



Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

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 company has a single sign-on (SSO) identity provider that supports Security Assertion Markup Language (SAML) integration with service providers. Your company has users in Cloud Identity. You would like users to authenticate using your company's SSO provider. What should you do?

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

Answer: B

Explanation:

https://support.google.com/cloudidentity/answer/6262987?hl=en&ref_topic=7558767

NEW QUESTION 3

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

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

Answer: A

NEW QUESTION 4

You have a developer laptop with the Cloud SDK installed on Ubuntu. The Cloud SDK was installed from the Google Cloud Ubuntu package repository. You want to test your application locally on your laptop with Cloud Datastore. What should you do?

- A. Export Cloud Datastore data using `gcloud datastore export`.
- B. Create a Cloud Datastore index using `gcloud datastore indexes create`.
- C. Install the `google-cloud-sdk-datastore-emulator` component using the `apt get install` command.
- D. Install the `cloud-datastore-emulator` component using the `gcloud components install` command.

Answer: D

Explanation:

- The Datastore emulator provides local emulation of the production Datastore environment. You can use the emulator to develop and test your application locally. Ref: <https://cloud.google.com/datastore/docs/tools/datastore-emulator>

NEW QUESTION 5

You have an on-premises data analytics set of binaries that processes data files in memory for about 45 minutes every midnight. The sizes of those data files range from 1 gigabyte to 16 gigabytes. You want to migrate this application to Google Cloud with minimal effort and cost. What should you do?

- A. Upload the code to Cloud Function
- B. Use Cloud Scheduler to start the application.
- C. Create a container for the set of binarie
- D. Use Cloud Scheduler to start a Cloud Run job for the container.
- E. Create a container for the set of binaries Deploy the container to Google Kubernetes Engine (GKE) and use the Kubernetes scheduler to start the application.
- F. Lift and shift to a VM on Compute Engin
- G. Use an instance schedule to start and stop the instance.

Answer: B

NEW QUESTION 6

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

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

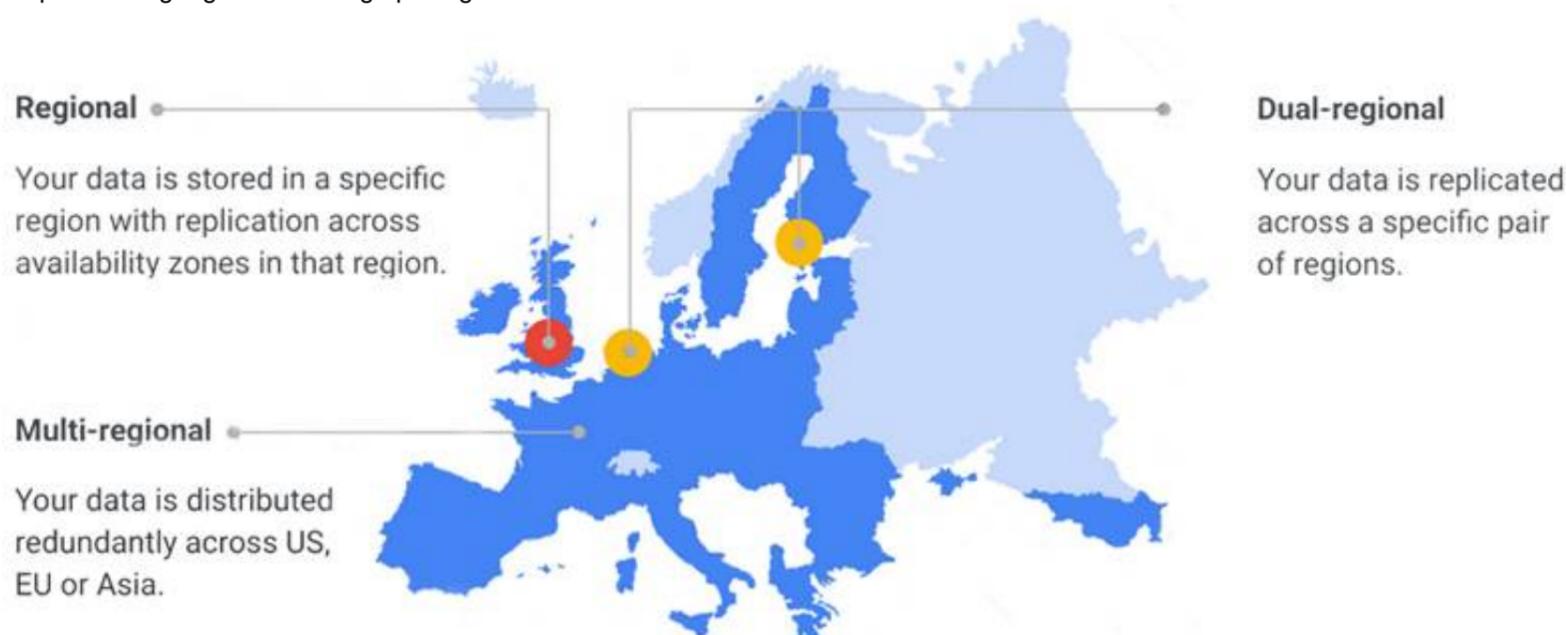
Answer: C

Explanation:

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

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

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



NEW QUESTION 7

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 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

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

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

Answer: D

Explanation:

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

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

Email/password

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

OIDC

Phone number Custom Anonymous

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

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

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

NEW QUESTION 10

You have an application on a general-purpose Compute Engine instance that is experiencing excessive disk read throttling on its Zonal SSD Persistent Disk. The application primarily reads large files from disk. The disk size is currently 350 GB. You want to provide the maximum amount of throughput while minimizing costs. What should you do?

- A. Increase the size of the disk to 1 TB.
- B. Increase the allocated CPU to the instance.
- C. Migrate to use a Local SSD on the instance.
- D. Migrate to use a Regional SSD on the instance.

Answer: C

Explanation:

Standard persistent disks are efficient and economical for handling sequential read/write operations, but they aren't optimized to handle high rates of random input/output operations per second (IOPS). If your apps require high rates of random IOPS, use SSD persistent disks. SSD persistent disks are designed for single-digit millisecond latencies. Observed latency is application specific.

NEW QUESTION 15

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

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

Answer: C

Explanation:

"Billing Account Viewer access would usually be granted to finance teams, it provides access to spend information, but does not confer the right to link or unlink projects or otherwise manage the properties of the billing account." <https://cloud.google.com/billing/docs/how-to/billing-access>

NEW QUESTION 16

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 18

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

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

Answer: A

NEW QUESTION 21

You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

- A. * 1. Create a consumer Gmail account.* 2. Write a script that monitors the CPU usage.* 3. When the CPU usage exceeds the threshold, have that script send an email using the Gmail account and smtp.gmail.com on port 25 as SMTP server.
- B. * 1. Create a Stackdriver Workspace, and associate your Google Cloud Platform (GCP) project with it.* 2. Create an Alerting Policy in Stackdriver that uses the threshold as a trigger condition
- C. 3. Configure your email address in the notification channel.
- D. * 1. Create a Stackdriver Workspace, and associate your GCP project with it.* 2. Write a script that monitors the CPU usage and sends it as a custom metric to Stackdriver
- E. 3. Create an uptime check for the instance in Stackdriver.
- F. * 1. In Stackdriver Logging, create a logs-based metric to extract the CPU usage by using this regular expression: CPU Usage: ([0-9]{1,3}) %* 2. In Stackdriver Monitoring, create an Alerting Policy based on this metric
- G. 3. Configure your email address in the notification channel.

Answer: B

Explanation:

Specifying conditions for alerting policies This page describes how to specify conditions for alerting policies. The conditions for an alerting policy define what is monitored and when to trigger an alert. For example, suppose you want to define an alerting policy that emails you if the CPU utilization of a Compute Engine VM instance is above 80% for more than 3 minutes. You use the conditions dialog to specify that you want to monitor the CPU utilization of a Compute Engine VM instance, and that you want an alerting policy to trigger when that utilization is above 80% for 3 minutes. <https://cloud.google.com/monitoring/alerts/ui-conditions-ga> <https://cloud.google.com/monitoring/alerts/using-alerting-ui> <https://cloud.google.com/monitoring/support/notification-options>

NEW QUESTION 22

You have files in a Cloud Storage bucket that you need to share with your suppliers. You want to restrict the time that the files are available to your suppliers to 1 hour. You want to follow Google recommended practices. What should you do?

- A. Create a service account with just the permissions to access files in the bucket
- B. Create a JSON key for the service account
- C. Execute the command `gsutil signurl -m 1h gs://**`.
- D. Create a service account with just the permissions to access files in the bucket
- E. Create a JSON key for the service account
- F. Execute the command `gsutil signurl -d 1h gs://**`.
- G. Create a service account with just the permissions to access files in the bucket
- H. Create a JSON key for the service account
- I. Execute the command `gsutil signurl -p 60m gs://`.
- J. Create a JSON key for the Default Compute Engine Service Account
- K. Execute the command `gsutil signurl -t 60m gs://***`

Answer: B

Explanation:

This command correctly specifies the duration that the signed url should be valid for by using the `-d` flag. The default is 1 hour so omitting the `-d` flag would have also resulted in the same outcome. Times may be specified with no suffix (default hours), or with `s` = seconds, `m` = minutes, `h` = hours, `d` = days. The max duration allowed is 7d. Ref: <https://cloud.google.com/storage/docs/gsutil/commands/signurl>

NEW QUESTION 25

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 27

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

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

Answer: B

Explanation:

Best practices for persistent disk snapshots

You can create persistent disk snapshots at any time, but you can create snapshots more quickly and with greater reliability if you use the following best practices.

Creating frequent snapshots efficiently

Use snapshots to manage your data efficiently.

Create a snapshot of your data on a regular schedule to minimize data loss due to unexpected failure. Improve performance by eliminating excessive snapshot downloads and by creating an image and reusing it. Set your snapshot schedule to off-peak hours to reduce snapshot time.

Snapshot frequency limits

Creating snapshots from persistent disks

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

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

NEW QUESTION 28

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

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

Answer: C

Explanation:

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

NEW QUESTION 30

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 34

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

- A. Configure username and password by using gcloud configure set proxy/username and gcloud configure set proxy/ proxy/password commands.
- B. Encode username and password in sha256 encoding, and save it to a text fil
- C. Use filename as a value in the gcloud configure set core/custom_ca_certs_file command.

- D. Provide values for CLOUDSDK_USERNAME and CLOUDSDK_PASSWORD in the gcloud CLI tool configure file.
- E. Set the CLOUDSDK_PROXY_USERNAME and CLOUDSDK_PROXY_PASSWORD properties by using environment variables in your command line tool.

Answer: D

NEW QUESTION 36

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 41

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 43

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

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

Answer: C

NEW QUESTION 47

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

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

Answer: A

Explanation:

VMs cannot communicate over the internet without a public IP address. Private Google Access permits access to Google APIs and services in Google's production infrastructure.

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

NEW QUESTION 49

Your team is running an on-premises ecommerce application. The application contains a complex set of microservices written in Python, and each microservice is running on Docker containers. Configurations are injected by using environment variables. You need to deploy your current application to a serverless Google Cloud solution. What should you do?

- A. Use your existing CI/CD pipeline. Use the generated Docker images and deploy them to Cloud Run. Update the configurations and the required endpoints.
- B. Use your existing continuous integration and delivery (CI/CD) pipeline.
- C. Use the generated Docker images and deploy them to Cloud Functions.
- D. Use the same configuration as on-premises.
- E. Use the existing codebase and deploy each service as a separate Cloud Function. Update the configurations and the required endpoints.
- F. Use your existing codebase and deploy each service as a separate Cloud Run. Use the same configurations as on-premises.

Answer: A

NEW QUESTION 54

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 57

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 58

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 63

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

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

Answer: C

Explanation:

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

* 1. Create a service account. 2. Give the Cloud Run Invoker role to that service account for your Cloud Run application. 3. Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint.

NEW QUESTION 66

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 67

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

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

Answer: C

Explanation:

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

Data Access audit logs Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data.

<https://cloud.google.com/logging/docs/audit#data-access>

NEW QUESTION 68

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 isnt going to work.

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

NEW QUESTION 73

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

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

Answer: A

Explanation:

IP 10.0.3.21 is internal by default, and to ensure that it will be static non-changing it should be selected as static internal ip address.

NEW QUESTION 78

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 project
- F. Create two new VPCs and add all instances.
- G. Verify that you are the Project Administrator of both project
- 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 83

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 85

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 89

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

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

Answer: C

NEW QUESTION 92

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

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

Answer: A

Explanation:

What is Google Cloud Memorystore?

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

NEW QUESTION 95

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:

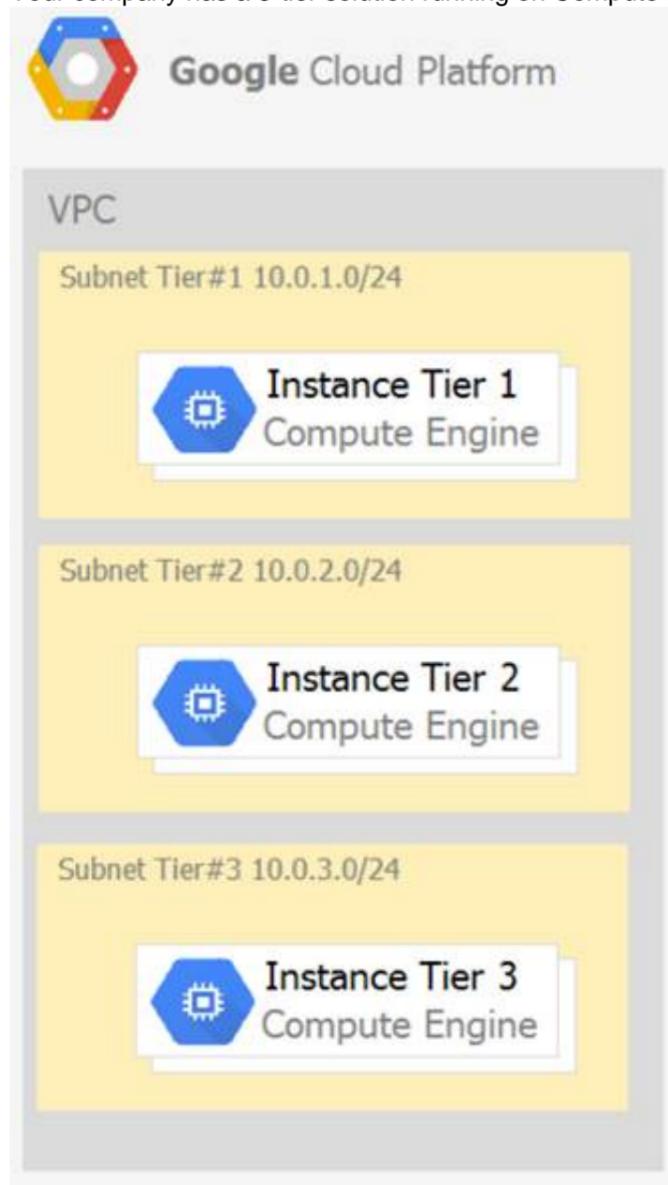
From this source:

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 100

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



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

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

What should you do?

- A. 1. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)•Protocols: allow all
- B. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow TCP:80802. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter:

all instances with tier #2 service account• Protocols: allow TCP: 8080

- C. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow all
2. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow all
- D. 1. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow TCP: 8080
2. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)• Protocols: allow TCP: 8080

Answer: B

Explanation:

* 1. Create an ingress firewall rule with the following settings: "☐ Targets: all instances with tier #2 service account "☐ Source filter: all instances with tier #1 service account "☐ Protocols: allow TCP:8080
2. Create an ingress firewall rule with the following settings: "☐ Targets: all instances with tier #3 service account "☐ Source filter: all instances with tier #2 service account "☐ Protocols: allow TCP: 8080

NEW QUESTION 103

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

- A. Fill in local SS
- B. Fill in persistent disk storage and snapshot storage.
- C. Fill in local SS
- D. Add estimated cost for cluster management.
- E. Select Add GPU
- F. Fill in persistent disk storage and snapshot storage.
- G. Select Add GPU
- H. Add estimated cost for cluster management.

Answer: A

Explanation:

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

NEW QUESTION 104

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

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

Answer: B

Explanation:

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

NEW QUESTION 106

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

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

Answer: B

NEW QUESTION 108

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

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

Answer: C

NEW QUESTION 112

You are building a new version of an application hosted in an App Engine environment. You want to test the new version with 1% of users before you completely switch your application over to the new version. What should you do?

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

- B. Deploy a new version of your application in a Compute Engine instance instead of App Engine and then use GCP Console to split traffic.
- C. Deploy a new version as a separate app in App Engine.
- D. Then configure App Engine using GCP Console to split traffic between the two apps.
- E. Deploy a new version of your application in App Engine.
- F. Then go to App Engine settings in GCP Console and split traffic between the current version and newly deployed versions accordingly.

Answer: D

Explanation:

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

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

NEW QUESTION 117

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

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

Answer: A

NEW QUESTION 121

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 to 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 125

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

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

Answer: D

Explanation:

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

NEW QUESTION 127

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 130

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 132

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

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

Answer: D

Explanation:

"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

<https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4>

NEW QUESTION 137

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 138

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

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

Answer: D

Explanation:

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

NEW QUESTION 143

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 146

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

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

Answer: A

Explanation:

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

NEW QUESTION 149

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 154

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

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

Answer: B

Explanation:

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

NEW QUESTION 157

You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

- A. Enable the Identity Aware Proxy API on the project.
- B. Scan the bucket using the Data Loss Prevention API.
- C. Allow only a single Service Account access to read the data.
- D. Enable Data Access audit logs for the Cloud Storage API.

Answer: D

Explanation:

Logged information Within Cloud Audit Logs, there are two types of logs: Admin Activity logs: Entries for operations that modify the configuration or metadata of a project, bucket, or object. Data Access logs: Entries for operations that modify objects or read a project, bucket, or object. There are several sub-types of data access logs: ADMIN_READ: Entries for operations that read the configuration or metadata of a project, bucket, or object. DATA_READ: Entries for operations that read an object. DATA_WRITE: Entries for operations that create or modify an object. <https://cloud.google.com/storage/docs/audit-logs#types>

NEW QUESTION 160

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 163

The DevOps group in your organization needs full control of Compute Engine resources in your development project. However, they should not have permission to create or update any other resources in the project. You want to follow Google's recommendations for setting permissions for the DevOps group. What should you do?

- A. Grant the basic role roles/viewer and the predefined role roles/compute.admin to the DevOps group.
- B. Create an IAM policy and grant all compute instanceAdmin permissions to the policy. Attach the policy to the DevOps group.
- C. Create a custom role at the folder level and grant all compute instanceAdmin permissions to the role. Grant the custom role to the DevOps group.
- D. Grant the basic role roles/editor to the DevOps group.

Answer: A

NEW QUESTION 166

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

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

Answer: D

Explanation:

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

NEW QUESTION 171

You created a cluster.YAML file containing

- > resources:
- > name: cluster
- > type: container.v1.cluster
- > properties:
- > zone: europe-west1-b
- > cluster:
- > description: My GCP ACE cluster

> initialNodeCount: 2

You want to use Cloud Deployment Manager to create this cluster in GKE. What should you do?

- A. gcloud deployment-manager deployments create my-gcp-ace-cluster --config cluster.yaml
- B. gcloud deployment-manager deployments create my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- C. gcloud deployment-manager deployments apply my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- D. gcloud deployment-manager deployments apply my-gcp-ace-cluster --config cluster.yaml

Answer: D

Explanation:

gcloud deployment-manager deployments create creates deployments based on the configuration file. (Infrastructure as code). All the configuration related to the artifacts is in the configuration file. This command correctly creates a cluster based on the provided cluster.yaml configuration file.

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

NEW QUESTION 175

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 threshold
- 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 threshold
- 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 threshold
- 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 180

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

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

Answer: A

Explanation:

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

NEW QUESTION 184

You want to configure a solution for archiving data in a Cloud Storage bucket. The solution must be cost-effective. Data with multiple versions should be archived after 30 days. Previous versions are accessed once a month for reporting. This archive data is also occasionally updated at month-end. What should you do?

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

Answer: B

NEW QUESTION 185

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

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

Answer: C

Explanation:

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

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

BigQuery supports the following external data sources: Amazon S3

Azure Storage Cloud Bigtable Cloud Spanner Cloud SQL Cloud Storage

Drive

NEW QUESTION 186

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

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

Answer: A

NEW QUESTION 189

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 191

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

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

Answer: D

Explanation:

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

NEW QUESTION 196

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%>

<https://www.sap.com/india/products/hana.html#:~:text=is%20SAP%20HANA-,in%2Dmemory,-database%3F>

NEW QUESTION 201

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 205

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

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

Answer: A

Explanation:

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it can't be changed. The only way is to create a new project and create an App Engine instance in europe-west3, send all user traffic to this instance and delete the app engine instance in us-central.
 Ref: <https://cloud.google.com/appengine/docs/locations>

NEW QUESTION 210

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

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

Answer: C

NEW QUESTION 215

You are building a product on top of Google Kubernetes Engine (GKE). You have a single GKE cluster. For each of your customers, a Pod is running in that cluster, and your customers can run arbitrary code inside their Pod. You want to maximize the isolation between your customers' Pods. What should you do?

- A. Use Binary Authorization and whitelist only the container images used by your customers' Pods.
- B. Use the Container Analysis API to detect vulnerabilities in the containers used by your customers' Pods.
- C. Create a GKE node pool with a sandbox type configured to gvisor
- D. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.
- E. Use the cos_containerd image for your GKE node
- F. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Answer: C

NEW QUESTION 219

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

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

Answer: D

Explanation:

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

NEW QUESTION 224

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

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

Answer: C

Explanation:

<https://blog.doit-intl.com/the-truth-behind-google-cloud-egress-traffic-6e8f57b5c2f8>

NEW QUESTION 227

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

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

Answer: A

Explanation:

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

NEW QUESTION 232

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

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

Answer: D

NEW QUESTION 236

Your company requires all developers to have the same permissions, regardless of the Google Cloud project they are working on. Your company's security policy also restricts developer permissions to Compute Engine.

Cloud Functions, and Cloud SQL. You want to implement the security policy with minimal effort. What should you do?

- A. • Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions in one project within the Google Cloud organization. • Copy the role across all projects created within the organization with the `gcloud iam roles copy` command. • Assign the role to developers in those projects.
- B. • Add all developers to a Google group in Google Groups for Workspace. • Assign the predefined role of Compute Admin to the Google group at the Google Cloud organization level.
- C. • Add all developers to a Google group in Cloud Identity. • Assign predefined roles for Compute Engine, Cloud Functions, and Cloud SQL permissions to the Google group for each project in the Google Cloud organization.
- D. • Add all developers to a Google group in Cloud Identity. • Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions at the Google Cloud organization level. • Assign the custom role to the Google group.

Answer: D

Explanation:

<https://www.cloudskillsboost.google/focuses/1035?parent=catalog#:~:text=custom%20role%20at%20the%20or>

NEW QUESTION 239

You have deployed an application on a single Compute Engine instance. The application writes logs to disk. Users start reporting errors with the application. You want to diagnose the problem. What should you do?

- A. Navigate to Cloud Logging and view the application logs.
- B. Connect to the instance's serial console and read the application logs.
- C. Configure a Health Check on the instance and set a Low Healthy Threshold value.

D. Install and configure the Cloud Logging Agent and view the logs from Cloud Logging.

Answer: D

NEW QUESTION 240

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 243

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 247

An application generates daily reports in a Compute Engine virtual machine (VM). The VM is in the project corp-iot-insights. Your team operates only in the project corp-aggregate-reports and needs a copy of the daily exports in the bucket corp-aggregate-reports-storage. You want to configure access so that the daily reports from the VM are available in the bucket corp-aggregate-reports-storage and use as few steps as possible while following Google-recommended practices. What should you do?

- A. Move both projects under the same folder.
- B. Grant the VM Service Account the role Storage Object Creator on corp-aggregate-reports-storage.
- C. Create a Shared VPC network between both projects
- D. Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
- E. Make corp-aggregate-reports-storage public and create a folder with a pseudo-randomized suffix name. Share the folder with the IoT team.

Answer: B

Explanation:

Predefined roles

The following table describes Identity and Access Management (IAM) roles that are associated with Cloud Storage and lists the permissions that are contained in each role. Unless otherwise noted, these roles can be applied either to entire projects or specific buckets.

Storage Object Creator (roles/storage.objectCreator) Allows users to create objects. Does not give permission to view, delete, or overwrite objects.

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

NEW QUESTION 249

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 251

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 252

You are creating an application that will run on Google Kubernetes Engine. You have identified MongoDB as the most suitable database system for your application and want to deploy a managed MongoDB environment that provides a support SLA. What should you do?

- A. Create a Cloud Bigtable cluster and use the HBase API
- B. Deploy MongoDB Alias from the Google Cloud Marketplace
- C. Download a MongoDB installation package and run it on Compute Engine instances
- D. Download a MongoDB installation package, and run it on a Managed Instance Group

Answer: B

Explanation:

<https://console.cloud.google.com/marketplace/details/gc-launcher-for-mongodb-atlas/mongodb-atlas>

NEW QUESTION 256

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

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

Answer: A

Explanation:

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

NEW QUESTION 260

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 262

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 263

Your continuous integration and delivery (CI/CD) server can't execute Google Cloud actions in a specific project because of permission issues. You need to validate whether the used service account has the appropriate roles in the specific project. What should you do?

- A. Open the Google Cloud console, and run a query to determine which resources this service account can access.
- B. Open the Google Cloud console, and run a query of the audit logs to find permission denied errors for this service account.
- C. Open the Google Cloud console, and check the organization policies.
- D. Open the Google Cloud console, and check the Identity and Access Management (IAM) roles assigned to the service account at the project or inherited from the folder or organization levels.

Answer: D

Explanation:

This answer is the most effective way to validate whether the service account used by the CI/CD server has the appropriate roles in the specific project. By checking the IAM roles assigned to the service account, you can see which permissions the service account has and which resources it can access. You can also check if the service account inherits any roles from the folder or organization levels, which may affect its access to the project. You can use the Google Cloud console, the gcloud command-line tool, or the IAM API to view the IAM roles of a service account.

NEW QUESTION 266

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

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

Answer: B

NEW QUESTION 268

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

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

Answer: C

Explanation:

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

<https://cloud.google.com/billing/docs/how-to/budgets#budget-scop>

NEW QUESTION 270

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

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

Answer: C

Explanation:

Using primitive roles The following table lists the primitive roles that you can grant to access a project, the description of what the role does, and the permissions bundled within that role. Avoid using primitive roles except when absolutely necessary. These roles are very powerful, and include a large number of permissions across all Google Cloud services. For more details on when you should use primitive roles, see the Identity and Access Management FAQ. IAM predefined roles

are much more granular, and allow you to carefully manage the set of permissions that your users have access to. See Understanding Roles for a list of roles that can be granted at the project level. Creating custom roles can further increase the control you have over user permissions. https://cloud.google.com/resource-manager/docs/access-control-proj#using_primitive_roles
<https://cloud.google.com/iam/docs/understanding-custom-roles>

NEW QUESTION 271

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

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

Answer: B

Explanation:

Google Cloud Storage Triggers

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

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

Event types

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

Supported trigger type values are: google.storage.object.finalize google.storage.object.delete google.storage.object.archive google.storage.object.metadataUpdate

Object Finalize

Trigger type value: google.storage.object.finalize

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

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

NEW QUESTION 272

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 in 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 health check 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 275

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 277

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

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

Answer: A

Explanation:

Adding an API as a type provider

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

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

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

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

NEW QUESTION 279

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 283

You want to configure autohealing for network load balancing for a group of Compute Engine instances that run in multiple zones, using the fewest possible steps. You need to configure re-creation of VMs if they are unresponsive after 3 attempts of 10 seconds each. What should you do?

- A. Create an HTTP load balancer with a backend configuration that references an existing instance group. Set the health check to healthy (HTTP).
- B. Create an HTTP load balancer with a backend configuration that references an existing instance group. Define a balancing mode and set the maximum RPS to 10.
- C. Create a managed instance group
- D. Set the Autohealing health check to healthy (HTTP).
- E. Create a managed instance group
- F. Verify that the autoscaling setting is on.

Answer: C

Explanation:

<https://cloud.google.com/compute/docs/instance-groups>

<https://cloud.google.com/load-balancing/docs/network/transition-to-backend-services#console>

➤ In order to enable auto-healing, you need to group the instances into a managed instance group.

Managed instance groups (MIGs) maintain the high availability of your applications by proactively keeping your virtual machine (VM) instances available. An auto-healing policy on the MIG relies on an application-based health check to verify that an application is responding as expected. If the auto-healer determines that an application isn't responding, the managed instance group automatically recreates that instance.

It is important to use separate health checks for load balancing and for auto-healing. Health checks for load balancing can and should be more aggressive because these health checks determine whether an instance receives user traffic. You want to catch non-responsive instances quickly, so you can redirect traffic if necessary. In contrast, health checking for auto-healing causes Compute Engine to proactively replace failing instances, so this health check should be more conservative than a load balancing health check.

NEW QUESTION 284

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 289

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

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

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