

Certification: Associate Cloud Engineer
Certification Full Name: Associate Cloud Engineer



Exam A

QUESTION 1

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

- A. Check the app.yaml file for your application and check project settings.
- B. Check the web-application.xml file for your application and check project settings.
- C. Go to Deployment Manager and review settings for deployment of applications.
- D. Go to Cloud Shell and run gcloud config list to review the Google Cloud configuration used for deployment.

Correct Answer: D

Section:

Explanation:

C:\GCP\appeng>gcloud config list

[core]

account = xxx@gmail.com

disable_usage_reporting = False

project = my-first-demo-xxxx

<https://cloud.google.com/endpoints/docs/openapi/troubleshoot-gce-deployment>

QUESTION 2

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: A

Section:

Explanation:

Ref:<https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>

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

QUESTION 3

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:

Explanation:

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

QUESTION 4

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:

Explanation:

In Google compute engine, if predefined machine types don't meet your needs, you can create an instance with custom virtualized hardware settings. Specifically, you can create an instance with a custom number of vCPUs and custom memory, effectively using a custom machine type. Custom machine types are ideal for the following scenarios: 1. Workloads that aren't a good fit for the predefined machine types that are available to you. 2. Workloads that require more processing power or more memory but don't need all of the upgrades that are provided by the next machine type level. In our scenario, we only need a memory upgrade. Moving to a bigger instance would also bump up the CPU which we don't need so we have to use a custom machine type. It is not possible to change memory while the instance is running so you need to first stop the instance, change the memory and then start it again. See below a screenshot that shows how CPU/Memory can be customized for an instance that has been stopped.

Ref:<https://cloud.google.com/compute/docs/instances/creating-instance-with-custom-machine-type>

QUESTION 5

You have production and test workloads that you want to deploy on Compute Engine. Production VMs need to be in a different subnet than the test VMs. All the VMs must be able to reach each other over internal IP without creating additional routes. You need to set up VPC and the 2 subnets. Which configuration meets these requirements?

- A. Create a single custom VPC with 2 subnets. Create each subnet in a different region and with a different CIDR range.
- B. Create a single custom VPC with 2 subnets. Create each subnet in the same region and with the same CIDR range.
- C. Create 2 custom VPCs, each with a single subnet. Create each subnet in a different region and with a different CIDR range.
- D. Create 2 custom VPCs, each with a single subnet. Create each subnet in the same region and with the same CIDR range.

Correct Answer: A

Section:

Explanation:

When we create subnets in the same VPC with different CIDR ranges, they can communicate automatically within VPC. Resources within a VPC network can communicate with one another by using internal (private) IPv4 addresses, subject to applicable network firewall rules

Ref:<https://cloud.google.com/vpc/docs/vpc>

QUESTION 6

You need to create an autoscaling managed instance group for an HTTPS web application. You want to make sure that unhealthy VMs are recreated. What should you do?

- A. Create a health check on port 443 and use that when creating the Managed Instance Group.

- B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.
- C. In the Instance Template, add the label 'health-check'.
- D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Correct Answer: A

Section:

Explanation:

https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs#setting_up_an_autohealing_policy

QUESTION 7

Your company has a Google Cloud Platform project that uses BigQuery for data warehousing. Your data science team changes frequently and has few members. You need to allow members of this team to perform queries. You want to follow Google-recommended practices. What should you do?

- A. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery jobUser role to the group.
- B. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery dataViewer user role to the group.
- C. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery jobUser role to the group.
- D. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery dataViewer user role to the group.

Correct Answer: C

Section:

Explanation:

Read the dataset's metadata and to list tables in the dataset. Read data and metadata from the dataset's tables. When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

BigQuery Data Viewer

(roles/bigquery.dataViewer)

When applied to a table or view, this role provides permissions to:

Read data and metadata from the table or view.

This role cannot be applied to individual models or routines.

When applied to a dataset, this role provides permissions to:

Read the dataset's metadata and list tables in the dataset.

Read data and metadata from the dataset's tables.

When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

Lowest-level resources where you can grant this role:

Table

View

BigQuery Job User

(roles/bigquery.jobUser)

Provides permissions to run jobs, including queries, within the project.

Lowest-level resources where you can grant this role:

Project

to run jobs <https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUser> databaseUser needs additional role permission to run jobs <https://cloud.google.com/spanner/docs/iam#spanner.databaseUser>

QUESTION 8

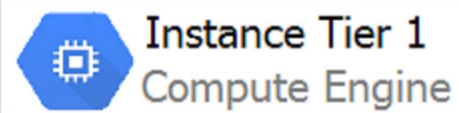
Your company has a 3-tier solution running on Compute Engine. The configuration of the current infrastructure is shown below.




Google Cloud Platform

VPC

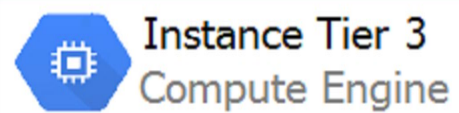
Subnet Tier#1 10.0.1.0/24



Subnet Tier#2 10.0.2.0/24



Subnet Tier#3 10.0.3.0/24



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 all2. 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:80802. 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 all2. 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: 80802. 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:

Explanation:

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

QUESTION 9

You are given a project with a single virtual private cloud (VPC) and a single subnetwork in the us-central1 region. There is a Compute Engine instance hosting an application in this subnetwork. You need to deploy a new instance in the same project in the europe-west1 region. This new instance needs access to the application. You want to follow Google-recommended practices. What should you do?

1. Create a subnetwork in the same VPC, in europe-west1.2. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
1. Create a VPC and a subnetwork in europe-west1.2. Expose the application with an internal load balancer.3. Create the new instance in the new subnetwork and use the load balancer's address as the endpoint.
1. Create a subnetwork in the same VPC, in europe-west1.2. Use Cloud VPN to connect the two subnetworks.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
1. Create a VPC and a subnetwork in europe-west1.2. Peer the 2 VPCs.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

Correct Answer: C

Section:

Explanation:

Given that the new instance wants to access the application on the existing compute engine instance, these applications seem to be related so they should be within the same VPC. It is possible to have them in different VPCs and peer the VPCs but this is a lot of additional work and we can simplify this by choosing the option below (which is the answer)

1. Create a subnet in the same VPC, in europe-west1.
2. Create the new instance in the new subnet and use the first instance subnets private address as the endpoint. is the right answer.

We can create another subnet in the same VPC and this subnet is located in europe-west1. We can then spin up a new instance in this subnet. We also have to set up a firewall rule to allow communication between the two subnets. All instances in the two subnets with the same VPC can communicate through the internal IP Address

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

QUESTION 10

Your projects incurred more costs than you expected last month. Your research reveals that a development GKE container emitted a huge number of logs, which resulted in higher costs. You want to disable the logs quickly using the minimum number of steps. What should you do?

1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE container resource.
1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE Cluster Operations resource.
1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Logging.
1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Monitoring.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/logging/docs/api/v2/resource-list>

GKE Containers have more log than GKE Cluster Operations:

.-GKE Containe:

cluster_name: An immutable name for the cluster the container is running in.

namespace_id: Immutable ID of the cluster namespace the container is running in.

instance_id: Immutable ID of the GCE instance the container is running in.

pod_id: Immutable ID of the pod the container is running in.

container_name: Immutable name of the container.
zone: The GCE zone in which the instance is running.
VS
.-GKE Cluster Operations
project_id: The identifier of the GCP project associated with this resource, such as 'my-project'.
cluster_name: The name of the GKE Cluster.
location: The location in which the GKE Cluster is running.

QUESTION 11

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 --splits 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:

Explanation:

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

QUESTION 12

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: B

Section:

Explanation:

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

QUESTION 13

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:

Explanation:

Cloud Spanner is a relational database and is highly scalable. Cloud Spanner is a highly scalable, enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with a non-relational horizontal scale. This combination delivers high-performance transactions and strong consistency across rows, regions, and continents with an industry-leading 99.999% availability SLA, no planned downtime, and enterprise-grade security

Ref: <https://cloud.google.com/spanner>

	CLOUD SPANNER	TRADITIONAL RELATIONAL	TRADITIONAL NON-RELATIONAL
Schema	✓ Yes	✓ Yes	✗ No
SQL	✓ Yes	✓ Yes	✗ No
Consistency	✓ Strong	✓ Strong	✗ Eventual
Availability	✓ High	✗ Failover	✓ High
Scalability	✓ Horizontal	✗ Vertical	✓ Horizontal
Replication	✓ Automatic	🔄 Configurable	🔄 Configurable

QUESTION 14

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: C

Section:

Explanation:

From this source: https://cloud.google.com/billing/docs/how-to/custom-roles#permission_association_and_inheritance

'For example, associating a project with a billing account requires the billing.resourceAssociations.create permission on the billing account and also the resourcemanager.projects.createBillingAssignment permission on the project. This is because project permissions are required for actions where project owners control access, while billing account permissions are required for actions where billing account administrators control access. When both should be involved, both permissions are necessary.'

QUESTION 15

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: Internal3. 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:

Explanation:

performs a peering between the two VPC's (the statement makes sure that this option is feasible since it clearly specifies that there is no overlapping between the ip ranges of both vpc's), deploy the LoadBalancer as internal with the annotation, and configure the endpoint so that the compute engine instance can access the application internally, that is, without the need to have a public ip at any time and therefore, without the need to go outside the google network. The traffic, therefore, never crosses the public internet.

<https://medium.com/pablo-perez/k8s-externaltrafficpolicy-local-or-cluster-40b259a19404>

<https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing>

clients in a VPC network connected to the LoadBalancer network using VPC Network Peering can also access the Service

<https://cloud.google.com/kubernetes-engine/docs/how-to/service-parameters>

QUESTION 16

Your organization is a financial company that needs to store audit log files for 3 years. Your organization has hundreds of Google Cloud projects. You need to implement a cost-effective approach for log file retention. What should you do?

- A. Create an export to the sink that saves logs from Cloud Audit to BigQuery.
- B. Create an export to the sink that saves logs from Cloud Audit to a Coldline Storage bucket.
- C. Write a custom script that uses logging API to copy the logs from Stackdriver logs to BigQuery.
- D. Export these logs to Cloud Pub/Sub and write a Cloud Dataflow pipeline to store logs to Cloud SQL.

Correct Answer: B

Section:

Explanation:

Coldline Storage is the perfect service to store audit logs from all the projects and is very cost-efficient as well. Coldline Storage is a very low-cost, highly durable storage service for storing infrequently accessed data.

QUESTION 17

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: A

Section:

Explanation:

What is Google Cloud Memorystore?

Overview. Cloud Memorystore for Redis is a fully managed Redis service for Google Cloud Platform. Applications running on Google Cloud Platform can achieve extreme performance by leveraging the highly scalable, highly available, and secure Redis service without the burden of managing complex Redis deployments.

QUESTION 18

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 Platform (GCP).2. In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.
- 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 GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

Correct Answer: D

Section:

Explanation:

Our requirement is to follow Google recommended practices to achieve the end result. Configuring Private Google Access for On-Premises Hosts is best achieved by VPN/Interconnect + Advertise Routes + Use restricted Google IP Range.

Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP

Using 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.

In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com is the right answer right, and it is what Google recommends.

Ref:<https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid>

You must configure routes so that Google API traffic is forwarded through your Cloud VPN or Cloud Interconnect connection, firewall rules on your on-premises firewall to allow the outgoing traffic, and DNS so that traffic to Google APIs resolves to the IP range youve added to your routes.

You can use Cloud Router Custom Route Advertisement to announce the Restricted Google APIs IP addresses through Cloud Router to your on-premises network. The Restricted Google APIs IP range is 199.36.153.4/30. While this is technically a public IP range, Google does not announce it publicly. This IP range is only accessible to hosts that can reach your Google Cloud projects through internal IP ranges, such as through a Cloud VPN or Cloud Interconnect connection. Without having a public IP address or access to the internet, the only way you could connect to cloud storage is if you have an internal route to it.

So Negotiate with the security team to be able to give public IP addresses to the servers is not right. Following Google recommended practices is synonymous with using Googles services (Not quite, but it is at least for the exam !!).

So In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance is not right.

Migrating the VM to Compute Engine is a bit drastic when Google says it is perfectly fine to have Hybrid Connectivity architectures<https://cloud.google.com/hybrid-connectivity>.

So,

Use Migrate for Compute Engine (formerly known as Velostrata) to migrate these servers to Compute Engine is not right.

QUESTION 19

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: C

Section:

Explanation:

<https://cloud.google.com/run/docs/tutorials/pubsub#integrating-pubsub>

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.

QUESTION 20

You need to deploy an application, which is packaged in a container image, in a new project. The application exposes an HTTP endpoint and receives very few requests per day. You want to minimize costs. What should you do?

- A. Deploy the container on Cloud Run.
- B. Deploy the container on Cloud Run on GKE.
- C. Deploy the container on App Engine Flexible.
- D. Deploy the container on Google Kubernetes Engine, with cluster autoscaling and horizontal pod autoscaling enabled.

Correct Answer: A

Section:

Explanation:

Cloud Run takes any container images and pairs great with the container ecosystem: Cloud Build, Artifact Registry, Docker. ... No infrastructure to manage: once deployed, Cloud Run manages your services so you can sleep well. Fast autoscaling. Cloud Run automatically scales up or down from zero to N depending on traffic.

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

QUESTION 21

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: A

Section:

Explanation:

https://cloud.google.com/resource-manager/docs/project-migration#oauth_consent_screen

<https://cloud.google.com/resource-manager/docs/project-migration>

QUESTION 22

You built an application on Google Cloud Platform that uses Cloud Spanner. Your support team needs to monitor the environment but should not have access to table data. You need a streamlined solution to grant the correct permissions to your support team, and you want to follow Google-recommended practices. What should you do?

- A. Add the support team group to the roles/monitoring.viewer role
- B. Add the support team group to the roles/spanner.databaseUser role.

- C. Add the support team group to the roles/spanner.databaseReader role.
- D. Add the support team group to the roles/stackdriver.accounts.viewer role.

Correct Answer: A

Section:

Explanation:

roles/monitoring.viewer provides read-only access to get and list information about all monitoring data and configurations. This role provides monitoring access and fits our requirements. roles/monitoring.viewer is the right answer.

Ref:<https://cloud.google.com/iam/docs/understanding-roles#cloud-spanner-roles>

QUESTION 23

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 Stackdriver 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 Stackdriver 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 Stackdriver 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.logsWHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)3. Use Cloud Scheduler to trigger this Cloud Function once a day.

Correct Answer: C

Section:

Explanation:

1. In Stackdriver 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.

QUESTION 24

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: A

Section:

Explanation:

Adding an API as a type provider

This page describes how to add an API to Google Cloud Deployment Manager as a type provider. To learn more about types and type providers, read the Types overview documentation.

A type provider exposes all of the resources of a third-party API to Deployment Manager as base types that you can use in your configurations. These types must be directly served by a RESTful API that supports Create, Read, Update, and Delete (CRUD).

If you want to use an API that is not automatically provided by Google with Deployment Manager, you must add the API as a type provider.

<https://cloud.google.com/deployment-manager/docs/configuration/type-providers/creating-type-provider>

QUESTION 25

You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML. You need to enable authentication to the APIs from your on-premises environment. What should you do?

- A. Use service account credentials in your on-premises application.
- B. Use gcloud to create a key file for the service account that has appropriate permissions.
- C. Set up direct interconnect between your data center and Google Cloud Platform to enable authentication for your on-premises applications.
- D. Go to the IAM & admin console, grant a user account permissions similar to the service account permissions, and use this user account for authentication from your data center.

Correct Answer: B

Section:

Explanation:

To use a service account outside of Google Cloud, such as on other platforms or on-premises, you must first establish the identity of the service account. Public/private key pairs provide a secure way of accomplishing this goal. You can create a service account key using the Cloud Console, the gcloud tool, the `serviceAccounts.keys.create()` method, or one of the client libraries.

Ref:<https://cloud.google.com/iam/docs/creating-managing-service-account-keys>

QUESTION 26

You are using Container Registry to centrally store your company's container images in a separate project. In another project, you want to create a Google Kubernetes Engine (GKE) cluster. You want to ensure that Kubernetes can download images from Container Registry. What should you do?

- A. In the project where the images are stored, grant the Storage Object Viewer IAM role to the service account used by the Kubernetes nodes.
- B. When you create the GKE cluster, choose the Allow full access to all Cloud APIs option under 'Access scopes'.
- C. Create a service account, and give it access to Cloud Storage. Create a P12 key for this service account and use it as an imagePullSecrets in Kubernetes.
- D. Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account.

Correct Answer: A

Section:

Explanation:

As mentioned above, Container Registry ignores permissions set on individual objects within the storage bucket so this isn't going to work.

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

QUESTION 27

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

```

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

```

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

```

kubect1 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:

Explanation:

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-application/#debugging-pods>

QUESTION 28

You are setting up a Windows VM on Compute Engine and want to make sure you can log in to the VM via RDP. What should you do?

- A. After the VM has been created, use your Google Account credentials to log in into the VM.
- B. After the VM has been created, use `gcloud compute reset-windows-password` to retrieve the login credentials for the VM.
- C. When creating the VM, add metadata to the instance using 'windows-password' as the key and a password as the value.

D. After the VM has been created, download the JSON private key for the default Compute Engine service account. Use the credentials in the JSON file to log in to the VM.

Correct Answer: B

Section:

Explanation:

You can generate Windows passwords using either the Google Cloud Console or the gcloud command-line tool. This option uses the right syntax to reset the windows password.

`gcloud compute reset-windows-password windows-instance`

Ref:<https://cloud.google.com/compute/docs/instances/windows/creating-passwords-for-windows-instances#gcloud>

QUESTION 29

You want to configure an SSH connection to a single Compute Engine instance for users in the dev1 group. This instance is the only resource in this particular Google Cloud Platform project that the dev1 users should be able to connect to. What should you do?

- A. Set metadata to enable-oslogin=true for the instance. Grant the dev1 group the compute.osLogin role. Direct them to use the Cloud Shell to ssh to that instance.
- B. Set metadata to enable-oslogin=true for the instance. Set the service account to no service account for that instance. Direct them to use the Cloud Shell to ssh to that instance.
- C. Enable block project wide keys for the instance. Generate an SSH key for each user in the dev1 group. Distribute the keys to dev1 users and direct them to use their third-party tools to connect.
- D. Enable block project wide keys for the instance. Generate an SSH key and associate the key with that instance. Distribute the key to dev1 users and direct them to use their third-party tools to connect.

Correct Answer: A

Section:

Explanation:

After you enable OS Login on one or more instances in your project, those VMs accept connections only from user accounts that have the necessary IAM roles in your project or organization. In this case, we are granting the group compute.osLogin which lets them log in as non-administrator account. And since we are directing them to use Cloud Shell to ssh, we dont need to add their SSH keys to the instance metadata.

Ref:https://cloud.google.com/compute/docs/instances/managing-instance-access#configure_users Ref:https://cloud.google.com/compute/docs/instances/managing-instance-access#add_oslogin_keys

QUESTION 30

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 .`
- 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` to verify the project value, and then run `gcloud services list --available`.

Correct Answer: A

Section:

Explanation:

`gcloud services list --available` returns not only the enabled services in the project but also services that CAN be enabled.`

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

Run the following command to list the enabled APIs and services in your current project:

`gcloud services list`

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

`gcloud services list --available`

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

`--available`

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

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

`gcloud services list --enabled`

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

gcloud services list --available

QUESTION 31

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:

Explanation:

GCP App Engine natively offers traffic splitting functionality between versions. You can use traffic splitting to specify a percentage distribution of traffic across two or more of the versions within a service. Splitting traffic allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.

Ref:<https://cloud.google.com/appengine/docs/standard/python/splitting-traffic>

QUESTION 32

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: A

Section:

Explanation:

<https://cloud.google.com/compute/docs/disks/local-ssd>

QUESTION 33

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:

Explanation:

Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service. is not right.

Kubernetes Service of type ClusterIP exposes the Service on a cluster-internal IP. Choosing this value makes the Service only reachable from within the cluster so you can not route external traffic to this IP.
Ref:<https://kubernetes.io/docs/concepts/services-networking/service/>

QUESTION 34

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:

Explanation:

Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network. When you use Shared VPC, you designate a project as a host project and attach one or more other service projects to it. The VPC networks in the host project are called Shared VPC networks. Eligible resources from service projects can use subnets in the Shared VPC network

<https://cloud.google.com/vpc/docs/shared-vpc>

'For example, an existing instance in a service project cannot be reconfigured to use a Shared VPC network, but a new instance can be created to use available subnets in a Shared VPC network.'

QUESTION 35

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:

Explanation:

The primitive role roles/viewer provides read access to all resources in the project. The permissions in this role are limited to Get and list access for all resources. As we have an out of the box role that exactly fits our requirement, we should use this.

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

It is advisable to use the existing GCP provided roles over creating custom roles with similar permissions as this becomes a maintenance overhead. If GCP modifies how permissions are handled or adds/removes permissions, the default GCP provided roles are automatically updated by Google whereas if they were custom roles, the responsibility is with us and this adds to the operational overhead and needs to be avoided.

QUESTION 36

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: D

Section:

Explanation:

Ref:<https://cloud.google.com/kubernetes-engine/pricing>

Example:

apiVersion: v1

kind: Pod

metadata:

name: my-gpu-pod

spec:

containers:

name: my-gpu-container

image: nvidia/cuda:10.0-runtime-ubuntu18.04

command: ["/bin/bash"]

resources:

limits:

nvidia.com/gpu: 2

nodeSelector:

cloud.google.com/gke-accelerator: nvidia-tesla-k80 # or nvidia-tesla-p100 or nvidia-tesla-p4 or nvidia-tesla-v100 or nvidia-tesla-t4

QUESTION 37

Your VMs are running in a subnet that has a subnet mask of 255.255.255.240. The current subnet has no more free IP addresses and you require an additional 10 IP addresses for new VMs. The existing and new VMs should all be able to reach each other without additional routes. What should you do?

- A. Use `gcloud` to expand the IP range of the current subnet.
- B. Delete the subnet, and recreate it using a wider range of IP addresses.
- C. Create a new project. Use Shared VPC to share the current network with the new project.
- D. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

`gcloud compute networks subnets expand-ip-range - expand the IP range of a Compute Engine subnetwork gcloud compute networks subnets expand-ip-range NAME --prefix-length=PREFIX_LENGTH [--region=REGION]`

[GLOUD_WIDE_FLAG ...]

QUESTION 38

Your organization uses G Suite for communication and collaboration. All users in your organization have a G Suite account. You want to grant some G Suite users access to your Cloud Platform project. What should you do?

- A. Enable Cloud Identity in the GCP Console for your domain.
- B. Grant them the required IAM roles using their G Suite email address.
- C. Create a CSV sheet with all users' email addresses. Use the `gcloud` command line tool to convert them into Google Cloud Platform accounts.
- D. In the G Suite console, add the users to a special group called `cloud-console-users@yourdomain.com`. Rely on the default behavior of the Cloud Platform to grant users access if they are members of this group.

Correct Answer: B

Section:

Explanation:

Default behavior does not grant access to the 'your GCP Project' Default behavior allow only create billing account and project - When the organization is created, all users in your domain are automatically granted Project Creator and Billing Account Creator IAM roles at the organization level. This enables users in your domain to continue creating projects with no disruption.

QUESTION 39

You have a Google Cloud Platform account with access to both production and development projects. You need to create an automated process to list all compute instances in development and production projects on a daily basis. What should you do?

- A. Create two configurations using gcloud config. Write a script that sets configurations as active, individually. For each configuration, use gcloud compute instances list to get a list of compute resources.
- B. Create two configurations using gsutil config. Write a script that sets configurations as active, individually. For each configuration, use gsutil compute instances list to get a list of compute resources.
- C. Go to Cloud Shell and export this information to Cloud Storage on a daily basis.
- D. Go to GCP Console and export this information to Cloud SQL on a daily basis.

Correct Answer: A

Section:

Explanation:

You can create two configurations -- one for the development project and another for the production project. And you do that by running "gcloud config configurations create" command.

<https://cloud.google.com/sdk/gcloud/reference/config/configurations/create> In your custom script, you can load these configurations one at a time and execute gcloud compute instances list to list Google Compute Engine instances in the project that is active in the gcloud configuration. Ref:<https://cloud.google.com/sdk/gcloud/reference/compute/instances/list> Once you have this information, you can export it in a suitable format to a suitable target e.g. export as CSV or export to Cloud Storage/BigQuery/SQL, etc

QUESTION 40

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:

Explanation:

<https://cloud.google.com/bigquery/external-data-sources>

An external data source is a data source that you can query directly from BigQuery, even though the data is not stored in BigQuery storage.

BigQuery supports the following external data sources:

- Amazon S3
- Azure Storage
- Cloud Bigtable
- Cloud Spanner
- Cloud SQL
- Cloud Storage
- Drive

QUESTION 41

You need to verify that a Google Cloud Platform service account was created at a particular time. What should you do?

- A. Filter the Activity log to view the Configuration category. Filter the Resource type to Service Account.
- B. Filter the Activity log to view the Configuration category. Filter the Resource type to Google Project.
- C. Filter the Activity log to view the Data Access category. Filter the Resource type to Service Account.
- D. Filter the Activity log to view the Data Access category. Filter the Resource type to Google Project.

Correct Answer: A

Section:

Explanation:

<https://developers.google.com/cloud-search/docs/guides/audit-logging-manual>

QUESTION 42

You deployed an LDAP server on Compute Engine that is reachable via TLS through port 636 using UDP. You want to make sure it is reachable by clients over that port. What should you do?

- A. Add the network tag allow-udp-636 to the VM instance running the LDAP server.
- B. Create a route called allow-udp-636 and set the next hop to be the VM instance running the LDAP server.
- C. Add a network tag of your choice to the instance. Create a firewall rule to allow ingress on UDP port 636 for that network tag.
- D. Add a network tag of your choice to the instance running the LDAP server. Create a firewall rule to allow egress on UDP port 636 for that network tag.

Correct Answer: C

Section:

Explanation:

A tag is simply a character string added to a tags field in a resource, such as Compute Engine virtual machine (VM) instances or instance templates. A tag is not a separate resource, so you cannot create it separately. All resources with that string are considered to have that tag. Tags enable you to make firewall rules and routes applicable to specific VM instances.

QUESTION 43

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

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

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/iam/docs/understanding-roles#billing-roles>

QUESTION 44

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: D

Section:

Explanation:

'The Cloud Spanner to Cloud Storage Text template is a batch pipeline that reads in data from a Cloud Spanner table, optionally transforms the data via a JavaScript User Defined Function (UDF) that you provide, and writes it to Cloud Storage as CSV text files.'

<https://cloud.google.com/dataflow/docs/guides/templates/provided-batch#cloudspannertogcstext>

'The Dataflow connector for Cloud Spanner lets you read data from and write data to Cloud Spanner in a Dataflow pipeline'

<https://cloud.google.com/spanner/docs/dataflow-connector>

<https://cloud.google.com/bigquery/external-data-sources>

QUESTION 45

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:

Explanation:

Choosing a region and zone You choose which region or zone hosts your resources, which controls where your data is stored and used. Choosing a region and zone is important for several reasons:

Handling failures

Distribute your resources across multiple zones and regions to tolerate outages. Google designs zones to be independent from each other: a zone usually has power, cooling, networking, and control planes that are isolated from other zones, and most single failure events will affect only a single zone. Thus, if a zone becomes unavailable, you can transfer traffic to another zone in the same region to keep your services running.

Similarly, if a region experiences any disturbances, you should have backup services running in a different region. For more information about distributing your resources and designing a robust system, see

Designing Robust Systems. Decreased network latency To decrease network latency, you might want to choose a region or zone that is close to your point of service.

https://cloud.google.com/compute/docs/regions-zones#choosing_a_region_and_zone

QUESTION 46

A colleague handed over a Google Cloud Platform project for you to maintain. As part of a security checkup, you want to review who has been granted the Project Owner role. What should you do?

- A. In the console, validate which SSH keys have been stored as project-wide keys.
- B. Navigate to Identity-Aware Proxy and check the permissions for these resources.
- C. Enable Audit Logs on the IAM & admin page for all resources, and validate the results.
- D. Use the command `gcloud projects get-iam-policy` to view the current role assignments.

Correct Answer: D

Section:

Explanation:

A simple approach would be to use the command flags available when listing all the IAM policy for a given project. For instance, the following command: `gcloud projects get-iam-policy $PROJECT_ID --flatten='bindings[].members' --format='table(bindings.members)' --filter='bindings.role:roles/owner'` outputs all the users and service accounts associated with the role 'roles/owner' in the project in question. <https://groups.google.com/g/google-cloud-dev/c/Z6sZs7TvygQ?pli=1>

QUESTION 47

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:****Explanation:**

`gcloud compute networks subnets expand-ip-range NAME gcloud compute networks subnets expand-ip-range - expand the IP range of a Compute Engine subnetwork`
<https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

QUESTION 48

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

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

Correct Answer: A**Section:****Explanation:**

Creating and starting a preemptible VM instance This page explains how to create and use a preemptible virtual machine (VM) instance. A preemptible instance is an instance you can create and run at a much lower price than normal instances. However, Compute Engine might terminate (preempt) these instances if it requires access to those resources for other tasks. Preemptible instances will always terminate after 24 hours. To learn more about preemptible instances, read the preemptible instances documentation. Preemptible instances are recommended only for fault-tolerant applications that can withstand instance preemptions. Make sure your application can handle preemptions before you decide to create a preemptible instance. To understand the risks and value of preemptible instances, read the preemptible instances documentation. <https://cloud.google.com/compute/docs/instances/create-start-preemptible-instance>

QUESTION 49

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: A

Section:

Explanation:

Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console: Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements. Implied deny ingress rule. An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them. A higher priority rule might allow incoming access. The default network includes some additional rules that override this one, allowing certain types of incoming connections. https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

QUESTION 50

Your company runs its Linux workloads on Compute Engine instances. Your company will be working with a new operations partner that does not use Google Accounts. You need to grant access to the instances to your operations partner so they can maintain the installed tooling. What should you do?

- A. Enable Cloud IAP for the Compute Engine instances, and add the operations partner as a Cloud IAP Tunnel User.
- B. Tag all the instances with the same network tag. Create a firewall rule in the VPC to grant TCP access on port 22 for traffic from the operations partner to instances with the network tag.
- C. Set up Cloud VPN between your Google Cloud VPC and the internal network of the operations partner.
- D. Ask the operations partner to generate SSH key pairs, and add the public keys to the VM instances.

Correct Answer: D

Section:

Explanation:

IAP controls access to your App Engine apps and Compute Engine VMs running on Google Cloud. It leverages user identity and the context of a request to determine if a user should be allowed access. IAP is a building block toward BeyondCorp, an enterprise security model that enables employees to work from untrusted networks without using a VPN.

By default, IAP uses Google identities and IAM. By leveraging Identity Platform instead, you can authenticate users with a wide range of external identity providers, such as:

Email/password

OAuth (Google, Facebook, Twitter, GitHub, Microsoft, etc.)

SAML

OIDC

Phone number

Custom

Anonymous

This is useful if your application is already using an external authentication system, and migrating your users to Google accounts is impractical.

<https://cloud.google.com/iap/docs/using-tcp-forwarding#grant-permission>

QUESTION 51

You have created a code snippet that should be triggered whenever a new file is uploaded to a Cloud Storage bucket. You want to deploy this code snippet. What should you do?

- A. Use App Engine and configure Cloud Scheduler to trigger the application using Pub/Sub.
- B. Use Cloud Functions and configure the bucket as a trigger resource.
- C. Use Google Kubernetes Engine and configure a CronJob to trigger the application using Pub/Sub.
- D. Use Dataflow as a batch job, and configure the bucket as a data source.

Correct Answer: B

Section:

Explanation:

Google Cloud Storage Triggers

Cloud Functions can respond to change notifications emerging from Google Cloud Storage. These notifications can be configured to trigger in response to various events inside a bucket---object creation, deletion,

archiving and metadata updates.

Note: Cloud Functions can only be triggered by Cloud Storage buckets in the same Google Cloud Platform project.

Event types

Cloud Storage events used by Cloud Functions are based on Cloud Pub/Sub Notifications for Google Cloud Storage and can be configured in a similar way.

Supported trigger type values are:

google.storage.object.finalize

google.storage.object.delete

google.storage.object.archive

google.storage.object.metadataUpdate

Object Finalize

Trigger type value: google.storage.object.finalize

This event is sent when a new object is created (or an existing object is overwritten, and a new generation of that object is created) in the bucket.

https://cloud.google.com/functions/docs/calling/storage#event_types

QUESTION 52

You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost. What should you do?

- A. Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
- B. Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.
- C. Set up a policy that uses Nearline storage for 30 days, then moves the Coldline for one year, and then moves to Archive storage for two years.
- D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

Correct Answer: B

Section:

Explanation:

The key to understand the requirement is : 'The objects are written once and accessed frequently for 30 days'

Standard Storage

Standard Storage is best for data that is frequently accessed ('hot' data) and/or stored for only brief periods of time.

Archive Storage

Archive Storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the 'coldest' storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days. Archive Storage is the best choice for data that you plan to access less than once a year.

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

QUESTION 53

You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data.

- A. You want to make sure you comply with these requirements. What should you do?
- B. Enable the Identity Aware Proxy API on the project.
- C. Scan the bucket using the Data Loss Prevention API.
- D. Allow only a single Service Account access to read the data.
- E. Enable Data Access audit logs for the Cloud Storage API.

Correct Answer: D

Section:

Explanation:

Logged information Within Cloud Audit Logs, there are two types of logs: Admin Activity logs: Entries for operations that modify the configuration or metadata of a project, bucket, or object. Data Access logs:



Entries for operations that modify objects or read a project, bucket, or object. There are several sub-types of data access logs: ADMIN_READ: Entries for operations that read the configuration or metadata of a project, bucket, or object. DATA_READ: Entries for operations that read an object. DATA_WRITE: Entries for operations that create or modify an object. <https://cloud.google.com/storage/docs/audit-logs#types>

QUESTION 54

You are the team lead of a group of 10 developers. You provided each developer with an individual Google Cloud Project that they can use as their personal sandbox to experiment with different Google Cloud solutions. You want to be notified if any of the developers are spending above \$500 per month on their sandbox environment. What should you do?

- A. Create a single budget for all projects and configure budget alerts on this budget.
- B. Create a separate billing account per sandbox project and enable BigQuery billing exports. Create a Data Studio dashboard to plot the spending per billing account.
- C. Create a budget per project and configure budget alerts on all of these budgets.
- D. Create a single billing account for all sandbox projects and enable BigQuery billing exports. Create a Data Studio dashboard to plot the spending per project.

Correct Answer: C

Section:

Explanation:

Set budgets and budget alerts Overview Avoid surprises on your bill by creating Cloud Billing budgets to monitor all of your Google Cloud charges in one place. A budget enables you to track your actual Google Cloud spend against your planned spend. After you've set a budget amount, you set budget alert threshold rules that are used to trigger email notifications. Budget alert emails help you stay informed about how your spend is tracking against your budget. 2. Set budget scope Set the budget Scope and then click Next. In the Projects field, select one or more projects that you want to apply the budget alert to. To apply the budget alert to all the projects in the Cloud Billing account, choose Select all. <https://cloud.google.com/billing/docs/how-to/budgets#budget-scop>

QUESTION 55

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: D

Section:

Explanation:

QUESTION 56

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: C

Section:

Explanation:

QUESTION 57

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: A, B

Section:

Explanation:

(A) Object Lifecycle Management

Delete

The Delete action deletes an object when the object meets all conditions specified in the lifecycle rule.

Exception: In buckets with Object Versioning enabled, deleting the live version of an object causes it to become a noncurrent version, while deleting a noncurrent version deletes that version permanently.

<https://cloud.google.com/storage/docs/lifecycle#delete>

(B) Signed URLs

This page provides an overview of signed URLs, which you use to give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account

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

QUESTION 58

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. Assign the appropriate permissions, and then use Cloud Monitoring to review metrics
- B. Use the export logs API to provide the Admin Activity Audit Logs in the format they want
- C. 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
- D. Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs

Correct Answer: C

Section:

Explanation:

Types of audit logs Cloud Audit Logs provides the following audit logs for each Cloud project, folder, and organization: Admin Activity audit logs Data Access audit logs System Event audit logs Policy Denied audit logs ***Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data.

<https://cloud.google.com/logging/docs/audit#types>

<https://cloud.google.com/logging/docs/audit#data-access> Cloud Storage: When Cloud Storage usage logs are enabled, Cloud Storage writes usage data to the Cloud Storage bucket, which generates Data Access audit logs for the

bucket. The generated Data Access audit log has its caller identity redacted.

QUESTION 59

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:

Explanation:

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0#:~:text=Go%20to%20the%20destination%20project,Voila!>

QUESTION 60

You are working for a hospital that stores its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?

- A. Deploy a Dataflow job from the batch template 'Datastore to Cloud Storage'. Schedule the batch job on the desired interval
- B. In the Cloud Console, go to Cloud Storage. Upload the relevant images to the appropriate bucket
- C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage. Schedule the script as a cron job
- D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub topic

Correct Answer: C

Section:

Explanation:

they require cloud storage for archival and they want to automate the process to upload new medical images to cloud storage, hence we go for gsutil to copy on-prem images to cloud storage and automate the process via cron job. whereas Pub/Sub listens to the changes in the Cloud Storage bucket and triggers the pub/sub topic, which is not required.

QUESTION 61

You have developed a containerized web application that will serve internal colleagues during business hours. You want to ensure that no costs are incurred outside of the hours the application is used. You have just created a new Google Cloud project and want to deploy the application. What should you do?

- A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
- B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.
- C. Deploy the container on App Engine flexible environment with autoscaling, and set the value min_instances to zero in the app.yaml
- D. Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app.yaml

Correct Answer: B

Section:

Explanation:

https://cloud.google.com/kubern/docs/architecture-overview#components_in_the_default_installation

QUESTION 62

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

Correct Answer: B

Section:

Explanation:

You can use Google Cloud Deployment Manager to create a set of Google Cloud resources and manage them as a unit, called a deployment. For example, if your team's development environment needs two virtual machines (VMs) and a BigQuery database, you can define these resources in a configuration file, and use Deployment Manager to create, change, or delete these resources. You can make the configuration file part of your team's code repository, so that anyone can create the same environment with consistent results. <https://cloud.google.com/deployment-manager/docs/quickstart>

QUESTION 63

You are using Data Studio to visualize a table from your data warehouse that is built on top of BigQuery. Data is appended to the data warehouse during the day. At night, the daily summary is recalculated by overwriting the table. You just noticed that the charts in Data Studio are broken, and you want to analyze the problem. What should you do?

- A. Use the BigQuery interface to review the nightly Job and look for any errors
- B. Review the Error Reporting page in the Cloud Console to find any errors.
- C. In Cloud Logging create a filter for your Data Studio report
- D. Use Cloud Debugger to find out why the data was not refreshed correctly

Correct Answer: C

Section:

Explanation:

Cloud Debugger helps inspect the state of an application, at any code location, without stopping or slowing down the running app // <https://cloud.google.com/debugger/docs>

QUESTION 64

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

- A. Save file automatic first-of-the-month backup for three years Store the backup file in an Archive class Cloud Storage bucket
- B. Convert the automatic first-of-the-month backup to an export file Write the export file to a Coldline class Cloud Storage bucket
- C. Set up an export job for the first of the month Write the export file to an Archive class Cloud Storage bucket
- D. Set up an on-demand backup for the first of the month Write the backup to an Archive class Cloud Storage bucket

Correct Answer: C

Section:

Explanation:

https://cloud.google.com/sql/docs/mysql/backup-recovery/backups#can_i_export_a_backup

https://cloud.google.com/sql/docs/mysql/import-export#automating_export_operations

QUESTION 65

You are developing a new web application that will be deployed on Google Cloud Platform. As part of your release cycle, you want to test updates to your application on a small portion of real user traffic. The majority of the users should still be directed towards a stable version of your application. What should you do?



- A. Deploy the application on App Engine For each update, create a new version of the same service Configure traffic splitting to send a small percentage of traffic to the new version
- B. Deploy the application on App Engine For each update, create a new service Configure traffic splitting to send a small percentage of traffic to the new service.
- C. Deploy the application on Kubernetes Engine For a new release, update the deployment to use the new version
- D. Deploy the application on Kubernetes Engine For a new release, create a new deployment for the new version Update the service to use the new deployment.

Correct Answer: A

Section:

Explanation:

Keyword, Version, traffic splitting, App Engine supports traffic splitting for versions before releasing.

QUESTION 66

Your organization has three existing Google Cloud projects. You need to bill the Marketing department for only their Google Cloud services for a new initiative within their group. What should you do?

- A. 1. Verify that you are assigned the Billing Administrator IAM role for your organization's Google Cloud Project for the Marketing department 2. Link the new project to a Marketing Billing Account
- B. 1. Verify that you are assigned the Billing Administrator IAM role for your organization's Google Cloud account 2. Create a new Google Cloud Project for the Marketing department 3. Set the default key-value project labels to department marketing for all services in this project
- C. 1. Verify that you are assigned the Organization Administrator IAM role for your organization's Google Cloud account 2. Create a new Google Cloud Project for the Marketing department 3. Link the new project to a Marketing Billing Account.
- D. 1. Verify that you are assigned the Organization Administrator IAM role for your organization's Google Cloud account 2. Create a new Google Cloud Project for the Marketing department 3. Set the default key value project labels to department marketing for all services in this project

Correct Answer: A

Section:



QUESTION 67

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

- A. Use a custom mode VPC network, configure static routes, and use active/passive routing
- B. Use an automatic mode VPC network, configure static routes, and use active/active routing
- C. Use a custom mode VPC network use Cloud Router border gateway protocol (BGP) routes, and use active/passive routing
- D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Correct Answer: C

Section:

Explanation:

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/best-practices>

QUESTION 68

You need to manage a third-party application that will run on a Compute Engine instance. Other Compute Engine instances are already running with default configuration. Application installation files are hosted on Cloud Storage. You need to access these files from the new instance without allowing other virtual machines (VMs) to access these files. What should you do?

- A. Create the instance with the default Compute Engine service account Grant the service account permissions on Cloud Storage.
- B. Create the instance with the default Compute Engine service account Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

- C. Create a new service account and assign this service account to the new instance Grant the service account permissions on Cloud Storage.
- D. Create a new service account and assign this service account to the new instance Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

Correct Answer: B

Section:

Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-using-and-managing-service-accounts>

If an application uses third-party or custom identities and needs to access a resource, such as a BigQuery dataset or a Cloud Storage bucket, it must perform a transition between principals. Because Google Cloud APIs don't recognize third-party or custom identities, the application can't propagate the end-user's identity to BigQuery or Cloud Storage. Instead, the application has to perform the access by using a different Google identity.

QUESTION 69

You manage three Google Cloud projects with the Cloud Monitoring API enabled. You want to follow Google-recommended practices to visualize CPU and network metrics for all three projects together. What should you do?

- A. 1. Create a Cloud Monitoring Dashboard 2. Collect metrics and publish them into the Pub/Sub topics 3. Add CPU and network Charts (or each of (he three projects
- B. 1. Create a Cloud Monitoring Dashboard. 2. Select the CPU and Network metrics from the three projects. 3. Add CPU and network Charts lot each of the three protects.
- C. 1 Create a Service Account and apply roles/viewer on the three projects 2. Collect metrics and publish them lo the Cloud Monitoring API 3. Add CPU and network Charts for each of the three projects.
- D. 1. Create a fourth Google Cloud project 2 Create a Cloud Workspace from the fourth project and add the other three projects

Correct Answer: B

Section:

QUESTION 70

Your organization uses Active Directory (AD) to manage user identities. Each user uses this identity for federated access to various on-premises systems. Your security team has adopted a policy that requires users to log into Google Cloud with their AD identity instead of their own login. You want to follow the Google-recommended practices to implement this policy. What should you do?

- A. Sync Identities with Cloud Directory Sync, and then enable SAML for single sign-on
- B. Sync Identities in the Google Admin console, and then enable Oauth for single sign-on
- C. Sync identities with 3rd party LDAP sync, and then copy passwords to allow simplified login with (he same credentials
- D. Sync identities with Cloud Directory Sync, and then copy passwords to allow simplified login with the same credentials.

Correct Answer: A

Section:

QUESTION 71

You have deployed multiple Linux instances on Compute Engine. You plan on adding more instances in the coming weeks. You want to be able to access all of these instances through your SSH client over me Internet without having to configure specific access on the existing and new instances. You do not want the Compute Engine instances to have a public IP. What should you do?

- A. Configure Cloud Identity-Aware Proxy (or HTTPS resources
- B. Configure Cloud Identity-Aware Proxy for SSH and TCP resources.
- C. Create an SSH keypair and store the public key as a project-wide SSH Key
- D. Create an SSH keypair and store the private key as a project-wide SSH Key

Correct Answer: B

Section:

Explanation:

<https://cloud.google.com/iap/docs/using-tcp-forwarding>

QUESTION 72

You need to immediately change the storage class of an existing Google Cloud bucket. You need to reduce service cost for infrequently accessed files stored in that bucket and for all files that will be added to that bucket in the future. What should you do?

- A. Use the gsutil to rewrite the storage class for the bucket Change the default storage class for the bucket
- B. Use the gsutil to rewrite the storage class for the bucket Set up Object Lifecycle management on the bucket
- C. Create a new bucket and change the default storage class for the bucket Set up Object Lifecycle management on lite bucket
- D. Create a new bucket and change the default storage class for the bucket import the files from the previous bucket into the new bucket

Correct Answer: B

Section:

QUESTION 73

You have been asked to set up the billing configuration for a new Google Cloud customer. Your customer wants to group resources that share common IAM policies. What should you do?

- A. Use labels to group resources that share common IAM policies
- B. Use folders to group resources that share common IAM policies
- C. Set up a proper billing account structure to group IAM policies
- D. Set up a proper project naming structure to group IAM policies

Correct Answer: B

Section:

Explanation:

Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. Organizations can use folders to group projects under the organization node in a hierarchy. For example, your organization might contain multiple departments, each with its own set of Google Cloud resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain multiple folders or resources, a given folder or resource can have exactly one parent. <https://cloud.google.com/resource-manager/docs/creating-managing-folders>

QUESTION 74

An external member of your team needs list access to compute images and disks in one of your projects. You want to follow Google-recommended practices when you grant the required permissions to this user. What should you do?

- A. Create a custom role, and add all the required compute.disks.list and compute, images.list permissions as includedPermissions. Grant the custom role to the user at the project level.
- B. Create a custom role based on the Compute Image User role Add the compute.disks, list to the includedPermissions field Grant the custom role to the user at the project level
- C. Grant the Compute Storage Admin role at the project level.
- D. Create a custom role based on the Compute Storage Admin role. Exclude unnecessary permissions from the custom role. Grant the custom role to the user at the project level.

Correct Answer: B

Section:

QUESTION 75

Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises workloads consist of:

- * A Flask web application
- * A backend API



* A scheduled long-running background job for ETL and reporting.

You need to keep operational costs low. You want to follow Google-recommended practices to migrate these workloads to serverless solutions on Google Cloud. What should you do?

- A. Migrate the web application to App Engine and the backend API to Cloud Run. Use Cloud Tasks to run your background job on Compute Engine.
- B. Migrate the web application to App Engine and the backend API to Cloud Run. Use Cloud Tasks to run your background job on Cloud Run.
- C. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run. Use Cloud Tasks to run your background job on Cloud Run.
- D. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run. Use Cloud Tasks to run your background job on Compute Engine.

Correct Answer: B

Section:

QUESTION 76

You are building a data lake on Google Cloud for your Internet of Things (IoT) application. The IoT application has millions of sensors that are constantly streaming structured and unstructured data to your backend in the cloud. You want to build a highly available and resilient architecture based on Google-recommended practices. What should you do?

- A. Stream data to Pub/Sub, and use Dataflow to send data to Cloud Storage.
- B. Stream data to Pub/Sub, and use Storage Transfer Service to send data to BigQuery.
- C. Stream data to Dataflow, and use Storage Transfer Service to send data to BigQuery.
- D. Stream data to Dataflow, and use Dataprep by Trifacta to send data to Bigtable.

Correct Answer: A

Section:

QUESTION 77

You installed the Google Cloud CLI on your workstation and set the proxy configuration. However, you are worried that your proxy credentials will be recorded in the gcloud CLI logs. You want to prevent your proxy credentials from being logged. What should you do?

- A. Configure username and password by using `gcloud configure set proxy/username` and `gcloud configure set proxy/proxy/password` commands.
- B. Encode username and password in sha256 encoding, and save it to a text file. Use filename as a value in the `gcloud configure set core/custom_ca_certs_file` command.
- C. Provide values for `CLOUDSDK_USERNAME` and `CLOUDSDK_PASSWORD` in the gcloud CLI tool configure file.
- D. Set the `CLOUDSDK_PROXY_USERNAME` and `CLOUDSDK_PROXY_PASSWORD` properties by using environment variables in your command line tool.

Correct Answer: D

Section:

