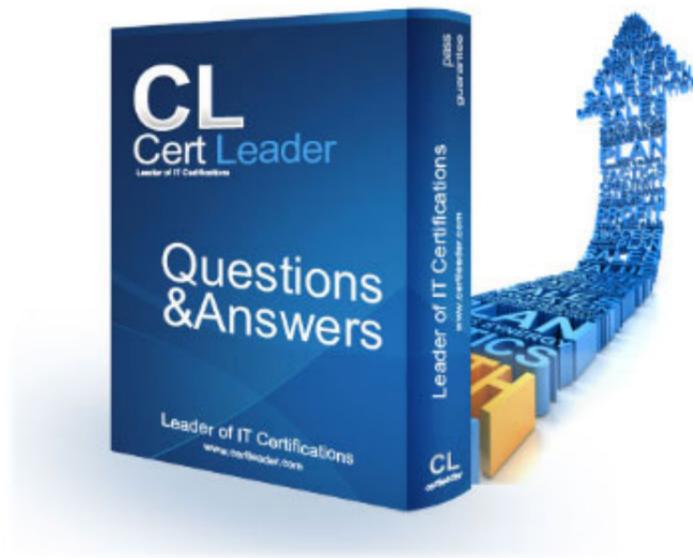


Associate-Cloud-Engineer Dumps

Google Cloud Certified - Associate Cloud Engineer

<https://www.certleader.com/Associate-Cloud-Engineer-dumps.html>



NEW QUESTION 1

You recently discovered that your developers are using many service account keys during their development process. While you work on a long term improvement, you need to quickly implement a process to enforce short-lived service account credentials in your company. You have the following requirements:

- All service accounts that require a key should be created in a centralized project called pj-sa.
- Service account keys should only be valid for one day.

You need a Google-recommended solution that minimizes cost. What should you do?

- A. Implement a Cloud Run job to rotate all service account keys periodically in pj-s
- B. Enforce an org policy to deny service account key creation with an exception to pj-sa.
- C. Implement a Kubernetes Cronjob to rotate all service account keys periodically
- D. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.
- E. Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hours
- F. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.
- G. Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hours
- H. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

Answer: C

Explanation:

According to the Google Cloud documentation, you can use organization policy constraints to control the creation and expiration of service account keys. The constraints are:

- `constraints/iam.allowServiceAccountKeyCreation`: This constraint allows you to specify which projects or folders can create service account keys. You can set the value to true or false, or use a condition to apply the constraint to specific service accounts. By setting this constraint to false for the organization and adding an exception for the pj-sa project, you can prevent developers from creating service account keys in other projects.
- `constraints/iam.serviceAccountKeyMaxLifetime`: This constraint allows you to specify the maximum lifetime of service account keys. You can set the value to a duration in seconds, such as 86400 for one day. By setting this constraint to 86400 for the organization, you can ensure that all service account keys expire after one day.

These constraints are recommended by Google Cloud as best practices to minimize the risk of service account key misuse or compromise. They also help you reduce the cost of managing service account keys, as you do not need to implement a custom solution to rotate or delete them.

References:

- 1: Associate Cloud Engineer Certification Exam Guide | Learn - Google Cloud
- 5: Create and delete service account keys - Google Cloud
- Organization policy constraints for service accounts

NEW QUESTION 2

Your coworker has helped you set up several configurations for gcloud. You've noticed that you're running commands against the wrong project. Being new to the company, you haven't yet memorized any of the projects. With the fewest steps possible, what's the fastest way to switch to the correct configuration?

- A. Run `gcloud configurations list` followed by `gcloud configurations activate .`
- B. Run `gcloud config list` followed by `gcloud config activate.`
- C. Run `gcloud config configurations list` followed by `gcloud config configurations activate.`
- D. Re-authenticate with the `gcloud auth login` command and select the correct configurations on login.

Answer: C

Explanation:

as `gcloud config configurations list` can help check for the existing configurations and `activate` can help switch to the configuration.

`gcloud config configurations list` lists existing named configurations

`gcloud config configurations activate` activates an existing named configuration

Obtains access credentials for your user account via a web-based authorization flow. When this command completes successfully, it sets the active account in the current configuration to the account specified. If no configuration exists, it creates a configuration named `default`.

NEW QUESTION 3

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

What should you do?

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

Answer: B

Explanation:

Forwarding zones Cloud DNS forwarding zones let you configure target name servers for specific private zones. Using a forwarding zone is one way to implement outbound DNS forwarding from your VPC network. A Cloud DNS forwarding zone is a special type of Cloud DNS private zone. Instead of creating records within the zone, you specify a set of forwarding targets. Each forwarding target is an IP address of a DNS server, located in your VPC network, or in an on-premises network connected to your VPC network by Cloud VPN or Cloud Interconnect.

<https://cloud.google.com/nat/docs/overview>

DNS configuration Your on-premises network must have DNS zones and records configured so that Google domain names resolve to the set of IP addresses for either `private.googleapis.com` or `restricted.googleapis.com`. You can create Cloud DNS managed private zones and use a Cloud DNS inbound server policy, or you can configure on-premises name servers. For example, you can use BIND or Microsoft Active Directory DNS.

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

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.

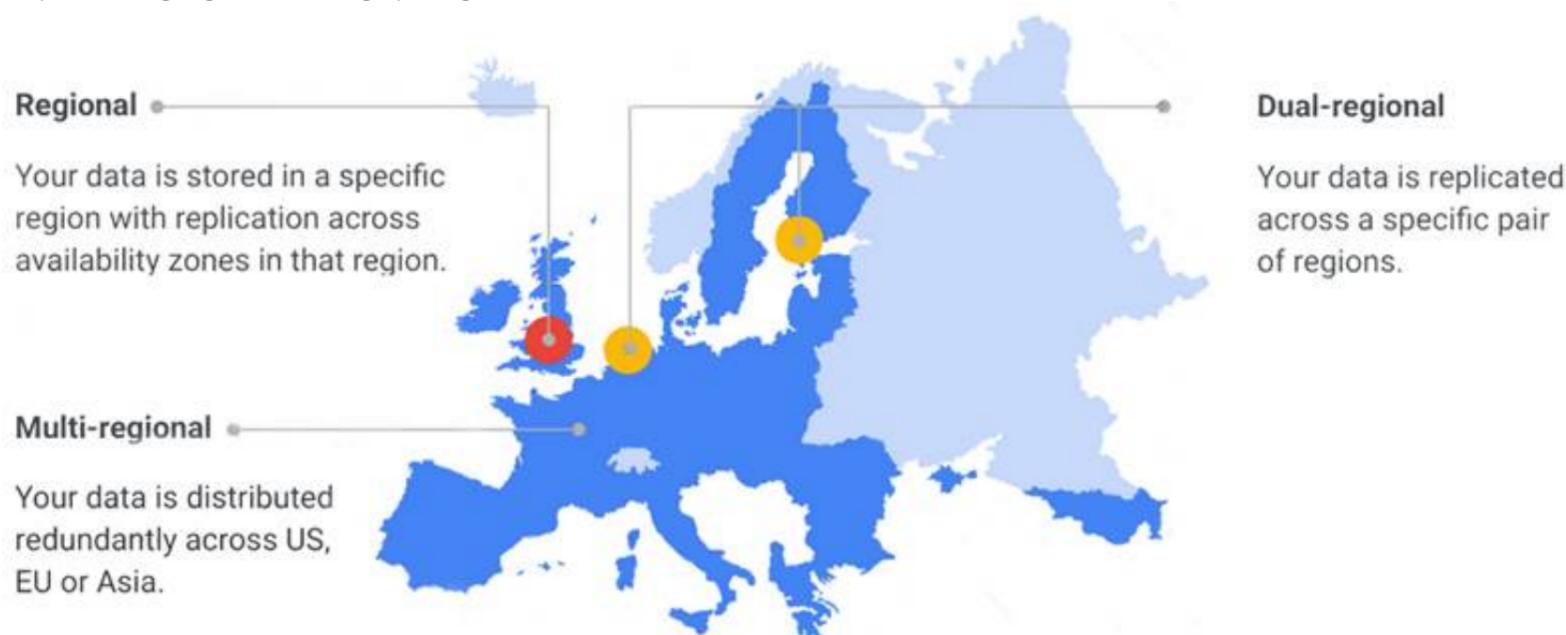
Answer: C

Explanation:

Cloud Storage allows world-wide storage and retrieval of any amount of data at any time. We don't 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 VM's 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

Your company has developed a new application that consists of multiple microservices. You want to deploy the application to Google Kubernetes Engine (GKE), and you want to ensure that the cluster can scale as more applications are deployed in the future. You want to avoid manual intervention when each new application is deployed. What should you do?

- A. Deploy the application on GKE, and add a HorizontalPodAutoscaler to the deployment.
- B. Deploy the application on GKE, and add a VerticalPodAutoscaler to the deployment.
- C. Create a GKE cluster with autoscaling enabled on the node pool
- D. Set a minimum and maximum for the size of the node pool.
- E. Create a separate node pool for each application, and deploy each application to its dedicated node pool.

Answer: C

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/cluster-autoscaler#adding_a_node_pool_with_autoscal

NEW QUESTION 7

Your company is moving its entire workload to Compute Engine. Some servers should be accessible through the Internet, and other servers should only be accessible over the internal network. All servers need to be able to talk to each other over specific ports and protocols. The current on-premises network relies on a demilitarized zone (DMZ) for the public servers and a Local Area Network (LAN) for the private servers. You need to design the networking infrastructure on

Google Cloud to match these requirements. What should you do?

- A. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LA
- B. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public ingress traffic for the DMZ.
- C. 1. Create a single VPC with a subnet for the DMZ and a subnet for the LA
- D. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.
- E. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LA
- F. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public ingress traffic for the DMZ.
- G. 1. Create a VPC with a subnet for the DMZ and another VPC with a subnet for the LA
- H. 2. Set up firewall rules to open up relevant traffic between the DMZ and the LAN subnets, and another firewall rule to allow public egress traffic for the DMZ.

Answer: C

Explanation:

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

NEW QUESTION 8

Your team is using Linux instances on Google Cloud. You need to ensure that your team logs in to these instances in the most secure and cost efficient way. What should you do?

- A. Attach a public IP to the instances and allow incoming connections from the internet on port 22 for SSH.
- B. Use a third party tool to provide remote access to the instances.
- C. Use the `gcloud compute ssh` command with the `--tunnel-through-iap` flag
- D. Allow ingress traffic from the IP range 35.235.240.0/20 on port 22.
- E. Create a bastion host with public internet access
- F. Create the SSH tunnel to the instance through the bastion host.

Answer: D

NEW QUESTION 9

Every employee of your company has a Google account. Your operational team needs to manage a large number of instances on Compute Engine. Each member of this team needs only administrative access to the servers. Your security team wants to ensure that the deployment of credentials is operationally efficient and must be able to determine who accessed a given instance. What should you do?

- A. Generate a new SSH key pair
- B. Give the private key to each member of your team
- C. Configure the public key in the metadata of each instance.
- D. Ask each member of the team to generate a new SSH key pair and to send you their public key
- E. Use a configuration management tool to deploy those keys on each instance.
- F. Ask each member of the team to generate a new SSH key pair and to add the public key to their Google account
- G. Grant the "compute.osAdminLogin" role to the Google group corresponding to this team.
- H. Generate a new SSH key pair
- I. Give the private key to each member of your team
- J. Configure the public key as a project-wide public SSH key in your Cloud Platform project and allow project-wide public SSH keys on each instance.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instances/managing-instance-access>

NEW QUESTION 10

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 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 12

Your application is running on Google Cloud in a managed instance group (MIG). You see errors in Cloud Logging for one VM that one of the processes is not responsive. You want to replace this VM in the MIG quickly. What should you do?

- A. Select the MIG from the Compute Engine console and, in the menu, select Replace VMs.
- B. Use the `gcloud compute instance-groups managed recreate-instances` command to recreate the VM.
- C. Use the `gcloud compute instances update` command with a REFRESH action for the VM.
- D. Update and apply the instance template of the MIG.

Answer: A

NEW QUESTION 14

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 16

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

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

Answer: D

Explanation:

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

NEW QUESTION 20

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 23

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 27

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

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

Answer: A

Explanation:

➤ roles/monitoring.viewer provides read-only access to get and list information about all monitoring data and configurations. This role provides monitoring access and fits our requirements. roles/monitoring.viewer. is the right answer.
Ref: <https://cloud.google.com/iam/docs/understanding-roles#cloud-spanner-roles>

NEW QUESTION 29

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

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

Answer: D

NEW QUESTION 32

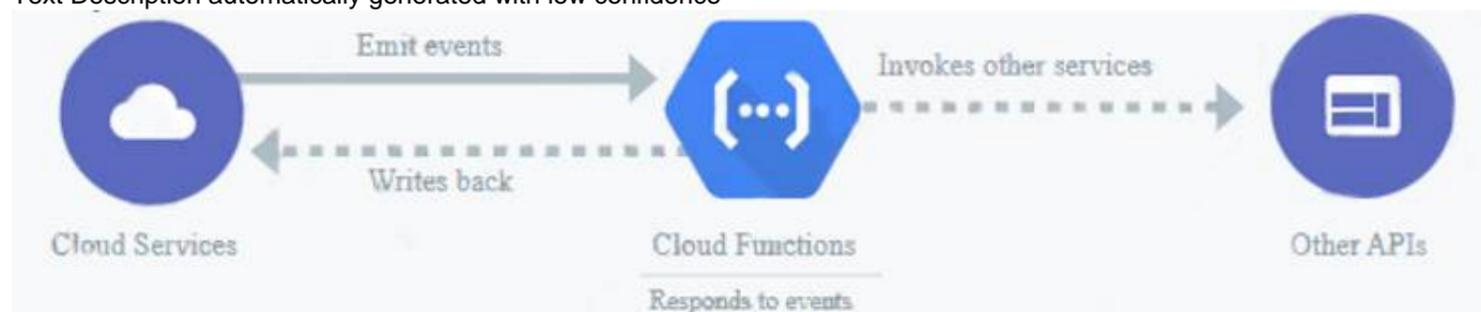
A company wants to build an application that stores images in a Cloud Storage bucket and wants to generate thumbnails as well as resize the images. They want to use a google managed service that can scale up and scale down to zero automatically with minimal effort. You have been asked to recommend a service. Which GCP service would you suggest?

- A. Google Compute Engine
- B. Google App Engine
- C. Cloud Functions
- D. Google Kubernetes Engine

Answer: C

Explanation:

Text Description automatically generated with low confidence



Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.

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

While all other options i.e. Google Compute Engine, Google Kubernetes Engine, Google App Engine support autoscaling, it needs to be configured explicitly based on the load and is not as trivial as the scale up or scale down offered by Google's cloud functions.

NEW QUESTION 35

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

- A. In the console, validate which SSH keys have been stored as project-wide keys.
- B. Navigate to Identity-Aware Proxy and check the permissions for these resources.
- C. Enable Audit Logs on the IAM & admin page for all resources, and validate the results.

D. Use the command `gcloud projects get-iam-policy` to view the current role assignments.

Answer: D

Explanation:

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

NEW QUESTION 40

You are deploying an application to App Engine. You want the number of instances to scale based on request rate. You need at least 3 unoccupied instances at all times. Which scaling type should you use?

- A. Manual Scaling with 3 instances.
- B. Basic Scaling with `min_instances` set to 3.
- C. Basic Scaling with `max_instances` set to 3.
- D. Automatic Scaling with `min_idle_instances` set to 3.

Answer: D

NEW QUESTION 44

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 48

You need to create a Compute Engine instance in a new project that doesn't exist yet. What should you do?

- A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.
- B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the `—project` flag to specify a new project.
- C. Using the Cloud SDK, create the new instance, and use the `—project` flag to specify the new project. Answer yes when prompted by Cloud SDK to enable the Compute Engine API.
- D. Enable the Compute Engine API in the Cloud Console
- E. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/projects/create> Quickstart: Creating a New Instance Using the Command Line Before you begin

- * 1. In the Cloud Console, on the project selector page, select or create a Cloud project.
- * 2. Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project.

To use the `gcloud` command-line tool for this quickstart, you must first install and initialize the Cloud SDK:

- * 1. Download and install the Cloud SDK using the instructions given on Installing Google Cloud SDK.
- * 2. Initialize the SDK using the instructions given on Initializing Cloud SDK.

To use `gcloud` in Cloud Shell for this quickstart, first activate Cloud Shell using the instructions given on Starting Cloud Shell.

<https://cloud.google.com/ai-platform/deep-learning-vm/docs/quickstart-cli#before-you-begin>

NEW QUESTION 50

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 objec
- C. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
- D. Store the database password inside a ConfigMap objec
- 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 51

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 56

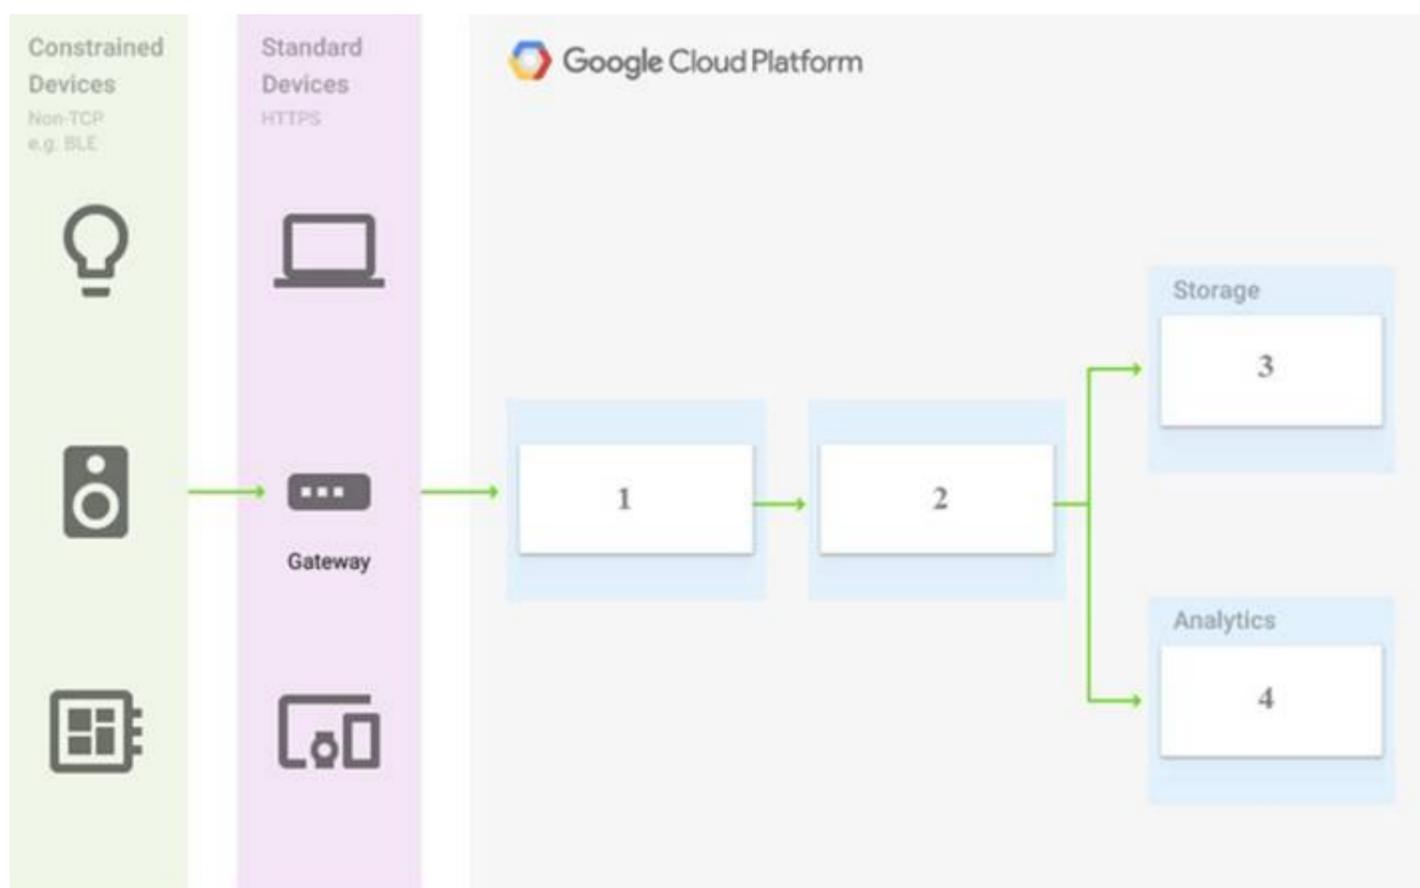
Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

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

Answer: D

NEW QUESTION 60

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 63

You are in charge of provisioning access for all Google Cloud users in your organization. Your company recently acquired a startup company that has their own Google Cloud organization. You need to ensure that your Site Reliability Engineers (SREs) have the same project permissions in the startup company's organization as in your own organization. What should you do?

- A. In the Google Cloud console for your organization, select Create role from selection, and choose destination as the startup company's organization
- B. In the Google Cloud console for the startup company, select Create role from selection and choose source as the startup company's Google Cloud organization.
- C. Use the `gcloud iam roles copy` command, and provide the Organization ID of the startup company's Google Cloud Organization as the destination.
- D. Use the `gcloud iam roles copy` command, and provide the project IDs of all projects in the startup company's organization as the destination.

Answer: C

Explanation:

<https://cloud.google.com/architecture/best-practices-vpc-design#shared-service> Cloud VPN is another alternative. Because Cloud VPN establishes reachability through managed IPsec tunnels, it doesn't have the aggregate limits of VPC Network Peering. Cloud VPN uses a VPN Gateway for connectivity and doesn't consider the aggregate resource use of the IPsec peer. The drawbacks of Cloud VPN include increased costs (VPN tunnels and traffic egress), management overhead required to maintain tunnels, and the performance overhead of IPsec.

NEW QUESTION 67

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 72

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 73

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

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

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

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

Answer: AB

Explanation:

(A) Object Lifecycle Management Delete

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

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

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

(B) Signed URLs

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

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

NEW QUESTION 75

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 rolestage 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 79

You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

- A. 1. Use nslookup to get the IP address for storage.googleapis.com.2. Negotiate with the security team to be able to give a public IP address to the servers.3. Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.
- B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP).2. In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.
- C. 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.2. Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.3. Configure your new instances to use this ILB as proxy.
- D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

Answer: D

Explanation:

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

- Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP
- Using Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.
- In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com is the right answer right, and it is what Google recommends.
Ref: <https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid>
- You must configure routes so that Google API traffic is forwarded through your Cloud VPN or Cloud Interconnect connection, firewall rules on your on-premises firewall to allow the outgoing traffic, and DNS so that traffic to Google APIs resolves to the IP range youve added to your routes.
- You can use Cloud Router Custom Route Advertisement to announce the Restricted Google APIs IP addresses through Cloud Router to your on-premises network. The Restricted Google APIs IP range is 199.36.153.4/30. While this is technically a public IP range, Google does not announce it publicly. This IP range is only accessible to hosts that can reach your Google Cloud projects through internal IP ranges, such as through a Cloud VPN or Cloud Interconnect connection. Without having a public IP address or access to the internet, the only way you could connect to cloud storage is if you have an internal route to it.
- So Negotiate with the security team to be able to give public IP addresses to the servers is not right.
Following Google recommended practices is synonymous with using Googles services (Not quite, but it is at least for the exam !!).
- So In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance is not right.
- Migrating the VM to Compute Engine is a bit drastic when Google says it is perfectly fine to have Hybrid Connectivity architectures
<https://cloud.google.com/hybrid-connectivity>.
- So,
- Use Migrate for Compute Engine (formerly known as Velostrata) to migrate these servers to Compute Engine is not right.

NEW QUESTION 84

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 ngine 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 87

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

- A. Verify that you are Project Billing Manager for the GCP projec
- B. Update the existing project to link it to the existing billing account.
- C. Verify that you are Project Billing Manager for the GCP projec
- D. Create a new billing account and link the new billing account to the existing project.
- E. Verify that you are Billing Administrator for the billing accoun
- F. Create a new project and link the new project to the existing billing account.
- G. Verify that you are Billing Administrator for the billing accoun
- H. Update the existing project to link it to the existing billing account.

Answer: B

Explanation:

Billing Administrators can not create a new billing account, and the project is presumably already created. Project Billing Manager allows you to link the created billing account to the project. It is vague on how the billing account gets created but by process of elimination

NEW QUESTION 90

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

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

Answer: A

Explanation:

Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account. is not right. As mentioned above, Container Registry ignores permissions set on individual objects within the storage bucket so this isn't going to work.
Ref: <https://cloud.google.com/container-registry/docs/access-control>

NEW QUESTION 92

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 the 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 97

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

- A. Verify that both projects are in a GCP Organization
- B. Create a new VPC and add all instances.
- C. Verify that both projects are in a GCP Organization
- D. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.
- E. Verify that you are the Project Administrator of both projects
- F. Create two new VPCs and add all instances.
- G. Verify that you are the Project Administrator of both projects
- H. Create a new VPC and add all instances.

Answer: B

Explanation:

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

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

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

NEW QUESTION 98

Your company developed an application to deploy on Google Kubernetes Engine. Certain parts of the application are not fault-tolerant and are allowed to have downtime. Other parts of the application are critical and must always be available. You need to configure a Google Kubernetes Engine cluster while optimizing for cost. What should you do?

- A. Create a cluster with a single node-pool by using standard VM
- B. Label the fault-tolerant Deployments as spot-true.
- C. Create a cluster with a single node-pool by using Spot VM
- D. Label the critical Deployments as spot-false.
- E. Create a cluster with both a Spot VM node pool and a node pool by using standard VMs. Deploy the critical
- F. deployments on the Spot VM node pool and the fault-tolerant deployments on the node pool by using standard VMs.
- G. Create a cluster with both a Spot VM node pool and by using standard VM
- H. Deploy the critical deployments on the node pool by using standard VMs and the fault-tolerant deployments on the Spot VM node pool.

Answer: C

NEW QUESTION 101

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

- A. Create two configurations using gcloud config
- B. Write a script that sets configurations as active, individually
- C. For each configuration, use gcloud compute instances list to get a list of compute resources.
- D. Create two configurations using gsutil config
- E. Write a script that sets configurations as active, individually
- 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 103

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 day
- 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 104

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 107

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

- A. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Set the service's externalTrafficPolicy to Cluster.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- B. 1. In GKE, create a Service of type NodePort that uses the application's Pods as backend.2. Create a Compute Engine instance called proxy with 2 network interfaces, one in each VPC.3. Use iptables on this instance to forward traffic from gce-network to the GKE nodes.4. Configure the Compute Engine instance to use the address of proxy in gce-network as endpoint.
- C. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add an annotation to this service: cloud.google.com/load-balancer-type: Internal3. Peer the two VPCs together.4. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- D. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG'sinstances.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 111

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 Engin
- 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 Engin

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 115

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

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

Answer: B

Explanation:

The key to understand the requirement is : "The objects are written once and accessed frequently for 30 days" Standard Storage
Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

Archive Storage

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

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

NEW QUESTION 117

You manage an App Engine Service that aggregates and visualizes data from BigQuery. The application is deployed with the default App Engine Service account. The data that needs to be visualized resides in a different project managed by another team. You do not have access to this project, but you want your application to be able to read data from the BigQuery dataset. What should you do?

- A. Ask the other team to grant your default App Engine Service account the role of BigQuery Job User.
- B. Ask the other team to grant your default App Engine Service account the role of BigQuery Data Viewer.
- C. In Cloud IAM of your project, ensure that the default App Engine service account has the role of BigQuery Data Viewer.
- D. In Cloud IAM of your project, grant a newly created service account from the other team the role of BigQuery Job User in your project.

Answer: B

Explanation:

The resource that you need to get access is in the other project. roles/bigquery.dataViewer BigQuery Data Viewer

When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.

This role cannot be applied to individual models or routines. When applied to a dataset, this role provides permissions to:

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

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

NEW QUESTION 119

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

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

Answer: C

Explanation:

Understanding IAM custom roles

Key Point: Custom roles enable you to enforce the principle of least privilege, ensuring that the user and service accounts in your organization have only the permissions essential to performing their intended functions.

Basic concepts

Custom roles are user-defined, and allow you to bundle one or more supported permissions to meet your specific needs. Custom roles are not maintained by Google; when new permissions, features, or services are added to Google Cloud, your custom roles will not be updated automatically.

When you create a custom role, you must choose an organization or project to create it in. You can then grant the custom role on the organization or project, as well as any resources within that organization or project.

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

NEW QUESTION 123

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 124

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 INTOdataset.platform-logs (timestamp, log)SELECT timestamp, log FROM compute.logsWHERE timestamp> DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)3. Use Cloud Scheduler to trigger this Cloud Function once a day.

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 127

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 pag
- 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 pag
- 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 wee
- G. Navigate to the Billing Report page in the Google Cloud Platform Consol
- H. Multiply the 1 week cost to determine the monthly costs.
- I. Provision the solution on GC
- J. Leave the solution provisioned for 1 wee
- 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 dont incur any charges for doing so.

Ref: <https://cloud.google.com/products/calculator>

NEW QUESTION 132

You are developing a financial trading application that will be used globally. Data is stored and queried using a relational structure, and clients from all over the world should get the exact identical state of the data. The application will be deployed in multiple regions to provide the lowest latency to end users. You need to select a storage option for the application data while minimizing latency. What should you do?

- A. Use Cloud Bigtable for data storage.
- B. Use Cloud SQL for data storage.
- C. Use Cloud Spanner for data storage.
- D. Use Firestore for data storage.

Answer: C

Explanation:

Keywords, Financial data (large data) used globally, data stored and queried using relational structure (SQL), clients should get exact identical copies(Strong Consistency), Multiple region, low latency to end user, select storage option to minimize latency.

NEW QUESTION 135

You just installed the Google Cloud CLI on your new corporate laptop. You need to list the existing instances of your company on Google Cloud. What must you do before you run the gcloud compute instances list command?

Choose 2 answers

- A. Run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI.
- B. Create a Google Cloud service account, and download the service account ke

- C. Place the key file in a folder on your machine where gcloud CLI can find it.
- D. Download your Cloud Identity user account ke
- E. Place the key file in a folder on your machine where gcloud CLI can find it.
- F. Run gcloud config set compute/zone \$my_zone to set the default zone for gcloud CLI.
- G. Run gcloud config set project \$my_project to set the default project for gcloud CLI.

Answer: AE

Explanation:

Before you run the gcloud compute instances list command, you need to do two things: authenticate with your user account and set the default project for gcloud CLI.

To authenticate with your user account, you need to run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI. This will authorize the gcloud CLI to access Google Cloud resources on your behalf1.

To set the default project for gcloud CLI, you need to run gcloud config set project \$my_project, where

\$my_project is the ID of the project that contains the instances you want to list. This will save you from having to specify the project flag for every gcloud command2.

Option B is not recommended, because using a service account key increases the risk of credential leakage and misuse. It is also not necessary, because you can use your user account to authenticate to the gcloud CLI3. Option C is not correct, because there is no such thing as a Cloud Identity user account key. Cloud Identity is a service that provides identity and access management for Google Cloud users and groups4. Option D is not required, because the gcloud compute instances list command does not depend on the default zone. You can list instances from all zones or filter by a specific zone using the --filter flag.

References:

- > 1: <https://cloud.google.com/sdk/docs/authorizing>
- > 2: <https://cloud.google.com/sdk/gcloud/reference/config/set>
- > 3: <https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys>
- > 4: <https://cloud.google.com/identity/docs/overview>
- > : <https://cloud.google.com/sdk/gcloud/reference/compute/instances/list>

NEW QUESTION 138

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 141

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

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

Answer: A

Explanation:

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

NEW QUESTION 146

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

- A. Create an instance template for the instance
- B. Set the 'Automatic Restart' to o
- C. Set the 'On-host maintenance' to Migrate VM instanc
- D. Add the instance template to an instance group.
- E. Create an instance template for the instance
- F. Set 'Automatic Restart' to of
- G. Set 'On-host maintenance' to Terminate VM instance
- H. Add the instance template to an instance group.
- I. Create an instance group for the instance
- J. Set the 'Autohealing' health check to healthy (HTTP).
- K. Create an instance group for the instanc

L. Verify that the 'Advanced creation options' setting for 'do not retry machine creation' is set to off.

Answer: A

Explanation:

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

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

Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash. And Enabling Migrate VM Instance enables live migrates i.e. compute instances are migrated during system maintenance and remain running during the migration.

Automatic Restart If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the automaticRestart field to true. This setting does not apply if the instance is taken offline through a user action, such as calling sudo shutdown, or during a zone outage. Ref: <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>

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

NEW QUESTION 149

You are working in a team that has developed a new application that needs to be deployed on Kubernetes. The production application is business critical and should be optimized for reliability. You need to provision a Kubernetes cluster and want to follow Google-recommended practices. What should you do?

- A. Create a GKE Autopilot cluste
- B. Enroll the cluster in the rapid release channel.
- C. Create a GKE Autopilot cluste
- D. Enroll the cluster in the stable release channel.
- E. Create a zonal GKE standard cluste
- F. Enroll the cluster in the stable release channel.
- G. Create a regional GKE standard cluste
- H. Enroll the cluster in the rapid release channel.

Answer: B

Explanation:

Autopilot is more reliable and stable release gives more time to fix issues in new version of GKE

NEW QUESTION 151

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

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

Answer: D

Explanation:

As to mexblood1's point, CPU utilization is a recommended proxy for traffic when it comes to Cloud Spanner. See: Alerts for high CPU utilization The following table specifies our recommendations for maximum CPU usage for both single-region and multi-region instances. These numbers are to ensure that your instance has enough compute capacity to continue to serve your traffic in the event of the loss of an entire zone (for single-region instances) or an entire region (for multi-region instances). - <https://cloud.google.com/spanner/docs/cpu-utilization>

NEW QUESTION 153

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

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

Answer: D

Explanation:

Preventing Accidental VM Deletion This document describes how to protect specific VM instances from deletion by setting the deletionProtection property on an Instance resource. To learn more about VM instances, read the Instances documentation. 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.
<https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

NEW QUESTION 154

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

- A. Install a RDP client on your deskto
- B. Verify that a firewall rule for port 3389 exists.
- C. Install a RDP client in your deskto
- D. Set a Windows username and password in the GCP Consol
- E. Use the credentials to log in to the instance.
- F. Set a Windows password in the GCP Consol
- G. Verify that a firewall rule for port 22 exist
- H. Click the RDP button in the GCP Console and supply the credentials to log in.
- I. Set a Windows username and password in the GCP Consol
- J. Verify that a firewall rule for port 3389 exist
- K. Click the RDP button in the GCP Console, and supply the credentials to log in.

Answer: D

Explanation:

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

NEW QUESTION 157

You need to deploy an application in Google Cloud using savorless 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 lo Clou
- B. Ru
- C. Use gradual rollouts for traffic spelling.
- D. Deploy the application lo Google Kubemetes Engin
- E. Use Anthos Service Mesh for traffic splitting.
- F. Deploy the application to Cloud function
- G. Saucily the version number in the functions name.
- H. Deploy the application to App Engin
- I. For each new version, create a new service.

Answer: A

NEW QUESTION 161

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

- A. Filter the Activity log to view the Configuration categor
- B. Filter the Resource type to Service Account.
- C. Filter the Activity log to view the Configuration categor
- D. Filter the Resource type to Google Project.
- E. Filter the Activity log to view the Data Access categor
- F. Filter the Resource type to Service Account.
- G. Filter the Activity log to view the Data Access categor
- H. Filter the Resource type to Google Project.

Answer: A

Explanation:

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

NEW QUESTION 162

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 BigQuer
- 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 NDFS by compressing i
- 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 163

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 168

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 170

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

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

Answer: D

Explanation:

(https://cloud.google.com/resource-manager/docs/project-migration#change_billing_account) <https://cloud.google.com/billing/docs/concepts>
<https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION 175

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%20memory,-database%3F>

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

NEW QUESTION 180

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 183

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

```

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

```

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

```

kubectl get pods -l app=myapp
NAME                                READY   STATUS    RESTART   AGE
myapp-deployment-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 184

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

```
$ kubectl get pods -l app=myapp
NAME                                READY   STATUS    RESTART   AGE
myapp-deployment-58ddb995-1p86m    0/1     Pending   0         9m
myapp-deployment-58ddb995-qjpkg    1/1     Running   0         9m
```

What is the most likely cause?

- A. The pending Pod's resource requests are too large to fit on a single node of the cluster.
- B. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.
- C. The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.
- D. The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' status.
- E. It is currently being rescheduled on a new node.

Answer: B

Explanation:

- > The pending Pods resource requests are too large to fit on a single node of the cluster. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod. is the right answer.
- > When you have a deployment with some pods in running and other pods in the pending state, more often than not it is a problem with resources on the nodes. Here's a sample output of this use case. We see that the problem is with insufficient CPU on the Kubernetes nodes so we have to either enable auto-scaling or manually scale up the nodes.

NEW QUESTION 187

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 191

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 195

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 198

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 203

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

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

Answer: C

Explanation:

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

NEW QUESTION 204

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 206

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

- A. Run `gcloud iam roles list`
- B. Review the output section.
- C. Run `gcloud iam service-accounts list`
- D. Review the output section.
- E. Navigate to the project and then to the IAM section in the GCP Console
- F. Review the members and roles.
- G. Navigate to the project and then to the Roles section in the GCP Console
- H. Review the roles and status.

Answer: C

Explanation:

Logged onto console and followed the steps and was able to see all the assigned users and roles.

NEW QUESTION 210

You have a project for your App Engine application that serves a development environment. The required testing has succeeded and you want to create a new project to serve as your production environment. What should you do?

- A. Use `gcloud` to create the new project, and then deploy your application to the new project.
- B. Use `gcloud` to create the new project and to copy the deployed application to the new project.
- C. Create a Deployment Manager configuration file that copies the current App Engine deployment into a new project.
- D. Deploy your application again using `gcloud` and specify the project parameter with the new project name to create the new project.

Answer: A

Explanation:

You can deploy to a different project by using `--project` flag.

By default, the service is deployed the current project configured via:

```
$ gcloud config set core/project PROJECT
```

To override this value for a single deployment, use the `-project` flag:

```
$ gcloud app deploy ~/my_app/app.yaml -project=PROJECT Ref: https://cloud.google.com/sdk/gcloud/reference/app/deploy
```

NEW QUESTION 213

After a recent security incident, your startup company wants better insight into what is happening in the Google Cloud environment. You need to monitor unexpected firewall changes and instance creation. Your company prefers simple solutions. What should you do?

- A. Use Cloud Logging filters to create log-based metrics for firewall and instance action
- B. Monitor the changes and set up reasonable alerts.
- C. Install Kibana on a compute Instance
- D. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Pub/Sub
- E. Target the Pub/Sub topic to push messages to the Kibana instance
- F. Analyze the logs on Kibana in real time.
- G. Turn on Google Cloud firewall rules logging, and set up alerts for any insert, update, or delete events.
- H. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Cloud Storage. Use BigQuery to periodically analyze log events in the storage bucket.

Answer: A

Explanation:

This answer is the simplest and most effective way to monitor unexpected firewall changes and instance creation in Google Cloud. Cloud Logging filters allow you to specify the criteria for the log entries that you want to view or export. You can use the Logging query language to write filters based on the LogEntry fields, such as `resource.type`, `severity`, or `protoPayload.methodName`. For example, you can filter for firewall-related events by using the following query:

```
resource.type="gce_subnetwork" logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Ffirewall"
```

You can filter for instance-related events by using the following query: `resource.type="gce_instance"`

```
logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Factivity_log"
```

You can create log-based metrics from these filters to measure the rate or count of log entries that match the filter. Log-based metrics can be used to create charts and dashboards in Cloud Monitoring, or to set up alerts based on the metric values. For example, you can create an alert policy that triggers when the log-based metric for firewall changes exceeds a certain threshold in a given time interval. This way, you can get notified of any unexpected or malicious changes to your firewall rules.

Option B is incorrect because it is unnecessarily complex and costly. Installing Kibana on a compute instance requires additional configuration and maintenance. Creating a log sink to forward Cloud Audit Logs to Pub/Sub also incurs additional charges for the Pub/Sub service. Analyzing the logs on Kibana in real time may not be feasible or efficient, as it requires constant monitoring and manual intervention.

Option C is incorrect because Google Cloud firewall rules logging is a different feature from Cloud Audit Logs. Firewall rules logging allows you to audit, verify, and analyze the effects of your firewall rules by creating connection records for each rule that applies to traffic. However, firewall rules logging does not log the insert, update, or delete events for the firewall rules themselves. Those events are logged by Cloud Audit Logs, which record the administrative activities in your Google Cloud project.

Option D is incorrect because it is not a real-time solution. Creating a log sink to forward Cloud Audit Logs to Cloud Storage requires additional storage space and charges. Using BigQuery to periodically analyze log events in the storage bucket also incurs additional costs for the BigQuery service. Moreover, this option does not provide any alerting mechanism to notify you of any unexpected or malicious changes to your firewall rules or instances.

NEW QUESTION 217

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 221

Your Dataproc cluster runs in a single Virtual Private Cloud (VPC) network in a single subnet with range 172.16.20.128/25. There are no private IP addresses available in the VPC network. You want to add new VMs to communicate with your cluster using the minimum number of steps. What should you do?

- A. Modify the existing subnet range to 172.16.20.0/24.
- B. Create a new Secondary IP Range in the VPC and configure the VMs to use that range.
- C. Create a new VPC network for the VM
- D. Enable VPC Peering between the VMs' VPC network and the Dataproc cluster VPC network.
- E. Create a new VPC network for the VMs with a subnet of 172.32.0.0/16. Enable VPC network Peering between the Dataproc VPC network and the VMs VPC network
- F. Configure a custom Route exchange.

Answer: A

Explanation:

/25:

CIDR to IP Range Result

CIDR Range 172.16.20.128/25 Netmask 255.255.255.128

Wildcard Bits 0.0.0.127

First IP 172.16.20.128

First IP (Decimal) 2886734976 Last IP 172.16.20.255

Last IP (Decimal) 2886735103 Total Host 128

CIDR 172.16.20.128/25

/24:

CIDR to IP Range Result

CIDR Range 172.16.20.128/24 Netmask 255.255.255.0
Wildcard Bits 0.0.0.255
First IP 172.16.20.0
First IP (Decimal) 2886734848 Last IP 172.16.20.255
Last IP (Decimal) 2886735103 Total Host 256
CIDR 172.16.20.128/24

NEW QUESTION 223

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

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

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud> https://cloud.google.com/sdk/docs/configurations#multiple_configurations

NEW QUESTION 228

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 command.
- D. Then use the gcloud compute ssh command.
- E. Create a key with the ssh-keygen command.
- F. Upload the key to the instance.
- 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 229

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

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

Answer: C

Explanation:

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

BigQuery Data Viewer (roles/bigquery.dataViewer)

When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.

This role cannot be applied to individual models or routines. When applied to a dataset, this role provides permissions to: Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables.

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

Lowest-level resources where you can grant this role: Table

View

BigQuery Job User (roles/bigquery.jobUser)

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

Lowest-level resources where you can grant this role:

Project

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

<https://cloud.google.com/spanner/docs/iam#spanner.databaseUser>

NEW QUESTION 230

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

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

Answer: D

NEW QUESTION 234

You have a managed instance group comprised of preemptible VM's. All of the VM's keep deleting and recreating themselves every minute. What is a possible cause of this behavior?

- A. Your zonal capacity is limited, causing all preemptible VM's to be shutdown to recover capacity.
- B. Try deploying your group to another zone.
- C. You have hit your instance quota for the region.
- D. Your managed instance group's VM's are toggled to only last 1 minute preemptible settings.
- E. Your managed instance group's health check is repeatedly failing, either to a misconfigured health check or misconfigured firewall rules not allowing the healthcheck to access the instance.

Answer: D

Explanation:

as the instances (normal or preemptible) would be terminated and relaunched if the health check fails either due to application not configured properly or the instances firewall do not allow health check to happen.

GCP provides health check systems that connect to virtual machine (VM) instances on a configurable, periodic basis. Each connection attempt is called a probe. GCP records the success or failure of each probe.

Health checks and load balancers work together. Based on a configurable number of sequential successful or failed probes, GCP computes an overall health state for each VM in the load balancer. VMs that respond successfully for the configured number of times are considered healthy. VMs that fail to respond successfully for a separate number of times are unhealthy.

GCP uses the overall health state of each VM to determine its eligibility for receiving new requests. In addition to being able to configure probe frequency and health state thresholds, you can configure the criteria that define a successful probe.

NEW QUESTION 238

You need to select and configure compute resources for a set of batch processing jobs. These jobs take around 2 hours to complete and are run nightly. You want to minimize service costs. What should you do?

- A. Select Google Kubernetes Engine
- B. Use a single-node cluster with a small instance type.
- C. Select Google Kubernetes Engine
- D. Use a three-node cluster with micro instance types.
- E. Select Compute Engine
- F. Use preemptible VM instances of the appropriate standard machine type.
- G. Select Compute Engine
- H. Use VM instance types that support micro bursting.

Answer: C

Explanation:

If your apps are fault-tolerant and can withstand possible instance preemptions, then preemptible instances can reduce your Compute Engine costs significantly. For example, batch processing jobs can run on preemptible instances. If some of those instances stop during processing, the job slows but does not completely stop. Preemptible instances complete your batch processing tasks without placing additional workload on your existing instances and without requiring you to pay full price for additional normal instances.

<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 242

You are migrating a business critical application from your local data center into Google Cloud. As part of your high-availability strategy, you want to ensure that any data used by the application will be immediately available if a zonal failure occurs. What should you do?

- A. Store the application data on a zonal persistent disk.
- B. Create a snapshot schedule for the disk.
- C. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
- D. Store the application data on a zonal persistent disk.
- E. If an outage occurs, create an instance in another zone with this disk attached.
- F. Store the application data on a regional persistent disk.
- G. Create a snapshot schedule for the disk.
- H. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
- I. Store the application data on a regional persistent disk. If an outage occurs, create an instance in another zone with this disk attached.

Answer: A

NEW QUESTION 247

You want to set up a Google Kubernetes Engine cluster. Verifiable node identity and integrity are required for the cluster, and nodes cannot be accessed from the internet. You want to reduce the operational cost of managing your cluster, and you want to follow Google-recommended practices. What should you do?

- A. Deploy a private autopilot cluster.
- B. Deploy a public autopilot cluster.

- C. Deploy a standard public cluster and enable shielded nodes.
- D. Deploy a standard private cluster and enable shielded nodes.

Answer: D

NEW QUESTION 251

You are running an application on multiple virtual machines within a managed instance group and have autoscaling enabled. The autoscaling policy is configured so that additional instances are added to the group if the CPU utilization of instances goes above 80%. VMs are added until the instance group reaches its maximum limit of five VMs or until CPU utilization of instances lowers to 80%. The initial delay for HTTP health checks against the instances is set to 30 seconds. The virtual machine instances take around three minutes to become available for users. You observe that when the instance group autoscales, it adds more instances than necessary to support the levels of end-user traffic. You want to properly maintain instance group sizes when autoscaling. What should you do?

- A. Set the maximum number of instances to 1.
- B. Decrease the maximum number of instances to 3.
- C. Use a TCP health check instead of an HTTP health check.
- D. Increase the initial delay of the HTTP health check to 200 seconds.

Answer: D

Explanation:

The reason is that when you do health check, you want the VM to be working. Do the first check after initial setup time of 3 mins = 180 s < 200 s is reasonable.

➤ The reason why our autoscaling is adding more instances than needed is that it checks 30 seconds after launching the instance and at this point, the instance isn't up and isn't ready to serve traffic. So our autoscaling policy starts another instance again checks this after 30 seconds and the cycle repeats until it gets to the maximum instances or the instances launched earlier are healthy and start processing traffic which happens after 180 seconds (3 minutes). This can be easily rectified by adjusting the initial delay to be higher than the time it takes for the instance to become available for processing traffic. So setting this to 200 ensures that it waits until the instance is up (around 180-second mark) and then starts forwarding traffic to this instance. Even after a cool out period, if the CPU utilization is still high, the autoscaler can again scale up but this scale-up is genuine and is based on the actual load.

Initial Delay Seconds This setting delays autohealing from potentially prematurely recreating the instance if the instance is in the process of starting up. The initial delay timer starts when the currentAction of the instance is

VERIFYING. Ref: <https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs>

NEW QUESTION 254

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your Associate-Cloud-Engineer Exam with Our Prep Materials Via below:

<https://www.certleader.com/Associate-Cloud-Engineer-dumps.html>