

Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer



NEW QUESTION 1

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.

Answer: A

NEW QUESTION 2

You will have several applications running on different Compute Engine instances in the same project. You want to specify at a more granular level the service account each instance uses when calling Google Cloud APIs. What should you do?

- A. When creating the instances, specify a Service Account for each instance
- B. When creating the instances, assign the name of each Service Account as instance metadata
- C. After starting the instances, use `gcloud compute instances update` to specify a Service Account for each instance
- D. After starting the instances, use `gcloud compute instances update` to assign the name of the relevant Service Account as instance metadata

Answer: A

Explanation:

https://cloud.google.com/compute/docs/access/service-accounts#associating_a_service_account_to_an_instance

NEW QUESTION 3

Your learn wants to deploy a specific content management system (CMS) solution to Google Cloud. You need a quick and easy way to deploy and install the solution. What should you do?

- A. Search for the CMS solution in Google Cloud Marketplac
- B. Use `gcloud CLI` to deploy the solution.
- C. Search for the CMS solution in Google Cloud Marketplac
- D. Deploy the solution directly from Cloud Marketplace.
- E. Search for the CMS solution in Google Cloud Marketplac
- F. Use Terraform and the Cloud Marketplace ID to deploy the solution with the appropriate parameters.
- G. Use the installation guide of the CMS provide
- H. Perform the installation through your configuration management system.

Answer: B

NEW QUESTION 4

You are designing an application that lets users upload and share photos. You expect your application to grow really fast and you are targeting a worldwide audience. You want to delete uploaded photos after 30 days. You want to minimize costs while ensuring your application is highly available. Which GCP storage solution should you choose?

- A. Persistent SSD on VM instances.
- B. Cloud Filestore.
- C. Multiregional Cloud Storage bucket.
- D. Cloud Datastore database.

Answer: C

Explanation:

Cloud Storage allows world-wide storage and retrieval of any amount of data at any time. We dont need to set up auto-scaling ourselves. Cloud Storage autoscaling is managed by GCP. Cloud Storage is an object store so it is suitable for storing photos. Cloud Storage allows world-wide storage and retrieval so cater well to our worldwide audience. Cloud storage provides us lifecycle rules that can be configured to automatically delete objects older than 30 days. This also fits our requirements. Finally, Google Cloud Storage offers several storage classes such as Nearline Storage (\$0.01 per GB per Month) Coldline Storage (\$0.007 per GB per Month) and Archive Storage (\$0.004 per GB per month) which are significantly cheaper than any of the options above.

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

Ref: <https://cloud.google.com/storage/pricing>



NEW QUESTION 5

You have one project called proj-sa where you manage all your service accounts. You want to be able to use a service account from this project to take snapshots of VMs running in another project called proj-vm. What should you do?

- A. Download the private key from the service account, and add it to each VMs custom metadata.
- B. Download the private key from the service account, and add the private key to each VM's SSH keys.
- C. Grant the service account the IAM Role of Compute Storage Admin in the project called proj-vm.
- D. When creating the VMs, set the service account's API scope for Compute Engine to read/write.

Answer: C

Explanation:

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0>

You create the service account in proj-sa and take note of the service account email, then you go to proj-vm in IAM > ADD and add the service account's email as new member and give it the Compute Storage Admin role.

<https://cloud.google.com/compute/docs/access/iam#compute.storageAdmin>

NEW QUESTION 6

You need to create a custom VPC with a single subnet. The subnet's range must be as large as possible. Which range should you use?

- A. 1.00.0.0/0
- B. 10.0.0.0/8
- C. 172.16.0.0/12
- D. 192.168.0.0/16

Answer: B

Explanation:

https://cloud.google.com/vpc/docs/vpc#manually_created_subnet_ip_ranges

NEW QUESTION 7

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

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

Answer: A

Explanation:

Installing Jenkins

In this section, you use Cloud Marketplace to provision a Jenkins instance. You customize this instance to use the agent image you created in the previous section.

Go to the Cloud Marketplace solution for Jenkins. Click Launch on Compute Engine.

Change the Machine Type field to 4 vCPUs 15 GB Memory, n1-standard-4.

Machine type selection for Jenkins deployment.

Click Deploy and wait for your Jenkins instance to finish being provisioned. When it is finished, you will see: Jenkins has been deployed.

https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine#installing_jenkins

NEW QUESTION 8

You want to send and consume Cloud Pub/Sub messages from your App Engine application. The Cloud Pub/Sub API is currently disabled. You will use a service

account to authenticate your application to the API. You want to make sure your application can use Cloud Pub/Sub. What should you do?

- A. Enable the Cloud Pub/Sub API in the API Library on the GCP Console.
- B. Rely on the automatic enablement of the Cloud Pub/Sub API when the Service Account accesses it.
- C. Use Deployment Manager to deploy your application
- D. Rely on the automatic enablement of all APIs used by the application being deployed.
- E. Grant the App Engine Default service account the role of Cloud Pub/Sub Admin
- F. Have your application enable the API on the first connection to Cloud Pub/Sub.

Answer: A

Explanation:

Quickstart: using the Google Cloud Console

This page shows you how to perform basic tasks in Pub/Sub using the Google Cloud Console. Note: If you are new to Pub/Sub, we recommend that you start with the interactive tutorial. Before you begin

Set up a Cloud Console project. Set up a project

Click to:

Create or select a project.

Enable the Pub/Sub API for that project.

You can view and manage these resources at any time in the Cloud Console. Install and initialize the Cloud SDK.

Note: You can run the gcloud tool in the Cloud Console without installing the Cloud SDK. To run the gcloud tool in the Cloud Console, use Cloud Shell .

<https://cloud.google.com/pubsub/docs/quickstart-console>

NEW QUESTION 9

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

Answer: B

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>

NEW QUESTION 10

Your company is moving its continuous integration and delivery (CI/CD) pipeline to Compute Engine instances. The pipeline will manage the entire cloud infrastructure through code. How can you ensure that the pipeline has appropriate permissions while your system is following security best practices?

- A. • Add a step for human approval to the CI/CD pipeline before the execution of the infrastructure provisioning. • Use the human approval IAM account for the provisioning.
- B. • Attach a single service account to the compute instances. • Add minimal rights to the service account. • Allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources.
- C. • Attach a single service account to the compute instances. • Add all required Identity and Access Management (IAM) permissions to this service account to create, update, or delete resources
- D. • Create multiple service accounts, one for each pipeline with the appropriate minimal Identity and Access Management (IAM) permissions. • Use a secret manager service to store the key files of the service accounts. • Allow the CI/CD pipeline to request the appropriate secrets during the execution of the pipeline.

Answer: B

Explanation:

The best option is to attach a single service account to the compute instances and add minimal rights to the service account. Then, allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources. This way, the service account can use short-lived access tokens to authenticate to Google Cloud APIs without needing to manage service account keys. This option follows the principle of least privilege and reduces the risk of credential leakage and misuse.

Option A is not recommended because it requires human intervention, which can slow down the CI/CD pipeline and introduce human errors. Option C is not secure because it grants all required IAM permissions to a single service account, which can increase the impact of a compromised key. Option D is not cost-effective because it requires creating and managing multiple service accounts and keys, as well as using a secret manager service.

References:

- > 1: <https://cloud.google.com/iam/docs/impersonating-service-accounts>
- > 2: <https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys>
- > 3: <https://cloud.google.com/iam/docs/understanding-service-accounts>

NEW QUESTION 10

You have developed an application that consists of multiple microservices, with each microservice packaged in its own Docker container image. You want to deploy the entire application on Google Kubernetes Engine so that each microservice can be scaled individually. What should you do?

- A. Create and deploy a Custom Resource Definition per microservice.
- B. Create and deploy a Docker Compose File.
- C. Create and deploy a Job per microservice.
- D. Create and deploy a Deployment per microservice.

Answer: A

NEW QUESTION 13

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 17

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/Su
- C. and use Storage Transfer Service to send data to BigQuery.
- D. Stream data to Dataflow, and use Storage Transfer Service to send data to BigQuery.
- E. Stream data to Dataflow, and use Dataprep by Trifacta to send data to Bigtable.

Answer: B

NEW QUESTION 18

You are managing a Data Warehouse on BigQuery. An external auditor will review your company's processes, and multiple external consultants will need view access to the data. You need to provide them with view access while following Google-recommended practices. What should you do?

- A. Grant each individual external consultant the role of BigQuery Editor
- B. Grant each individual external consultant the role of BigQuery Viewer
- C. Create a Google Group that contains the consultants and grant the group the role of BigQuery Editor
- D. Create a Google Group that contains the consultants, and grant the group the role of BigQuery Viewer

Answer: D

NEW QUESTION 20

You have an object in a Cloud Storage bucket that you want to share with an external company. The object contains sensitive data. You want access to the content to be removed after four hours. The external company does not have a Google account to which you can grant specific user-based access privileges. You want to use the most secure method that requires the fewest steps. What should you do?

- A. Create a signed URL with a four-hour expiration and share the URL with the company.
- B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.
- C. Configure the storage bucket as a static website and furnish the object's URL to the compan
- D. Delete the object from the storage bucket after four hours.
- E. Create a new Cloud Storage bucket specifically for the external company to acces
- F. Copy the object to that bucke
- G. Delete the bucket after four hours have passed.

Answer: A

Explanation:

Signed URLs are used 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>

NEW QUESTION 23

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.

Answer: A

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

NEW QUESTION 28

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

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

Answer: A

Explanation:

auto-provisioning = Attaches and deletes node pools to cluster based on the requirements. Hence creating a GPU node pool, and auto-scaling would be better
<https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-provisioning>

NEW QUESTION 31

You have been asked to migrate a docker application from datacenter to cloud. Your solution architect has suggested uploading docker images to GCR in one project and running an application in a GKE cluster in a separate project. You want to store images in the project img-278322 and run the application in the project prod-278986. You want to tag the image as acme_track_n_trace:v1. You want to follow Google-recommended practices. What should you do?

- A. Run `gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace`
- B. Run `gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace:v1`
- C. Run `gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace`
- D. Run `gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace:v1`

Answer: B

Explanation:

➤ Run `gcloud builds submit tag gcr.io/img-278322/acme_track_n_trace:v1`. is the right answer.
 This command correctly tags the image as acme_track_n_trace:v1 and uploads the image to the img-278322 project.
 Ref: <https://cloud.google.com/sdk/gcloud/reference/builds/submit>

NEW QUESTION 34

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

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

Answer: B

Explanation:

<https://cloud.google.com/kuberun/docs/rollouts-rollbacks-traffic-migration>

NEW QUESTION 36

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

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

Answer: B

Explanation:

Best practices for persistent disk snapshots
 You can create persistent disk snapshots at any time, but you can create snapshots more quickly and with greater reliability if you use the following best practices.
 Creating frequent snapshots efficiently
 Use snapshots to manage your data efficiently.
 Create a snapshot of your data on a regular schedule to minimize data loss due to unexpected failure. Improve performance by eliminating excessive snapshot downloads and by creating an image and reusing it. Set your snapshot schedule to off-peak hours to reduce snapshot time.

Snapshot frequency limits

Creating snapshots from persistent disks

You can snapshot your disks at most once every 10 minutes. If you want to issue a burst of requests to snapshot your disks, you can issue at most 6 requests in 60 minutes.

If the limit is exceeded, the operation fails and returns the following error: <https://cloud.google.com/compute/docs/disks/snapshot-best-practices>

NEW QUESTION 41

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

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

Answer: C

Explanation:

https://cloud.google.com/iam/docs/understanding-service-accounts#using_service_accounts_with_compute_eng <https://cloud.google.com/storage/docs/access-control/iam-roles>

NEW QUESTION 42

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

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

Answer: D

Explanation:

cell phones are sending UDP packets and the only that can receive that type of traffic is a External Network TCP/UDP <https://cloud.google.com/load-balancing/docs/network>

<https://cloud.google.com/load-balancing/docs/choosing-load-balancer#lb-decision-tree>

NEW QUESTION 47

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 templat
- D. Add the group to the backend service for the load balance
- E. When all instances in the new managed instance group are healthy, delete the old managed instance group.
- F. Create a new instance template with the new application versio
- G. Update the existing managed instance group with the new instance templat
- H. Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

Answer: B

Explanation:

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

NEW QUESTION 51

You are the project owner of a GCP project and want to delegate control to colleagues to manage buckets and files in Cloud Storage. You want to follow Google-recommended practices. Which IAM roles should you grant your colleagues?

- A. Project Editor
- B. Storage Admin
- C. Storage Object Admin
- D. Storage Object Creator

Answer: B

Explanation:

Storage Admin (roles/storage.admin) Grants full control of buckets and objects.

When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.

`firebase.projects.get resourceManager.projects.get`

`resourceManager.projects.list storage.buckets.* storage.objects.*`

<https://cloud.google.com/storage/docs/access-control/iam-roles>

This role grants full control of buckets and objects. When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.

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

NEW QUESTION 54

The sales team has a project named Sales Data Digest that has the ID acme-data-digest. You need to set up similar Google Cloud resources for the marketing team but their resources must be organized independently of the sales team. What should you do?

- A. Grant the Project Editor role to the Marketing team for acme data digest
- B. Create a Project Lien on acme-data digest and then grant the Project Editor role to the Marketing team
- C. Create another project with the ID acme-marketing-data-digest for the Marketing team and deploy the resources there
- D. Create a new project named Meeting Data Digest and use the ID acme-data-digest. Grant the Project Editor role to the Marketing team.

Answer: C

NEW QUESTION 55

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 the open source CLI tool
- E. Snapshot Debugger, to find out why the data was not refreshed correctly.

Answer: D

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>

NEW QUESTION 57

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

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

Answer: A

Explanation:

VMs cannot communicate over the internet without a public IP address. Private Google Access permits access to Google APIs and services in Google's production infrastructure. <https://cloud.google.com/vpc/docs/private-google-access>

NEW QUESTION 58

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?

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

Answer: A

Explanation:

<https://cloud.google.com/logging/docs/api/v2/resource-list> GKE Containers have more log than GKE Cluster Operations:

-GKE Container:

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.

NEW QUESTION 60

You need to manage a Cloud Spanner Instance for best query performance. Your instance in production runs in a single Google Cloud region. You need to improve performance in the shortest amount of time. You want to follow Google best practices for service configuration. What should you do?

- A. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45%. If you exceed this threshold, add nodes to your instance.
- B. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45%. Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.
- C. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%. If you exceed this threshold, add nodes to your instance.
- D. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%. Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.

Answer: B

Explanation:

<https://cloud.google.com/spanner/docs/cpu-utilization#recommended-max>

NEW QUESTION 65

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.

Answer: D

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>

NEW QUESTION 69

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

- A. Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.
- B. Use the command line to run a dry run query to estimate the number of bytes read.
- C. Then convert that bytes estimate to dollars using the Pricing Calculator.
- D. Use the command line to run a dry run query to estimate the number of bytes returned.
- E. Then convert that bytes estimate to dollars using the Pricing Calculator.
- F. Run a select count (*) to get an idea of how many records your query will look through.
- G. Then convert that number of rows to dollars using the Pricing Calculator.

Answer: B

NEW QUESTION 73

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

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

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

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

Answer: B

Explanation:

<https://cloud.google.com/config-connector/docs/how-to/secrets#gcloud>

NEW QUESTION 76

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

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

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset> https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns
DaemonSets attempt to adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed.

In GKE, DaemonSets manage groups of replicated Pods and adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. So, this is a perfect fit for our monitoring pod.

Ref: <https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset>

DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluentd, and node monitoring daemons like collectd. For example, you could have DaemonSets for each type of daemon run on all of your nodes. Alternatively, you could run multiple DaemonSets for a single type of daemon, but have them use different configurations for different hardware types and resource needs.

NEW QUESTION 79

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.

Answer: B

NEW QUESTION 83

You need to create a copy of a custom Compute Engine virtual machine (VM) to facilitate an expected increase in application traffic due to a business acquisition. What should you do?

- A. Create a Compute Engine snapshot of your base V
- B. Create your images from that snapshot.
- C. Create a Compute Engine snapshot of your base V
- D. Create your instances from that snapshot.
- E. Create a custom Compute Engine image from a snapsho
- F. Create your images from that image.
- G. Create a custom Compute Engine image from a snapsho
- H. Create your instances from that image.

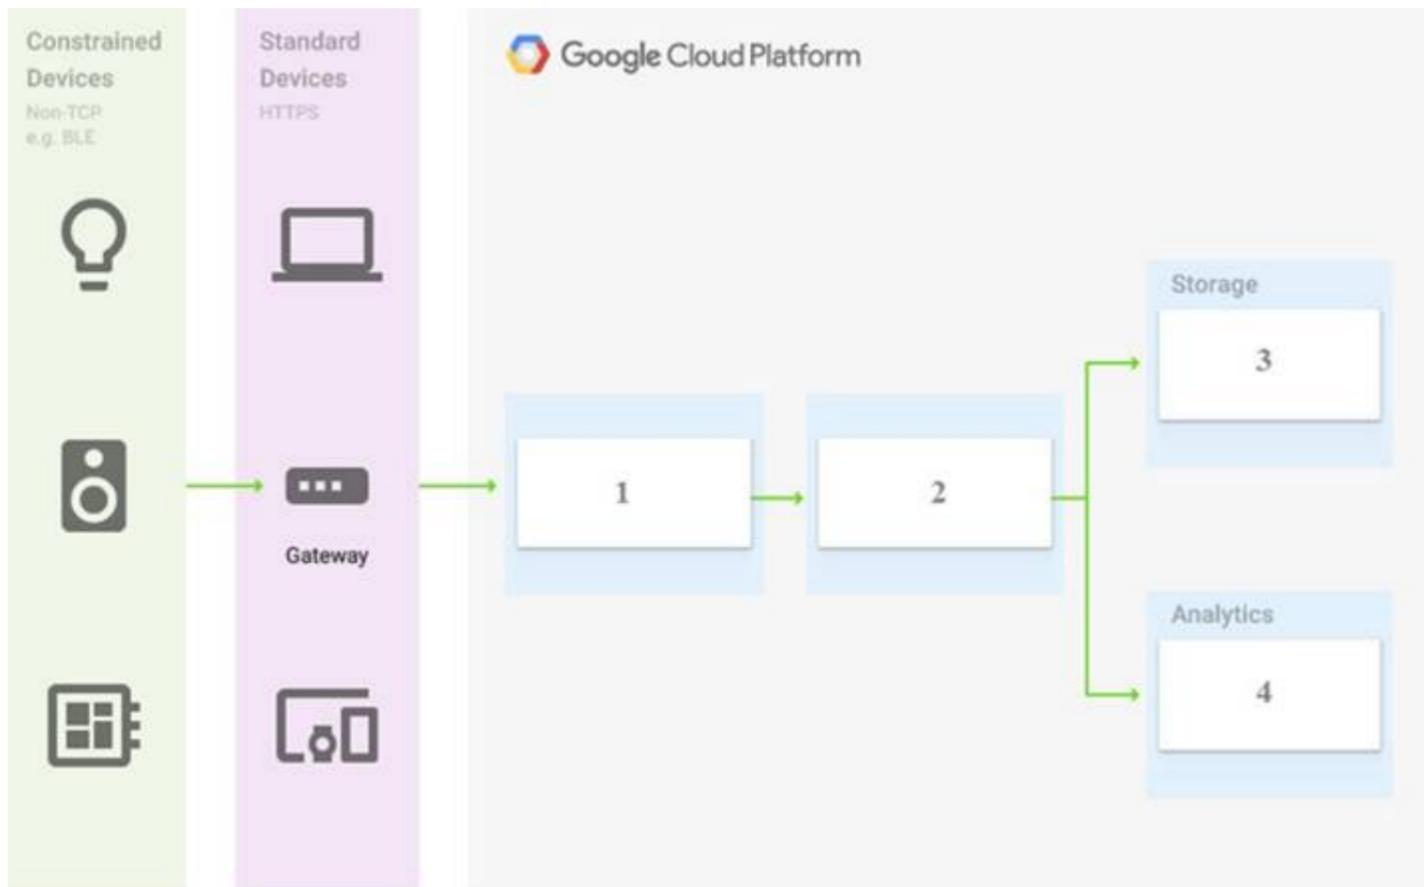
Answer: D

Explanation:

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

NEW QUESTION 85

You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?



- A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
- B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery
- C. Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable
- D. Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

Answer: D

NEW QUESTION 88

You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your organizational policies require you to use virtual machines directly. You need to ensure that the application scaling is operationally efficient and completed as quickly as possible. What should you do?

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

Answer: B

Explanation:

Managed instance groups offer autoscaling capabilities that let you automatically add or delete instances from a managed instance group based on increases or decreases in load (CPU Utilization in this case). Autoscaling helps your apps gracefully handle increases in traffic and reduce costs when the need for resources is lower. You define the autoscaling policy and the autoscaler performs automatic scaling based on the measured load (CPU Utilization in this case). Autoscaling works by adding more instances to your instance group when there is more load (upscaling), and deleting instances when the need for instances is lowered (downscaling). Ref: <https://cloud.google.com/compute/docs/autoscaler>

NEW QUESTION 91

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

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

Answer: B

Explanation:

Creating or upgrading a cluster by specifying the version as latest does not provide automatic upgrades. Enable node auto-upgrades to ensure that the nodes in your cluster are up-to-date with the latest stable version.

<https://cloud.google.com/kubernetes-engine/versioning-and-upgrades>

Node auto-upgrades help you keep the nodes in your cluster up to date with the cluster master version when your master is updated on your behalf. When you create a new cluster or node pool with Google Cloud Console or the gcloud command, node auto-upgrade is enabled by default.

Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades>

NEW QUESTION 96

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.
- E. Use the credentials in the JSON file to log in to the VM.

Answer: B

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#gc>

NEW QUESTION 99

You have deployed an application on a Compute Engine instance. An external consultant needs to access the Linux-based instance. The consultant is connected to your corporate network through a VPN connection, but the consultant has no Google account. What should you do?

- A. Instruct the external consultant to use the gcloud compute ssh command line tool by using Identity-Aware Proxy to access the instance.
- B. Instruct the external consultant to use the gcloud compute ssh command line tool by using the public IP address of the instance to access it.
- C. Instruct the external consultant to generate an SSH key pair, and request the public key from the consultant. Add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key.
- D. Instruct the external consultant to generate an SSH key pair, and request the private key from the consultant. Add the private key to the instance yourself, and have the consultant access the instance through SSH with their public key.

Answer: C

Explanation:

The best option is to instruct the external consultant to generate an SSH key pair, and request the public key from the consultant. Then, add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key. This way, you can grant the consultant access to the instance without requiring a Google account or exposing the instance's public IP address. This option also follows the best practice of using user-managed SSH keys instead of service account keys for SSH access¹.

Option A is not feasible because the external consultant does not have a Google account, and therefore cannot use Identity-Aware Proxy (IAP) to access the instance. IAP requires the user to authenticate with a Google account and have the appropriate IAM permissions to access the instance². Option B is not secure because it exposes the instance's public IP address, which can increase the risk of unauthorized access or attacks. Option D is not correct because it reverses the roles of the public and private keys. The public key should be added to the instance, and the private key should be kept by the consultant. Sharing the private key with anyone else can compromise the security of the SSH connection³.

References:

- > 1: <https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys>
- > 2: <https://cloud.google.com/iap/docs/using-tcp-forwarding>
- > 3: <https://cloud.google.com/compute/docs/instances/connecting-advanced#sshbetweeninstances>

NEW QUESTION 103

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

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

Answer: A

Explanation:

When setting support levels for permissions in custom roles, you can set to one of SUPPORTED, TESTING or NOT_SUPPORTED.

Ref: <https://cloud.google.com/iam/docs/custom-roles-permissions-support>

NEW QUESTION 106

An employee was terminated, but their access to Google Cloud Platform (GCP) was not removed until 2 weeks later. You need to find out this employee accessed any sensitive customer information after their termination. What should you do?

- A. View System Event Logs in Stackdrive
- B. Search for the user's email as the principal.
- C. View System Event Logs in Stackdrive
- D. Search for the service account associated with the user.
- E. View Data Access audit logs in Stackdrive
- F. Search for the user's email as the principal.
- G. View the Admin Activity log in Stackdrive
- H. Search for the service account associated with the user.

Answer: C

Explanation:

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

Data Access 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#data-access>

NEW QUESTION 111

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

- A. Create a file in Cloud Storage per device and append new data to that file.
- B. Create a file in Cloud Filestore per device and append new data to that file.
- C. Ingest the data into Datastor
- D. Store data in an entity group based on the device.
- E. Ingest the data into Cloud Bigtabl
- F. Create a row key based on the event timestamp.

Answer: D

Explanation:

Keyword need to look for

- "High Throughput",
 - "Consistent",
 - "Property based data insert/fetch like engine status, distance traveled, fuel level, and more." which can be designed in column,
 - "Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds" This will go for pera bytes situation,
 - Export data based on the time of the event.
 - Atomic
- o BigTable will fit all requirement. o DataStore is not fully Atomic
o CloudStorage is not a option where we can export data based on time of event. We need another solution to do that
o FireStore can be used with MobileSDK.

NEW QUESTION 116

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

- A. Use kubectl app deploy <dockerfilename>.
- B. Use gcloud app deploy <dockerfilename>.
- C. Create a docker image from the Dockerfile and upload it to Container Registr
- D. Create a Deployment YAML file to point to that imag
- E. Use kubectl to create the deployment with that file.
- F. Create a docker image from the Dockerfile and upload it to Cloud Storag
- G. Create a Deployment YAML file to point to that imag
- H. Use kubectl to create the deployment with that file.

Answer: C

NEW QUESTION 119

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

Answer: B

Explanation:

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

NEW QUESTION 120

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

Answer: A

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

NEW QUESTION 124

Your company completed the acquisition of a startup and is now merging the IT systems of both companies. The startup had a production Google Cloud project in their organization. You need to move this project into your organization and ensure that the project is billed to your organization. You want to accomplish this task with minimal effort. What should you do?

- A. Use the project
- B. move method to move the project to your organization
- C. Update the billing account of the project to that of your organization.
- D. Ensure that you have an Organization Administrator Identity and Access Management (IAM) role assigned to you in both organization
- E. Navigate to the Resource Manager in the startup's Google Cloud organization, and drag the project to your company's organization.
- F. Create a Private Catalog for the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalog
- G. Share the Catalog with your organization, and deploy the resources in your company's project.
- H. Create an infrastructure-as-code template for all resources in the project by using Terraform
- I. and deploy that template to a new project in your organization
- J. Delete the project from the startup's Google Cloud organization.

Answer: A

NEW QUESTION 128

You are configuring Cloud DNS. You want to create DNS records to point home.mydomain.com, mydomain.com. and www.mydomain.com to the IP address of your Google Cloud load balancer. What should you do?

- A. Create one CNAME record to point mydomain.com to the load balancer, and create two A records to point WWW and HOME to mydomain.com respectively.
- B. Create one CNAME record to point mydomain.com to the load balancer, and create two AAAA records to point WWW and HOME to mydomain.com respectively.
- C. Create one A record to point mydomain.com to the load balancer, and create two CNAME records to point WWW and HOME to mydomain.com respectively.
- D. Create one A record to point mydomain.com to the load balancer, and create two NS records to point WWW and HOME to mydomain.com respectively.

Answer: C

NEW QUESTION 133

You have a Compute Engine instance hosting an application used between 9 AM and 6 PM on weekdays. You want to back up this instance daily for disaster recovery purposes. You want to keep the backups for 30 days. You want the Google-recommended solution with the least management overhead and the least number of services. What should you do?

- A. * 1. Update your instances' metadata to add the following value: snapshot-schedule: 0 1 * * * * 2. Update your instances' metadata to add the following value: snapshot-retention: 30
- B. * 1. In the Cloud Console, go to the Compute Engine Disks page and select your instance's disk.* 2. In the Snapshot Schedule section, select Create Schedule and configure the following parameters:–Schedule frequency: Daily–Start time: 1:00 AM – 2:00 AM–Autodelete snapshots after 30 days
- C. * 1. Create a Cloud Function that creates a snapshot of your instance's disk.* 2.Create a Cloud Function that deletes snapshots that are older than 30 days
- D. 3.Use Cloud Scheduler to trigger both Cloud Functions daily at 1:00 AM.
- E. * 1. Create a bash script in the instance that copies the content of the disk to Cloud Storage.* 2. Create a bash script in the instance that deletes data older than 30 days in the backup Cloud Storage bucket.* 3. Configure the instance's crontab to execute these scripts daily at 1:00 AM.

Answer: B

Explanation:

Creating scheduled snapshots for persistent disk This document describes how to create a snapshot schedule to regularly and automatically back up your zonal and regional persistent disks. Use snapshot schedules as a best practice to back up your Compute Engine workloads. After creating a snapshot schedule, you can apply it to one or more persistent disks. <https://cloud.google.com/compute/docs/disks/scheduled-snapshots>

NEW QUESTION 135

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
- D. and set the value min_instances to zero in the app.yaml
- E. Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app.yaml

Answer: B

Explanation:

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

NEW QUESTION 139

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.

Answer: C

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."

NEW QUESTION 143

You need to configure optimal data storage for files stored in Cloud Storage for minimal cost. The files are used in a mission-critical analytics pipeline that is used continually. The users are in Boston, MA (United States). What should you do?

- A. Configure regional storage for the region closest to the users Configure a Nearline storage class
- B. Configure regional storage for the region closest to the users Configure a Standard storage class
- C. Configure dual-regional storage for the dual region closest to the users Configure a Nearline storage class
- D. Configure dual-regional storage for the dual region closest to the users Configure a Standard storage class

Answer: B

Explanation:

Keywords: - continually -> Standard - mission-critical analytics -> dual-regional

NEW QUESTION 147

You have an application running in Google Kubernetes Engine (GKE) with cluster autoscaling enabled. The application exposes a TCP endpoint. There are several replicas of this application. You have a Compute Engine instance in the same region, but in another Virtual Private Cloud (VPC), called `gce-network`, that has no overlapping IP ranges with the first VPC. This instance needs to connect to the application on GKE. You want to minimize effort. What should you do?

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

Answer: A

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>

NEW QUESTION 148

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

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

Answer: C

NEW QUESTION 153

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
- D. Then configure App Engine using GCP Console to split traffic between the two apps.
- E. Deploy a new version of your application in App Engine
- F. Then go to App Engine settings in GCP Console and split traffic between the current version and newly deployed versions accordingly.

Answer: D

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>

NEW QUESTION 154

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

Answer: A

NEW QUESTION 158

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

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

Answer: D

Explanation:

Install the workload in a compute engine VM, start and stop the instance as needed, because as per the question the VM runs for 30 hours, process can be performed offline and should not be interrupted, if interrupted we need to restart the batch process again. Preemptible VMs are cheaper, but they will not be available beyond 24hrs, and if the process gets interrupted the preemptible VM will restart.

NEW QUESTION 162

You created a Kubernetes deployment by running `kubectl run nginx --image=nginx --replicas=1`. After a few days, you decided you no longer want this deployment. You identified the pod and deleted it by running `kubectl delete pod`. You noticed the pod got recreated.

```
> $ kubectl get pods
NAME READY STATUS RESTARTS AGE
nginx-84748895c4-nqqmt 1/1 Running 0 9m41s
> $ kubectl delete pod nginx-84748895c4-nqqmt
pod nginx-84748895c4-nqqmt deleted
> $ kubectl get pods
NAME READY STATUS RESTARTS AGE
nginx-84748895c4-k6bz1 1/1 Running 0 25s
```

What should you do to delete the deployment and avoid pod getting recreated?

- A. `kubectl delete deployment nginx`
- B. `kubectl delete --deployment=nginx`
- C. `kubectl delete pod nginx-84748895c4-k6bz1 --no-restart 2`
- D. `kubectl delete nginx`

Answer: A

Explanation:

This command correctly deletes the deployment. Pods are managed by Kubernetes workloads (deployments). When a pod is deleted, the deployment detects the pod is unavailable and brings up another pod to maintain the replica count. The only way to delete the workload is by deleting the deployment itself using the `kubectl delete deployment` command.

```
> $ kubectl delete deployment nginx
deployment.apps/nginx deleted
```

Ref: <https://kubernetes.io/docs/reference/kubectl/cheatsheet/#deleting-resources>

NEW QUESTION 165

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.logs WHERE timestamp >`

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

Answer: C

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.

NEW QUESTION 169

You have two subnets (subnet-a and subnet-b) in the default VPC. Your database servers are running in subnet-a. Your application servers and web servers are running in subnet-b. You want to configure a firewall rule that only allows database traffic from the application servers to the database servers. What should you do?

- A. * Create service accounts sa-app and sa-db. • Associate service account: sa-app with the application servers and the service account sa-db with the database servers. • Create an ingress firewall rule to allow network traffic from source service account sa-app to target service account sa-db.
- B. • Create network tags app-server and db-server. • Add the app-server tag to the application servers and the db-server tag to the database servers. • Create an egress firewall rule to allow network traffic from source network tag app-server to target network tag db-server.
- C. * Create a service account sa-app and a network tag db-server. • Associate the service account sa-app with the application servers and the network tag db-server with the database servers. • Create an ingress firewall rule to allow network traffic from source VPC IP addresses and target the subnet-a IP addresses.
- D. • Create a network tag app-server and service account sa-db. • Add the tag to the application servers and associate the service account with the database servers. • Create an egress firewall rule to allow network traffic from source network tag app-server to target service account sa-db.

Answer: C

NEW QUESTION 170

You have designed a solution on Google Cloud Platform (GCP) that uses multiple GCP products. Your company has asked you to estimate the costs of the solution. You need to provide estimates for the monthly total cost. What should you do?

- A. For each GCP product in the solution, review the pricing details on the products pricing page
- B. Use the pricing calculator to total the monthly costs for each GCP product.
- C. For each GCP product in the solution, review the pricing details on the products pricing page
- D. Create a Google Sheet that summarizes the expected monthly costs for each product.
- E. Provision the solution on GC
- F. Leave the solution provisioned for 1 week
- G. Navigate to the Billing Report page in the Google Cloud Platform Console
- H. Multiply the 1 week cost to determine the monthly costs.
- I. Provision the solution on GC
- J. Leave the solution provisioned for 1 week
- K. Use Stackdriver to determine the provisioned and used resource amount
- L. Multiply the 1 week cost to determine the monthly costs.

Answer: A

Explanation:

You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You don't incur any charges for doing so.
Ref: <https://cloud.google.com/products/calculator>

NEW QUESTION 171

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

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

Answer: D

Explanation:

"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "
<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>
<https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4>

NEW QUESTION 174

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

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

Answer: A

Explanation:

<https://cloud.google.com/run/docs/tutorials/gcloud> <https://cloud.google.com/resource-manager/docs/creating-managing-projects>

https://cloud.google.com/iam/docs/understanding-roles#primitive_roles

You can shut down projects using the Cloud Console. When you shut down a project, this immediately happens: All billing and traffic serving stops, You lose access to the project, The owners of the project will be notified and can stop the deletion within 30 days, The project will be scheduled to be deleted after 30 days. However, some resources may be deleted much earlier.

NEW QUESTION 177

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 179

Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1-standard-2 nodes. You need to deploy additional pods requiring n2-highmem-16 nodes without any downtime. What should you do?

- A. Use gcloud container clusters upgrad
- B. Deploy the new services.
- C. Create a new Node Pool and specify machine type n2-highmem-16. Deploy the new pods.
- D. Create a new cluster with n2-highmem-16 node
- E. Redeploy the pods and delete the old cluster.
- F. Create a new cluster with both n1-standard-2 and n2-highmem-16 node
- G. Redeploy the pods and delete the old cluster.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/deployment>

NEW QUESTION 184

You deployed an application on a managed instance group in Compute Engine. The application accepts Transmission Control Protocol (TCP) traffic on port 389 and requires you to preserve the IP address of the client who is making a request. You want to expose the application to the internet by using a load balancer. What should you do?

- A. Expose the application by using an external TCP Network Load Balancer.
- B. Expose the application by using a TCP Proxy Load Balancer.
- C. Expose the application by using an SSL Proxy Load Balancer.
- D. Expose the application by using an internal TCP Network Load Balancer.

Answer: B

NEW QUESTION 188

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

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

Answer: A

Explanation:

<https://cloud.google.com/bigquery/docs/quickstarts/quickstart-web-ui?authuser=4>

NEW QUESTION 191

Your company has an internal application for managing transactional orders. The application is used exclusively by employees in a single physical location. The application requires strong consistency, fast queries, and ACID guarantees for multi-table transactional updates. The first version of the application is implemented in PostgreSQL, and you want to deploy it to the cloud with minimal code changes. Which database is most appropriate for this application?

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

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres>

NEW QUESTION 193

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

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

Answer: B

Explanation:

➤ Navigate to Stackdriver Logging and select resource.labels.project_id=*. is not right.

Log entries are held in Stackdriver Logging for a limited time known as the retention period which is 30 days (default configuration). After that, the entries are deleted. To keep log entries longer, you need to export them outside of Stackdriver Logging by configuring log sinks.

Ref: <https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging> ➤ Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days. is not right.

While this works, it makes no sense to use Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery when Google provides a feature (export sinks) that does exactly the same thing and works out of the box. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2

➤ Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days. is not right.

You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and

Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2

Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud

organization. Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks

Either way, we now have the data in Cloud Storage, but querying logs information from Cloud Storage is harder than Querying information from BigQuery dataset. For this reason, we should prefer Big Query over Cloud Storage.

➤ Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days. is the right answer.

You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and

Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2

Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud

organization. Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks

Either way, we now have the data in a BigQuery Dataset. Querying information from a Big Query dataset is easier and quicker than analyzing contents in Cloud Storage bucket. As our requirement is to Quickly analyze the log contents, we should prefer Big Query over Cloud Storage.

Also, You can control storage costs and optimize storage usage by setting the default table expiration for newly created tables in a dataset. If you set the property when the dataset is created, any table created in the dataset is deleted after the expiration period. If you set the property after the dataset is created, only new tables are deleted after the expiration period. For example, if you set the default table expiration to 7 days, older data is automatically deleted after 1 week. Ref: <https://cloud.google.com/bigquery/docs/best-practices-storage>

NEW QUESTION 196

You are designing an application that uses WebSockets and HTTP sessions that are not distributed across the web servers. You want to ensure the application runs properly on Google Cloud Platform. What should you do?

- A. Meet with the cloud enablement team to discuss load balancer options.
- B. Redesign the application to use a distributed user session service that does not rely on WebSockets and HTTP sessions.
- C. Review the encryption requirements for WebSocket connections with the security team.
- D. Convert the WebSocket code to use HTTP streaming.

Answer: A

Explanation:

➤ Google HTTP(S) Load Balancing has native support for the WebSocket protocol when you use HTTP or HTTPS, not HTTP/2, as the protocol to the backend. Ref: https://cloud.google.com/load-balancing/docs/https#websocket_proxy_support

➤ We dont need to convert WebSocket code to use HTTP streaming or Redesign the application, as

WebSocket support is offered by Google HTTP(S) Load Balancing. Reviewing the encryption requirements is a good idea but it has nothing to do with WebSockets.

NEW QUESTION 198

You have sensitive data stored in three Cloud Storage buckets and have enabled data access logging. You want to verify activities for a particular user for these buckets, using the fewest possible steps. You need to verify the addition of metadata labels and which files have been viewed from those buckets. What should you do?

- A. Using the GCP Console, filter the Activity log to view the information.
- B. Using the GCP Console, filter the Stackdriver log to view the information.
- C. View the bucket in the Storage section of the GCP Console.

D. Create a trace in Stackdriver to view the information.

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/audit-logs> https://cloud.google.com/compute/docs/logging/audit-logging#audited_operations

NEW QUESTION 200

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 202

You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

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

Answer: D

Explanation:

<https://cloud.google.com/monitoring/audit-logging>

NEW QUESTION 206

You are planning to migrate your on-premises data to Google Cloud. The data includes:

- 200 TB of video files in SAN storage
- Data warehouse data stored on Amazon Redshift
- 20 GB of PNG files stored on an S3 bucket

You need to load the video files into a Cloud Storage bucket, transfer the data warehouse data into BigQuery, and load the PNG files into a second Cloud Storage bucket. You want to follow Google-recommended practices and avoid writing any code for the migration. What should you do?

- A. Use gcloud storage for the video file
- B. Dataflow for the data warehouse data, and Storage Transfer Service for the PNG files.
- C. Use Transfer Appliance for the video
- D. BigQuery Data Transfer Service for the data warehouse data, and Storage Transfer Service for the PNG files.
- E. Use Storage Transfer Service for the video files, BigQuery Data Transfer Service for the data warehouse data, and Storage Transfer Service for the PNG files.
- F. Use Cloud Data Fusion for the video files, Dataflow for the data warehouse data, and Storage Transfer Service for the PNG files.

Answer: C

NEW QUESTION 207

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

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

Answer: A

Explanation:

"Run `gcloud configurations list` to start the Compute Engine instances". How the heck are you expecting to "start" GCE instances doing "configuration list". Each `gcloud configuration` has a 1 to 1 relationship with the region (if a region is defined). Since we have two different regions, we would need to create two separate configurations using `gcloud config configurations create`Ref: <https://cloud.google.com/sdk/gcloud/reference/config/configurations/create> Secondly, you can activate each configuration independently by running `gcloud config configurations activate [NAME]`Ref: <https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate>

Finally, while each configuration is active, you can run the `gcloud compute instances start [NAME]` command to start the instance in the configurations region. <https://cloud.google.com/sdk/gcloud/reference/compute/instances/start>

NEW QUESTION 208

You need to deploy an application in Google Cloud using serverless technology. You want to test a new version of the application with a small percentage of production traffic. What should you do?

- A. Deploy the application to Cloud
- B. Run
- C. Use gradual rollouts for traffic splitting.
- D. Deploy the application to Google Kubernetes Engine
- E. Use Anthos Service Mesh for traffic splitting.
- F. Deploy the application to Cloud Function
- G. Suffix the version number in the function name.
- H. Deploy the application to App Engine
- I. For each new version, create a new service.

Answer: A

NEW QUESTION 213

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

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

Answer: A

Explanation:

Ref: https://cloud.google.com/compute/docs/machine-types#n1_machine_type

NEW QUESTION 215

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

Answer: C

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

NEW QUESTION 217

You built an application on your development laptop that uses Google Cloud services. Your application uses Application Default Credentials for authentication and works fine on your development laptop. You want to migrate this application to a Compute Engine virtual machine (VM) and set up authentication using Google-recommended practices and minimal changes. What should you do?

- A. Assign appropriate access for Google services to the service account used by the Compute Engine VM.
- B. Create a service account with appropriate access for Google services, and configure the application to use this account.
- C. Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
- D. Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.

Answer: B

Explanation:

In general, Google recommends that each instance that needs to call a Google API should run as a service account with the minimum permissions necessary for that instance to do its job. In practice, this means you should configure service accounts for your instances with the following process: Create a new service account rather than using the Compute Engine default service account. Grant IAM roles to that service account for only the resources that it needs. Configure the instance to run as that service account. Grant the instance the <https://www.googleapis.com/auth/cloud-platform> scope to allow full access to all Google Cloud APIs, so that the IAM permissions of the instance are completely determined by the IAM roles of the service account. Avoid granting more access than necessary and regularly check your service account permissions to make sure they are up-to-date.

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

NEW QUESTION 218

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

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

Answer: D

Explanation:

As part of your workload, there might be certain VM instances that are critical to running your application or services, such as an instance running a SQL server, a server used as a license manager, and so on. These VM instances might need to stay running indefinitely so you need a way to protect these VMs from being deleted. By setting the deletionProtection flag, a VM instance can be protected from accidental deletion. If a user attempts to delete a VM instance for which you have set the deletionProtection flag, the request fails. Only a user that has been granted a role with compute.instances.create permission can reset the flag to allow the resource to be deleted. Ref: <https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

NEW QUESTION 222

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 GPU
- D. Dedicate this cluster to your ML team.
- E. Add a new, GPU-enabled, node pool to the GKE cluster
- F. Ask your ML team to add the cloud.google.com/gke -accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Answer: D

Explanation:

This is the most optimal solution. Rather than recreating all nodes, you create a new node pool with GPU enabled. You then modify the pod specification to target particular GPU types by adding node selector to your workloads Pod specification. You still have a single cluster so you pay Kubernetes cluster management fee for just one cluster thus minimizing the cost. Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/gpus> Ref: <https://cloud.google.com/kubern>

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
```

NEW QUESTION 226

You are about to deploy a new Enterprise Resource Planning (ERP) system on Google Cloud. The application holds the full database in-memory for fast data access, and you need to configure the most appropriate resources on Google Cloud for this application. What should you do?

- A. Provision preemptible Compute Engine instances.
- B. Provision Compute Engine instances with GPUs attached.
- C. Provision Compute Engine instances with local SSDs attached.
- D. Provision Compute Engine instances with M1 machine type.

Answer: D

Explanation:

M1 machine series Medium in-memory databases such as SAP HANA Tasks that require intensive use of memory with higher memory-to-vCPU ratios than the general-purpose high-memory machine types.

In-memory databases and in-memory analytics, business warehousing (BW) workloads, genomics analysis, SQL analysis services. Microsoft SQL Server and similar databases.

<https://cloud.google.com/compute/docs/machine-types>

<https://cloud.google.com/compute/docs/machine-types#:~:text=databases%20such%20as-,SAP%20HANA,-In%20>

<https://www.sap.com/india/products/hana.html#:~:text=is%20SAP%20HANA-,in%20memory,-database%3F>

NEW QUESTION 231

You need to add a group of new users to Cloud Identity. Some of the users already have existing Google accounts. You want to follow one of Google's recommended practices and avoid conflicting accounts. What should you do?

- A. Invite the user to transfer their existing account
- B. Invite the user to use an email alias to resolve the conflict
- C. Tell the user that they must delete their existing account

D. Tell the user to remove all personal email from the existing account

Answer: A

Explanation:

<https://cloud.google.com/architecture/identity/migrating-consumer-accounts>

NEW QUESTION 233

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-58ddb995-lp86m    0/1     Pending  0        9m
myapp-deployment-58ddb995-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-58ddb995-lp86m Pod and check for warning messages.
- D. View logs of the container in myapp-deployment-58ddb995-lp86m pod and check for warning messages.

Answer: C

Explanation:

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

NEW QUESTION 237

Users of your application are complaining of slowness when loading the application. You realize the slowness is because the App Engine deployment serving the application is deployed in us-central whereas all users of this application are closest to europe-west3. You want to change the region of the App Engine application to europe-west3 to minimize latency. What's the best way to change the App Engine region?

- A. Create a new project and create an App Engine instance in europe-west3
- B. Use the gcloud app region set command and supply the name of the new region.
- C. From the console, under the App Engine page, click edit, and change the region drop-down.
- D. Contact Google Cloud Support and request the change.

Answer: A

Explanation:

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it cant be changed. The only way is to create a new project and create an App Engine instance in europe-west3, send all user traffic to this instance and delete the app engine instance in us-central.

Ref: <https://cloud.google.com/appengine/docs/locations>

NEW QUESTION 239

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.

Answer: A

Explanation:

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

NEW QUESTION 244

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

Answer: A

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>

NEW QUESTION 248

You want to permanently delete a Pub/Sub topic managed by Config Connector in your Google Cloud project. What should you do?

- A. Use kubectl to delete the topic resource.
- B. Use gcloud CLI to delete the topic.
- C. Use kubectl to create the label deleted-by-cnrm and to change its value to true for the topic resource.
- D. Use gcloud CLI to update the topic label managed-by-cnrm to false.

Answer: A

NEW QUESTION 250

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

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

Answer: D

Explanation:

<https://cloud.google.com/monitoring/settings/multiple-projects> <https://cloud.google.com/monitoring/workspaces>

NEW QUESTION 253

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 me 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 now release, create a new deployment for the new version Update the service e to use the now deployment.

Answer: D

Explanation:

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

NEW QUESTION 257

You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command should you use?

- A. gcloud deployment-manager deployments create --config <deployment-config-path>
- B. gcloud deployment-manager deployments update --config <deployment-config-path>

- C. gcloud deployment-manager resources create --config <deployment-config-path>
- D. gcloud deployment-manager resources update --config <deployment-config-path>

Answer: B

NEW QUESTION 262

You have just created a new project which will be used to deploy a globally distributed application. You will use Cloud Spanner for data storage. You want to create a Cloud Spanner instance. You want to perform the first step in preparation of creating the instance. What should you do?

- A. Grant yourself the IAM role of Cloud Spanner Admin
- B. Create a new VPC network with subnetworks in all desired regions
- C. Configure your Cloud Spanner instance to be multi-regional
- D. Enable the Cloud Spanner API

Answer: A

Explanation:

<https://cloud.google.com/spanner/docs/getting-started/set-up>

NEW QUESTION 263

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
- D. Use the gcloud command line tool to convert them into Google Cloud Platform accounts.
- E. 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.

Answer: B

NEW QUESTION 264

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

Answer: C

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

NEW QUESTION 268

Your company implemented BigQuery as an enterprise data warehouse. Users from multiple business units run queries on this data warehouse. However, you notice that query costs for BigQuery are very high, and you need to control costs. Which two methods should you use? (Choose two.)

- A. Split the users from business units to multiple projects.
- B. Apply a user- or project-level custom query quota for BigQuery data warehouse.
- C. Create separate copies of your BigQuery data warehouse for each business unit.
- D. Split your BigQuery data warehouse into multiple data warehouses for each business unit.
- E. Change your BigQuery query model from on-demand to flat rate
- F. Apply the appropriate number of slots to each Project.

Answer: BE

Explanation:

<https://cloud.google.com/bigquery/docs/custom-quotas> https://cloud.google.com/bigquery/pricing#flat_rate_pricing

NEW QUESTION 273

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

Answer:

A

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>

NEW QUESTION 278

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

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

Answer: B

Explanation:

Google Cloud provides Cloud Audit Logs, which is an integral part of Cloud Logging. It consists of two log streams for each project: Admin Activity and Data Access, which are generated by Google Cloud services to help you answer the question of who did what, where, and when? within your Google Cloud projects.

Ref: https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors

NEW QUESTION 283

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 includedPermission
- B. Grant the custom role to the user at the project level.
- C. Create a custom role based on the Compute Image User role Add the compute.disks, list to theincludedPermissions field Grant the custom role to the user at the project level
- D. Grant the Compute Storage Admin role at the project level.
- E. Create a custom role based on the Compute Storage Admin rol
- F. Exclude unnecessary permissions from the custom rol
- G. Grant the custom role to the user at the project level.

Answer: B

NEW QUESTION 285

You are monitoring an application and receive user feedback that a specific error is spiking. You notice that the error is caused by a Service Account having insufficient permissions. You are able to solve the problem but want to be notified if the problem recurs. What should you do?

- A. In the Log Viewer, filter the logs on severity 'Error' and the name of the Service Account.
- B. Create a sink to BigQuery to export all the log
- C. Create a Data Studio dashboard on the exported logs.
- D. Create a custom log-based metric for the specific error to be used in an Alerting Policy.
- E. Grant Project Owner access to the Service Account.

Answer: C

NEW QUESTION 287

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

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

Answer: D

NEW QUESTION 291

You have a number of compute instances belonging to an unmanaged instances group. You need to SSH to one of the Compute Engine instances to run an ad hoc script. You've already authenticated gcloud, however, you don't have an SSH key deployed yet. In the fewest steps possible, what's the easiest way to SSH to the instance?

- A. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.
- B. Use the gcloud compute ssh command.
- C. Create a key with the ssh-keygen comman
- D. Then use the gcloud compute ssh command.
- E. Create a key with the ssh-keygen comman
- F. Upload the key to the instanc
- G. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.

Answer: B

Explanation:

gcloud compute ssh ensures that the user's public SSH key is present in the project's metadata. If the user does not have a public SSH key, one is generated using ssh-keygen and added to the project's metadata. This is similar to the other option where we copy the key explicitly to the project's metadata but here it is done automatically for us. There are also security benefits with this approach. When we use gcloud compute ssh to connect to Linux instances, we are adding a layer of security by storing your host keys as guest attributes. Storing SSH host keys as guest attributes improve the security of your connections by helping to protect against vulnerabilities such as man-in-the-middle (MITM) attacks. On the initial boot of a VM instance, if guest attributes are enabled, Compute Engine stores your generated host keys as guest attributes.

Compute Engine then uses these host keys that were stored during the initial boot to verify all subsequent connections to the VM instance.

Ref: <https://cloud.google.com/compute/docs/instances/connecting-to-instance>Ref: <https://cloud.google.com/s>

NEW QUESTION 294

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

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

Answer: A

Explanation:

<https://cloud.google.com/compute/docs/autoscaler#specifications>

Autoscaling works independently from autohealing. If you configure autohealing for your group and an instance fails the health check, the autohealer attempts to recreate the instance. Recreating an instance can cause the number of instances in the group to fall below the autoscaling threshold (minNumReplicas) that you specify.

- > Since we need the application running at all times, we need a minimum 1 instance.
- > Only a single instance of the VM should run, we need a maximum 1 instance.
- > We want the application running at all times. If the VM crashes due to any underlying hardware failure, we want another instance to be added to MIG so that application can continue to serve requests. We can achieve this by enabling autoscaling. The only option that satisfies these three is Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.

Ref: <https://cloud.google.com/compute/docs/autoscaler>

NEW QUESTION 299

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

- * Associate-Cloud-Engineer Questions and Answers Updated Frequently
- * Associate-Cloud-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Associate-Cloud-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Associate-Cloud-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Associate-Cloud-Engineer Practice Test Here](#)