storage bucket, and then import into Google Sheets for analysis.
- C. Export your transactions to a local file, and perform analysis with a desktop tool.
- D. Export your bill to a BigQuery dataset, and then write time window-based SQL queries for analysis.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 10

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?

- A. Install a RDP client on your desktop. Verify that a firewall rule for port 3389 exists.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://medium.com/falafel-software/sql-server-in-the-google-cloud-a17e8a1f11ce>

QUESTION 11

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?

- A. 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.
- B. Create two configurations using `gcloud config configurations create [NAME]`. Run `gcloud configurations list` to start the Compute Engine instances.
- C. Activate two configurations using `gcloud configurations activate [NAME]`. Run `gcloud config list` to start the Compute Engine instances.
- D. Activate two configurations using `gcloud configurations activate [NAME]`. Run `gcloud configurations list` to start the Compute Engine instances.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 12

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?

- A. Use granular logging statements within a Deployment Manager template authored in Python.
- B. Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.
- C. Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.
- D. Execute the Deployment Manager template using the `--preview` option in the same project, and observe the state of interdependent resources.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 13

You need to configure IAM access audit logging in BigQuery for external auditors. You want to follow Google-recommended practices. What should you do?

- A. Add the auditors group to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
- B. Add the auditors group to two new custom IAM roles.
- C. Add the auditor user accounts to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
- D. Add the auditor user accounts to two new custom IAM roles.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/iam/docs/roles-audit-logging>

QUESTION 14

You need to set up permissions for a set of Compute Engine instances to enable them to write data into a particular Cloud Storage bucket. You want to follow Google-recommended practices. What should you do?

- A. Create a service account with an access scope. Use the access scope 'https://www.googleapis.com/auth/devstorage.write_only'.
- B. Create a service account with an access scope. Use the access scope 'https://www.googleapis.com/auth/cloud-platform'.
- C. Create a service account and add it to the IAM role 'storage.objectCreator' for that bucket.
- D. Create a service account and add it to the IAM role 'storage.objectAdmin' for that bucket.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 15

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?

- A. Create a signed URL with a four-hour expiration and share the URL with the company.
- B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.
- C. 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.
- D. 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 16

You are creating a Google Kubernetes Engine (GKE) cluster with a cluster autoscaler feature enabled. You need to make sure that each node of the cluster will run a monitoring pod that sends container metrics to a third-party monitoring solution. What should you do?

- A. Deploy the monitoring pod in a StatefulSet object.
- B. Deploy the monitoring pod in a DaemonSet object.
- C. Reference the monitoring pod in a Deployment object.
- D. Reference the monitoring pod in a cluster initializer at the GKE cluster creation time.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 17

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?

- A. Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
- B. 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.
- C. Configure a single Stackdriver account, and link all projects to the same account.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 18

You are deploying an application to a Compute Engine VM in a managed instance group. The application must be running at all times, but only a single instance of the VM should run per GCP project. How should you configure the instance group?

- A. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- B. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- C. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 2.
- D. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 19

You want to verify the IAM users and roles assigned within a GCP project named `my-project`. What should you do?

- A. Run `gcloud iam roles list`. Review the output section.
- B. Run `gcloud iam service-accounts list`. Review the output section.
- C. Navigate to the project and then to the IAM section in the GCP Console. Review the members and roles.

D. Navigate to the project and then to the Roles section in the GCP Console. Review the roles and status.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 20

You need to grant access for three users so that they can view and edit table data on a Cloud Spanner instance. What should you do?

- A. Run `gcloud iam roles describe roles/spanner.databaseUser`. Add the users to the role.
- B. Run `gcloud iam roles describe roles/spanner.databaseUser`. Add the users to a new group. Add the group to the role.
- C. Run `gcloud iam roles describe roles/spanner.viewer - --project my-project`. Add the users to the role.
- D. Run `gcloud iam roles describe roles/spanner.viewer - --project my-project`. Add the users to a new group. Add the group to the role.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 21

You create a new Google Kubernetes Engine (GKE) cluster and want to make sure that it always runs a supported and stable version of Kubernetes. What should you do?

- A. Enable the Node Auto-Repair feature for your GKE cluster.
- B. Enable the Node Auto-Upgrades feature for your GKE cluster.
- C. Select the latest available cluster version for your GKE cluster.
- D. Select "Container-Optimized OS (cos)" as a node image for your GKE cluster.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 22

You have an instance group that you want to load balance. You want the load balancer to terminate the client SSL session. The instance group is used to serve a public web application over HTTPS. You want to follow Google-recommended practices. What should you do?

- A. Configure an HTTP(S) load balancer.
- B. Configure an internal TCP load balancer.
- C. Configure an external SSL proxy load balancer.
- D. Configure an external TCP proxy load balancer.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/load-balancing/docs/https/>

QUESTION 23

You've deployed a microservice called `myapp1` to a Google Kubernetes Engine cluster using the YAML file specified below:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp1-deployment
spec:
  selector:
    matchLabels:
      app: myapp1
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp1
    spec:
      containers:
        - name: main-container
          image: gcr.io/my-company-repo/myapp1:1.4
          env:
            - name: DB_PASSWORD
              value: "t0ugh2guess!"
          ports:
            - containerPort: 8080

```

You need to refactor this configuration so that the database password is not stored in plain text. You want to follow Google-recommended practices. What should you do?

- A. Store the database password inside the Docker image of the container, not in the YAML file.
- B. Store the database password inside a Secret object. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
- C. Store the database password inside a ConfigMap object. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
- D. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 24

You recently deployed a new version of an application to App Engine and then discovered a bug in the release. You need to immediately revert to the prior version of the application. What should you do?

- A. Run `gcloud app restore`.
- B. On the App Engine page of the GCP Console, select the application that needs to be reverted and click `Revert`.
- C. On the App Engine Versions page of the GCP Console, route 100% of the traffic to the previous version.
- D. Deploy the original version as a separate application. Then go to App Engine settings and split traffic between applications so that the original version serves 100% of the requests.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://medium.com/google-cloud/app-engine-project-cleanup-9647296e796a>

QUESTION 25

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?

- A. 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.
- B. 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.
- C. Create an instance group for the instances. Set the 'Autohealing' health check to `healthy (HTTP)`.
- D. Create an instance group for the instance. Verify that the 'Advanced creation options' setting for 'do not retry machine creation' is set to off.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 26

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?

- A. Enable Cloud CDN on the website frontend.
- B. Enable 'Share publicly' on the PDF file objects.

- C. Set Content-Type metadata to application/pdf on the PDF file objects.
- D. Add a label to the storage bucket with a key of Content-Type and value of application/pdf.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 27

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?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 28

Your company has a 3-tier solution running on Compute Engine. The configuration of the current infrastructure is shown below.



Google Cloud Platform

VPC

Subnet Tier#1 10.0.1.0/24



Instance Tier 1
Compute Engine

Subnet Tier#2 10.0.2.0/24



Instance Tier 2
Compute Engine

Subnet Tier#3 10.0.3.0/24



Instance Tier 3
Compute Engine

Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:

- Instances in tier #1 must communicate with tier #2.
- Instances in tier #2 must communicate with tier #3.

What should you do?

A. 1. Create an ingress firewall rule with the following settings:

- Targets: all instances
- Source filter: IP ranges (with the range set to 10.0.2.0/24)
- Protocols: allow all

2. Create an ingress firewall rule with the following settings:

- Targets: all instances
- Source filter: IP ranges (with the range set to 10.0.1.0/24)
- Protocols: allow all

B. 1. Create an ingress firewall rule with the following settings:

- Targets: all instances with tier #2 service account
- Source filter: all instances with tier #1 service account
- Protocols: allow TCP:8080

2. Create an ingress firewall rule with the following settings:

- Targets: all instances with tier #3 service account
- Source filter: all instances with tier #2 service account
- Protocols: allow TCP: 8080

C. 1. Create an ingress firewall rule with the following settings:

- Targets: all instances with tier #2 service account
- Source filter: all instances with tier #1 service account
- Protocols: allow all

2. Create an ingress firewall rule with the following settings:

- Targets: all instances with tier #3 service account
- Source filter: all instances with tier #2 service account
- Protocols: allow all

D. 1. Create an egress firewall rule with the following settings:

- Targets: all instances
- Source filter: IP ranges (with the range set to 10.0.2.0/24)
- Protocols: allow TCP: 8080

2. Create an egress firewall rule with the following settings:

- Targets: all instances
- Source filter: IP ranges (with the range set to 10.0.1.0/24)
- Protocols: allow TCP: 8080

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 29

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?

- A. Deploy the new version in the same application and use the `--migrate` option.
- B. Deploy the new version in the same application and use the `--split` option to give a weight of 99 to the current version and a weight of 1 to the new version.
- C. 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.
- D. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 30

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?

- A. Perform a `rolling-action start-update` with `maxSurge` set to 0 and `maxUnavailable` set to 1.
- B. Perform a `rolling-action start-update` with `maxSurge` set to 1 and `maxUnavailable` set to 0.
- C. 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.
- D. 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 31

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?

- A. Cloud SQL
- B. Cloud Spanner
- C. Cloud Firestore
- D. Cloud Datastore

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 32

You are the organization and billing administrator for your company. The engineering team has the Project Creator role on the organization. You do not want the engineering team to be able to link projects to the billing account. Only the finance team should be able to link a project to a billing account, but they should not be able to make any other changes to projects. What should you do?

- A. Assign the finance team only the Billing Account User role on the billing account.
- B. Assign the engineering team only the Billing Account User role on the billing account.
- C. Assign the finance team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.
- D. Assign the engineering team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 33

You have an application running in Google Kubernetes Engine (GKE) with cluster autoscaling enabled. The application exposes a TCP endpoint. There are several

replicas of this application. You have a Compute Engine instance in the same region, but in another Virtual Private Cloud (VPC), called `gce-network`, that has no overlapping IP ranges with the first VPC. This instance needs to connect to the application on GKE. You want to minimize effort. What should you do?

- A. 1. In GKE, create a `Service` of type `LoadBalancer` that uses the application's Pods as backend.
2. Set the service's `externalTrafficPolicy` to `Cluster`.
3. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- B. 1. In GKE, create a `Service` of type `NodePort` that uses the application's Pods as backend.
2. Create a Compute Engine instance called `proxy` with 2 network interfaces, one in each VPC.
3. Use iptables on this instance to forward traffic from `gce-network` to the GKE nodes.
4. Configure the Compute Engine instance to use the address of `proxy` in `gce-network` as endpoint.
- C. 1. In GKE, create a `Service` of type `LoadBalancer` that uses the application's Pods as backend.
2. Add an annotation to this service: `cloud.google.com/load-balancer-type: Internal`
3. Peer the two VPCs together.
4. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- D. 1. In GKE, create a `Service` of type `LoadBalancer` that uses the application's Pods as backend.
2. Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG's instances.
3. Configure the Compute Engine instance to use the address of the load balancer that has been created.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 34

You want to run a single caching HTTP reverse proxy on GCP for a latency-sensitive website. This specific reverse proxy consumes almost no CPU. You want to have a 30-GB in-memory cache, and need an additional 2 GB of memory for the rest of the processes. You want to minimize cost. How should you run this reverse proxy?

- A. Create a Cloud Memorystore for Redis instance with 32-GB capacity.
- B. Run it on Compute Engine, and choose a custom instance type with 6 vCPUs and 32 GB of memory.
- C. Package it in a container image, and run it on Kubernetes Engine, using `n1-standard-32` instances as nodes.
- D. Run it on Compute Engine, choose the instance type `n1-standard-1`, and add an SSD persistent disk of 32 GB.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 35

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?

- A. 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`.
- B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud.
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.
- C. 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.
- D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in Google Cloud.
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`.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 36

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?

- A. 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.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 37

Your company has an existing GCP organization with hundreds of projects and a billing account. Your company recently acquired another company that also has hundreds of projects and its own billing account. You would like to consolidate all GCP costs of both GCP organizations onto a single invoice. You would like to consolidate all costs as of tomorrow. What should you do?

- A. Link the acquired company's projects to your company's billing account.
- B. Configure the acquired company's billing account and your company's billing account to export the billing data into the same BigQuery dataset.
- C. Migrate the acquired company's projects into your company's GCP organization. Link the migrated projects to your company's billing account.
- D. Create a new GCP organization and a new billing account. Migrate the acquired company's projects and your company's projects into the new GCP organization and link the projects to the new billing account.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/resource-manager/docs/migrating-projects-billing>

QUESTION 38

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called `platform-logs`. You have already installed the Cloud Logging agent on all the instances. You want to minimize cost. What should you do?

- A.
 1. Give the BigQuery Data Editor role on the `platform-logs` dataset to the service accounts used by your instances.
 2. Update your instances' metadata to add the following value: `logs-destination: bq://platform-logs`.
- B.
 1. In Cloud Logging, create a logs export with a Cloud Pub/Sub topic called `logs` as a sink.
 2. Create a Cloud Function that is triggered by messages in the `logs` topic.
 3. Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the `platform-logs` dataset.
- C.
 1. In Cloud Logging, create a filter to view only Compute Engine logs.
 2. Click Create Export.
 3. Choose BigQuery as Sink Service, and the `platform-logs` dataset as Sink Destination.

- D. 1. Create a Cloud Function that has the BigQuery User role on the `platform-logs` dataset.
2. Configure this Cloud Function to create a BigQuery Job that executes this query:

```
INSERT INTO dataset.platform-logs (timestamp, log)
SELECT timestamp, log FROM compute.logs
WHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)
```

3. Use Cloud Scheduler to trigger this Cloud Function once a day.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 39

You are using Deployment Manager to create a Google Kubernetes Engine cluster. Using the same Deployment Manager deployment, you also want to create a DaemonSet in the `kube-system` namespace of the cluster. You want a solution that uses the fewest possible services. What should you do?

- A. Add the cluster's API as a new Type Provider in Deployment Manager, and use the new type to create the DaemonSet.
- B. Use the Deployment Manager Runtime Configurator to create a new Config resource that contains the DaemonSet definition.
- C. With Deployment Manager, create a Compute Engine instance with a startup script that uses `kubectl` to create the DaemonSet.
- D. In the cluster's definition in Deployment Manager, add a metadata that has `kube-system` as key and the DaemonSet manifest as value.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/kubernetes-engine/docs/how-to/cluster-access-for-kubectl>

QUESTION 40

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?

- A. Review details of the myapp-service Service object and check for error messages.
- B. Review details of the myapp-deployment Deployment object and check for error messages.
- C. Review details of myapp-deployment-58ddbbb995-lp86m Pod and check for warning messages.
- D. View logs of the container in myapp-deployment-58ddbbb995-lp86m pod and check for warning messages.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/run/docs/gke/troubleshooting>

QUESTION 41

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?

- A. Run `gcloud projects list` to get the project ID, and then run `gcloud services list --project <project ID>`.
- B. Run `gcloud init` to set the current project to my-project, and then run `gcloud services list --available`.
- C. Run `gcloud info` to view the account value, and then run `gcloud services list --account <Account>`.
- D. Run `gcloud projects describe <project ID>` to verify the project value, and then run `gcloud services list --available`.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 42

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?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 43

You need to provide a cost estimate for a Kubernetes cluster using the GCP pricing calculator for Kubernetes. Your workload requires high IOPs, and you will also be using disk snapshots. You start by entering the number of nodes, average hours, and average days. What should you do next?

- A. Fill in local SSD. Fill in persistent disk storage and snapshot storage.
- B. Fill in local SSD. Add estimated cost for cluster management.
- C. Select Add GPUs. Fill in persistent disk storage and snapshot storage.
- D. Select Add GPUs. Add estimated cost for cluster management.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/products/calculator#tab=container>

QUESTION 44

You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?

- A. Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
- B. Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service.
- C. Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluster. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.
- D. Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application. Forward the public traffic to HAProxy with an iptable rule. Configure the DNS name of your application using the public IP of the node HAProxy is running on.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer>

QUESTION 45

You need to enable traffic between multiple groups of Compute Engine instances that are currently running two different GCP projects. Each group of Compute Engine instances is running in its own VPC. What should you do?

- A. Verify that both projects are in a GCP Organization. Create a new VPC and add all instances.
- B. Verify that both projects are in a GCP Organization. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.

- C. Verify that you are the Project Administrator of both projects. Create two new VPCs and add all instances.
- D. Verify that you are the Project Administrator of both projects. Create a new VPC and add all instances.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 46

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?

- A. Create a custom role with view-only project permissions. Add the user's account to the custom role.
- B. Create a custom role with view-only service permissions. Add the user's account to the custom role.
- C. Select the built-in IAM project Viewer role. Add the user's account to this role.
- D. Select the built-in IAM service Viewer role. Add the user's account to this role.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/resource-manager/docs/access-control-proj>

QUESTION 47

You are operating a Google Kubernetes Engine (GKE) cluster for your company where different teams can run non-production workloads. Your Machine Learning (ML) team needs access to Nvidia Tesla P100 GPUs to train their models. You want to minimize effort and cost. What should you do?

- A. Ask your ML team to add the `accelerator: gpu` annotation to their pod specification.
- B. Recreate all the nodes of the GKE cluster to enable GPUs on all of them.
- C. Create your own Kubernetes cluster on top of Compute Engine with nodes that have GPUs. Dedicate this cluster to your ML team.
- D. Add a new, GPU-enabled, node pool to the GKE cluster. Ask your ML team to add the `cloud.google.com/gke -accelerator: nvidia-tesla-p100` `nodeSelector` to their pod specification.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 48

You have a large 5-TB AVRO file stored in a Cloud Storage bucket. Your analysts are proficient only in SQL and need access to the data stored in this file. You want to find a cost-effective way to complete their request as soon as possible. What should you do?

- A. Load data in Cloud Datastore and run a SQL query against it.
- B. Create a BigQuery table and load data in BigQuery. Run a SQL query on this table and drop this table after you complete your request.
- C. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.
- D. Create a Hadoop cluster and copy the AVRO file to NDfs by compressing it. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 49

You are migrating a production-critical on-premises application that requires 96 vCPUs to perform its task. You want to make sure the application runs in a similar environment on GCP. What should you do?

- A. When creating the VM, use machine type n1-standard-96.
- B. When creating the VM, use Intel Skylake as the CPU platform.
- C. Create the VM using Compute Engine default settings. Use gcloud to modify the running instance to have 96 vCPUs.
- D. Start the VM using Compute Engine default settings, and adjust as you go based on Rightsizing Recommendations.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 50

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?

- A. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Coldline Storage.
- B. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Nearline Storage.
- C. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Coldline Storage.
- D. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Nearline Storage.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/storage/docs/managing-lifecycles>

QUESTION 51

Your company's infrastructure is on-premises, but all machines are running at maximum capacity. You want to burst to Google Cloud. The workloads on Google Cloud must be able to directly communicate to the workloads on-premises using a private IP range. What should you do?

- A. In Google Cloud, configure the VPC as a host for Shared VPC.
- B. In Google Cloud, configure the VPC for VPC Network Peering.
- C. Create bastion hosts both in your on-premises environment and on Google Cloud. Configure both as proxy servers using their public IP addresses.
- D. Set up Cloud VPN between the infrastructure on-premises and Google Cloud.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 52

You want to select and configure a solution for storing and archiving data on Google Cloud Platform. You need to support compliance objectives for data from one geographic location. This data is archived after 30 days and needs to be accessed annually. What should you do?

- A. Select Multi-Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.
- B. Select Multi-Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
- C. Select Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
- D. Select Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 53

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?

- A. Go to Data Catalog and search for `employee_ssn` in the search box.
- B. Write a shell script that uses the `bq` command line tool to loop through all the projects in your organization.
- C. 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.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 54

You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use `kubectl` to examine the status of your Pod and observe that one of them is still in `Pending` status:

```
$ kubectl get pods -l app=myapp
NAME                                READY   STATUS    RESTARTS   AGE
myapp-deployment-58ddb995-lp86m    0/1     Pending   0           9m
myapp-deployment-58ddb995-qjpkq    1/1     Running   0           9m
```

What is the most likely cause?

- A. The pending Pod's resource requests are too large to fit on a single node of the cluster.
- B. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.

- C. The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.
- D. The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' status. It is currently being rescheduled on a new node.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 55

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?

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 56

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
)
```

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 57

Your finance team wants to view the billing report for your projects. You want to make sure that the finance team does not get additional permissions to the project. What should you do?

- A. Add the group for the finance team to roles/billing user role.
- B. Add the group for the finance team to roles/billing admin role.
- C. Add the group for the finance team to roles/billing viewer role.
- D. Add the group for the finance team to roles/billing project/Manager role.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 58

Your organization has strict requirements to control access to Google Cloud projects. You need to enable your Site Reliability Engineers (SREs) to approve requests from the Google Cloud support team when an SRE opens a support case. You want to follow Google-recommended practices. What should you do?

- A. Add your SREs to roles/iam.roleAdmin role.
- B. Add your SREs to roles/accessapproval.approver role.
- C. Add your SREs to a group and then add this group to roles/iam.roleAdmin.role.
- D. Add your SREs to a group and then add this group to roles/accessapproval.approver role.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 59

You need to host an application on a Compute Engine instance in a project shared with other teams. You want to prevent the other teams from accidentally causing downtime on that application. Which feature should you use?

- A. Use a Shielded VM.
- B. Use a Preemptible VM.
- C. Use a sole-tenant node.
- D. Enable deletion protection on the instance.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 60

Your organization needs to grant users access to query datasets in BigQuery but prevent them from accidentally deleting the datasets. You want a solution that follows Google-recommended practices. What should you do?

- A. Add users to roles/bigquery user role only, instead of roles/bigquery dataOwner.
- B. Add users to roles/bigquery dataEditor role only, instead of roles/bigquery dataOwner.
- C. Create a custom role by removing delete permissions, and add users to that role only.
- D. Create a custom role by removing delete permissions. Add users to the group, and then add the group to the custom role.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 61

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?

- A. In Cloud Identity, set up SSO with Google as an identity provider to access custom SAML apps.
- B. In Cloud Identity, set up SSO with a third-party identity provider with Google as a service provider.
- C. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Mobile & Desktop Apps.
- D. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Web Server Applications.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 62

You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

- A. Cold Storage
- B. Nearline Storage
- C. Regional Storage
- D. Multi-Regional Storage

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Nearline, Coldline, and Archive offer ultra low-cost, highly-durable, highly available archival storage. For data accessed less than once a year, Archive is a cost-

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

Reference: <https://cloud.google.com/storage/archival>

QUESTION 63

A team of data scientists infrequently needs to use a Google Kubernetes Engine (GKE) cluster that you manage. They require GPUs for some long-running, non-restartable jobs. You want to minimize cost. What should you do?

- A. Enable node auto-provisioning on the GKE cluster.
- B. Create a VerticalPodAutscaler for those workloads.
- C. Create a node pool with preemptible VMs and GPUs attached to those VMs.
- D. Create a node pool of instances with GPUs, and enable autoscaling on this node pool with a minimum size of 1.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/kubernetes-engine/docs/how-to/gpus>

QUESTION 64

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?

- A. Use Google Cloud Directory Sync (GCDS) to synchronize users into Cloud Identity.
- B. Use the cloud Identity APIs and write a script to synchronize users to Cloud Identity.
- C. Export users from Active Directory as a CSV and import them to Cloud Identity via the Admin Console.
- D. Ask each employee to create a Google account using self signup. Require that each employee use their company email address and password.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction>

QUESTION 65

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 has asked you to follow Google-recommended practices. What should you do?

- A. 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.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 66

Your management has asked an external auditor to review all the resources in a specific project. The security team has enabled the Organization Policy called Domain Restricted Sharing on the organization node by specifying only your Cloud Identity domain. You want the auditor to only be able to view, but not modify, the resources in that project. What should you do?

- A. Ask the auditor for their Google account, and give them the Viewer role on the project.
- B. Ask the auditor for their Google account, and give them the Security Reviewer role on the project.
- C. Create a temporary account for the auditor in Cloud Identity, and give that account the Viewer role on the project.
- D. Create a temporary account for the auditor in Cloud Identity, and give that account the Security Reviewer role on the project.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 67

You have a workload running on Compute Engine that is critical to your business. You want to ensure that the data on the boot disk of this workload is backed up regularly. You need to be able to restore a backup as quickly as possible in case of disaster. You also want older backups to be cleaned automatically to save on cost. You want to follow Google-recommended practices. What should you do?

- A. Create a Cloud Function to create an instance template.
- B. Create a snapshot schedule for the disk using the desired interval.
- C. Create a `cron` job to create a new disk from the disk using `gcloud`.
- D. Create a Cloud Task to create an image and export it to Cloud Storage.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 68

You need to assign a Cloud Identity and Access Management (Cloud IAM) role to an external auditor. The auditor needs to have permissions to review your Google Cloud Platform (GCP) Audit Logs and also to review your Data Access logs. What should you do?

- A. Assign the auditor the IAM role `roles/logging.privateLogViewer`. Perform the export of logs to Cloud Storage.
- B. Assign the auditor the IAM role `roles/logging.privateLogViewer`. Direct the auditor to also review the logs for changes to Cloud IAM policy.
- C. Assign the auditor's IAM user to a custom role that has `logging.privateLogEntries.list` permission. Perform the export of logs to Cloud Storage.
- D. Assign the auditor's IAM user to a custom role that has `logging.privateLogEntries.list` permission. Direct the auditor to also review the logs for changes to Cloud IAM policy.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 69

You are managing several Google Cloud Platform (GCP) projects and need access to all logs for the past 60 days. You want to be able to explore and quickly analyze the log contents. You want to follow Google-recommended practices to obtain the combined logs for all projects. What should you do?

- A. Navigate to Stackdriver Logging and select `resource.labels.project_id="*"`
- B. Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days.
- C. Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days.
- D. Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging>

QUESTION 70

You need to reduce GCP service costs for a division of your company using the fewest possible steps. You need to turn off all configured services in an existing GCP project. What should you do?

- A. 1. Verify that you are assigned the Project Owners IAM role for this project.
2. Locate the project in the GCP console, click **Shut down** and then enter the project ID.
- B. 1. Verify that you are assigned the Project Owners IAM role for this project.
2. Switch to the project in the GCP console, locate the resources and delete them.
- C. 1. Verify that you are assigned the Organizational Administrator IAM role for this project.
2. Locate the project in the GCP console, enter the project ID and then click **Shut down**.
- D. 1. Verify that you are assigned the Organizational Administrators IAM role for this project.
2. Switch to the project in the GCP console, locate the resources and delete them.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 71

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

- A. Give "project owner" for web-applications appropriate roles to crm-databases-proj.
- B. Give "project owner" role to crm-databases-proj and the web-applications project.
- C. Give "project owner" role to crm-databases-proj and bigquery.dataViewer role to web-applications.
- D. Give bigquery.dataViewer role to crm-databases-proj and appropriate roles to web-applications.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging>

QUESTION 72

You need to create a custom IAM role for use with a GCP service. All permissions in the role must be suitable for production use. You also want to clearly share with your organization the status of the custom role. This will be the first version of the custom role. What should you do?

- A. Use permissions in your role that use the 'supported' support level for role permissions. Set the role stage to ALPHA while testing the role permissions.
- B. Use permissions in your role that use the 'supported' support level for role permissions. Set the role stage to BETA while testing the role permissions.
- C. Use permissions in your role that use the 'testing' support level for role permissions. Set the role stage to ALPHA while testing the role permissions.
- D. Use permissions in your role that use the 'testing' support level for role permissions. Set the role stage to BETA while testing the role permissions.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:**QUESTION 73**

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?

- A. Upload the data to BigQuery using the `bq` command line tool.
- B. Upload the data to Cloud Storage using the `gsutil` command line tool.
- C. Upload the data into Cloud SQL using the import function in the console.
- D. Upload the data into Cloud Spanner using the import function in the console.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 74

You need to manage multiple Google Cloud projects in the fewest steps possible. You want to configure the Google Cloud SDK command line interface (CLI) so that you can easily manage multiple projects. What should you do?

- A. 1. Create a configuration for each project you need to manage.
2. Activate the appropriate configuration when you work with each of your assigned Google Cloud projects.
- B. 1. Create a configuration for each project you need to manage.
2. Use `gcloud init` to update the configuration values when you need to work with a non-default project
- C. 1. Use the default configuration for one project you need to manage.
2. Activate the appropriate configuration when you work with each of your assigned Google Cloud projects.
- D. 1. Use the default configuration for one project you need to manage.
2. Use `gcloud init` to update the configuration values when you need to work with a non-default project.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 75

Your managed instance group raised an alert stating that new instance creation has failed to create new instances. You need to maintain the number of running instances specified by the template to be able to process expected application traffic. What should you do?

- A. Create an instance template that contains valid syntax which will be used by the instance group. Delete any persistent disks with the same name as instance names.
- B. Create an instance template that contains valid syntax that will be used by the instance group. Verify that the instance name and persistent disk name values are not the same in the template.
- C. Verify that the instance template being used by the instance group contains valid syntax. Delete any persistent disks with the same name as instance names. Set the `disks.autoDelete` property to `true` in the instance template.
- D. Delete the current instance template and replace it with a new instance template. Verify that the instance name and persistent disk name values are not the same in the template. Set the `disks.autoDelete` property to `true` in the instance template.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 76

Your company is moving from an on-premises environment to Google Cloud. 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 Google Cloud quickly and with minimal support effort. What should you do?

- A. 1. Build an instruction guide to install Cassandra on Google Cloud.
2. Make the instruction guide accessible to your developers.
- B. 1. Advise your developers to go to Cloud Marketplace.
2. Ask the developers to launch a Cassandra image for their development work.
- C. 1. Build a Cassandra Compute Engine instance and take a snapshot of it.
2. Use the snapshot to create instances for your developers.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 77

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?

- A. 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.
- B. 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.
- C. 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.
- D. 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.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 78

You have an application that uses Cloud Spanner as a backend database. The application has a very predictable traffic pattern. You want to automatically scale up or down the number of Spanner nodes depending on traffic. What should you do?

- A. Create a `cron` job that runs on a scheduled basis to review Cloud Monitoring metrics, and then resize the Spanner instance accordingly.
- B. Create a Cloud Monitoring alerting policy to send an alert to oncall SRE emails when Cloud Spanner CPU exceeds the threshold. SREs would scale resources up or down accordingly.
- C. Create a Cloud Monitoring alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshold. Google support would scale resources up or down accordingly.
- D. Create a Cloud Monitoring alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 79

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. What should you do?

- A. Set up a budget alert on the project with an amount of 100 dollars, a threshold of 100%, and notification type of "email."
- B. Set up a budget alert on the billing account with an amount of 100 dollars, a threshold of 100%, and notification type of "email."
- C. 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.
- D. Use the Cloud Logging Agent to export the Apache web server logs to Cloud Logging. Create a Cloud Function that uses BigQuery to parse the HTTP response log data in Cloud Logging for the current month and sends an email if the size of all HTTP responses, multiplied by current Google Cloud egress prices, totals over 100 dollars. Schedule the Cloud Function using Cloud Scheduler to run hourly.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 80

You have an application that receives SSL-encrypted TCP traffic on port 443. Clients for this application are located all over the world. You want to minimize latency for the clients. Which load balancing option should you use?

- A. HTTPS Load Balancer
- B. Network Load Balancer
- C. SSL Proxy Load Balancer
- D. Internal TCP/UDP Load Balancer. Add a firewall rule allowing ingress traffic from 0.0.0.0/0 on the target instances.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/load-balancing/docs/ssl>

QUESTION 81

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?

- A. Create a Compute Engine snapshot of your base VM. Create your images from that snapshot.
- B. Create a Compute Engine snapshot of your base VM. Create your instances from that snapshot.
- C. Create a custom Compute Engine image from a snapshot. Create your images from that image.
- D. Create a custom Compute Engine image from a snapshot. Create your instances from that image.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

Reference: <https://cloud.google.com/compute/docs/instances/create-start-instance>

QUESTION 82

Your company runs one batch process in an on-premises server that takes around 30 hours to complete. The task runs monthly, can be performed offline, and must be restarted if interrupted. You want to migrate this workload to the cloud while minimizing cost. What should you do?

- A. Migrate the workload to a Compute Engine Preemptible VM.
- B. Migrate the workload to a Google Kubernetes Engine cluster with Preemptible nodes.
- C. Migrate the workload to a Compute Engine VM. Start and stop the instance as needed.
- D. Create an Instance Template with Preemptible VMs On. Create a Managed Instance Group from the template and adjust Target CPU Utilization. Migrate the workload.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/kubernetes-engine/docs/tutorials/migrating-node-pool>

QUESTION 83

You are developing a new application and are looking for a Jenkins installation to build and deploy your source code. You want to automate the installation as quickly and easily as possible. What should you do?

- A. Deploy Jenkins through the Google Cloud Marketplace.
- B. Create a new Compute Engine instance. Run the Jenkins executable.
- C. Create a new Kubernetes Engine cluster. Create a deployment for the Jenkins image.
- D. Create an instance template with the Jenkins executable. Create a managed instance group with this template.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/jenkins-on-kubernetes-engine>

QUESTION 84

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?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 85

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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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

QUESTION 86

Your company has workloads running on Compute Engine and on-premises. The Google Cloud Virtual Private Cloud (VPC) is connected to your WAN over a Virtual Private Network (VPN). You need to deploy a new Compute Engine instance and ensure that no public Internet traffic can be routed to it. What should you do?

- A. Create the instance without a public IP address.
- B. Create the instance with Private Google Access enabled.
- C. Create a deny-all egress firewall rule on the VPC network.
- D. Create a route on the VPC to route all traffic to the instance over the VPN tunnel.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Get private access to Google services, such as storage, big data, analytics, or machine learning, without having to give your service a public IP address.

Reference: <https://cloud.google.com/vpc>

QUESTION 87

Your team maintains the infrastructure for your organization. The current infrastructure requires changes. You need to share your proposed changes with the rest of the team. You want to follow Google's recommended best practices. What should you do?

- A. Use Deployment Manager templates to describe the proposed changes and store them in a Cloud Storage bucket.
- B. Use Deployment Manager templates to describe the proposed changes and store them in Cloud Source Repositories.
- C. Apply the changes in a development environment, run `gcloud compute instances list`, and then save the output in a shared Storage bucket.
- D. Apply the changes in a development environment, run `gcloud compute instances list`, and then save the output in Cloud Source Repositories.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 88

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?

- A. Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
- B. Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
- C. Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
- D. 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.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 89

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?

- A. – Create Compute Engine resources in `us-central1-b`.

- Balance the load across both us-central1-a and us-central1-b.
- B. – Create a Managed Instance Group and specify us-central1-a as the zone.
 - Configure the Health Check with a short Health Interval.
- C. – Create an HTTP(S) Load Balancer.
 - Create one or more global forwarding rules to direct traffic to your VMs.
- D. – Perform regular backups of your application.
 - Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.
 - Restore from backups when notified.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 90

You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

- A. Create a new subnet in the same region as the subnet being used.
- B. Add an alias IP range to the subnet used by the GKE clusters.
- C. Create a new VPC, and set up VPC peering with the existing VPC.
- D. Expand the CIDR range of the relevant subnet for the cluster.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 91

You are working with a user to set up an application in a new VPC behind a firewall. The user is concerned about data egress. You want to configure the fewest open egress ports. What should you do?

- A. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.
- B. Set up a high-priority (1000) rule that pairs both ingress and egress ports.
- C. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.

D. Set up a high-priority (1000) rule to allow the appropriate ports.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 92

You are deploying a production application on Compute Engine. You want to prevent anyone from accidentally destroying the instance by clicking the wrong button. What should you do?

- A. Disable the flag "Delete boot disk when instance is deleted."
- B. Enable delete protection on the instance.
- C. Disable Automatic restart on the instance.
- D. Enable Preemptibility on the instance.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://googlecloudplatform.uservoice.com/forums/302595-compute-engine/suggestions/14227521-set-delete-boot-disk-when-instance-is-deleted-to>

QUESTION 93

Your company uses a large number of Google Cloud services centralized in a single project. All teams have specific projects for testing and development. The DevOps team needs access to all of the production services in order to perform their job. You want to prevent Google Cloud product changes from broadening their permissions in the future. You want to follow Google-recommended practices. What should you do?

- A. Grant all members of the DevOps team the role of Project Editor on the organization level.
- B. Grant all members of the DevOps team the role of Project Editor on the production project.
- C. Create a custom role that combines the required permissions. Grant the DevOps team the custom role on the production project.
- D. Create a custom role that combines the required permissions. Grant the DevOps team the custom role on the organization level.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 94

You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:

- Restrict access so that suppliers can access only their own data.
- Give suppliers write access to data only for 30 minutes.
- Delete data that is over 45 days old.

You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

- A. Build a lifecycle policy to delete Cloud Storage objects after 45 days.
- B. Use signed URLs to allow suppliers limited time access to store their objects.
- C. Set up an SFTP server for your application, and create a separate user for each supplier.
- D. Build a Cloud function that triggers a timer of 45 days to delete objects that have expired.
- E. Develop a script that loops through all Cloud Storage buckets and deletes any buckets that are older than 45 days.

Correct Answer: AE

Section: (none)

Explanation

Explanation/Reference:

QUESTION 95

Your company wants to standardize the creation and management of multiple Google Cloud resources using Infrastructure as Code. You want to minimize the amount of repetitive code needed to manage the environment. What should you do?

- A. Develop templates for the environment using Cloud Deployment Manager.
- B. Use `curl` in a terminal to send a REST request to the relevant Google API for each individual resource.
- C. Use the Cloud Console interface to provision and manage all related resources.
- D. Create a bash script that contains all requirement steps as `gcloud` commands.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/deployment-manager/docs/fundamentals> (see templates)

QUESTION 96

You are performing a monthly security check of your Google Cloud environment and want to know who has access to view data stored in your Google Cloud Project. What should you?

- A. Enable Audit Logs for all APIs that are related to data storage.
- B. Review the IAM permissions for any role that allows for data access.
- C. Review the Identity-Aware Proxy settings for each resource.
- D. Create a Data Loss Prevention job.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/compute/docs/access>

QUESTION 97

Your company has embraced a hybrid cloud strategy where some of the applications are deployed on Google Cloud. A Virtual Private Network (VPN) tunnel connects your Virtual Private Cloud (VPC) in Google Cloud with your company's on-premises network. Multiple applications in Google Cloud need to connect to an on-premises database server, and you want to avoid having to change the IP configuration in all of your applications when the IP of the database changes. What should you do?

- A. Configure Cloud NAT for all subnets of your VPC to be used when egressing from the VM instances.
- B. Create a private zone on Cloud DNS, and configure the applications with the DNS name.
- C. Configure the IP of the database as custom metadata for each instance, and query the metadata server.
- D. Query the Compute Engine internal DNS from the applications to retrieve the IP of the database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 98

You have experimented with Google Cloud using your own credit card and expensed the costs to your company. Your company wants to streamline the billing process and charge the costs of your projects to their monthly invoice. What should you do?

- A. Grant the financial team the IAM role of “Billing Account User” on the billing account linked to your credit card.
- B. Set up BigQuery billing export and grant your financial department IAM access to query the data.
- C. Create a ticket with Google Billing Support to ask them to send the invoice to your company.
- D. Change the billing account of your projects to the billing account of your company.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 99

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?

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 100

Your web application has been running successfully on Cloud Run for Anthos. You want to evaluate an updated version of the application with a specific percentage of your production users (canary deployment). What should you do?

- A. Create a new service with the new version of the application. Split traffic between this version and the version that is currently running.
- B. Create a new revision with the new version of the application. Split traffic between this version and the version that is currently running.
- C. Create a new service with the new version of the application. Add HTTP Load Balancer in front of both services.
- D. Create a new revision with the new version of the application. Add HTTP Load Balancer in front of both revisions.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 101

Your company developed a mobile game that is deployed on Google Cloud. Gamers are connecting to the game with their personal phones over the Internet. The game sends UDP packets to update the servers about the gamers' actions while they are playing in multiplayer mode. Your game backend can scale over multiple virtual machines (VMs), and you want to expose the VMs over a single IP address. What should you do?

- A. Configure an SSL Proxy load balancer in front of the application servers.
- B. Configure an Internal UDP load balancer in front of the application servers.
- C. Configure an External HTTP(s) load balancer in front of the application servers.
- D. Configure an External Network load balancer in front of the application servers.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/solutions/connecting-securely>

QUESTION 102

Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

- A. Turn on Data Access Logs for the buckets they want to audit, and then build a query in the log viewer that filters on Cloud Storage.
- B. Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs.
- C. Assign the appropriate permissions, and the use Cloud Monitoring to review metrics.
- D. Use the export logs API to provide the Admin Activity Audit Logs in the format they want.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/storage/docs/audit-logging>

QUESTION 103

You received a JSON file that contained a private key of a Service Account in order to get access to several resources in a Google Cloud project. You downloaded and installed the Cloud SDK and want to use this private key for authentication and authorization when performing gcloud commands. What should you do?

- A. Use the command `gcloud auth login` and point it to the private key.
- B. Use the command `gcloud auth activate-service-account` and point it to the private key.
- C. Place the private key file in the installation directory of the Cloud SDK and rename it to "credentials.json".
- D. Place the private key file in your home directory and rename it to "GOOGLE_APPLICATION_CREDENTIALS".

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Reference: <https://cloud.google.com/sdk/docs/authorizing>

