Answer Sheet

Q1) You have two workloads on GKE (Google Kubernetes Engine) - create-order and dispatch-order. create-order handles the creation of customer orders, and dispatch-order handles dispatching orders to your shipping partner. Both create-order and dispatch-order workloads have cluster autoscaling enabled. The create-order deployment needs to access (i.e. invoke web service of) dispatch-order deployment. dispatch-order deployment cannot be exposed publicly. How should you define the services?

- Create a Service of type LoadBalancer for dispatch-order. Have create-order use the Service IP address.
- Create a Service of type LoadBalancer for dispatch-order and an Ingress Resource for that Service. Have create-order use the Ingress IP address.
- Create a Service of type NodePort for dispatch-order and an Ingress Resource for that Service. Have create-order use the Ingress IP address.
- ✅ Create a Service of type ClusterIP for dispatch-order. Have create-order use the Service IP address.

Explanation:-ClusterIP exposes the Service on a cluster-internal IP that is only reachable within the cluster. This satisfies our requirement that dispatch-order shouldn't be publicly accessible. create-order which is also located in the same GKE cluster can now access the ClusterIP of the service to reach dispatch-order.
Ref: https://kubernetes.io/docs/concepts/services-networking/service/

Q2) Your team is working towards using the desired state configuration for your application deployed on the GKE cluster. You have YAML files for the Kubernetes Deployment and Service objects. Your application is designed to have 2 pods, which is defined by the replicas parameter in app-deployment.yaml. Your service uses GKE Load Balancer which is defined in app-service.yaml
You created the Kubernetes resources by running
kubectl apply -f app-deployment.yaml
kubectl apply -f app-service.yaml
Your deployment is now serving live traffic but is suffering from performance issues. You want to increase the number of replicas to 5. What should you do in order to update the replicas in existing Kubernetes deployment objects?

- Modify the current configuration of the deployment by using kubectl edit to open the YAML file of the current configuration, modify and save the configuration. kubectl edit deployment/app-deployment -o yaml --save-config
- Disregard the YAML file. Enable autoscaling on the deployment to trigger on CPU usage and set max pods to 5. kubectl autoscale myapp --max=5 --cpu-percent=80
- Disregard the YAML file. Use the kubectl scale command to scale the replicas to 5. kubectl scale --replicas=5 -f app-deployment.yaml
- ✅ Edit the number of replicas in the YAML file and rerun the kubectl apply. kubectl apply -f app-deployment.yaml

Explanation:-This is the only approach that guarantees that you use desired state configuration. By updating the YAML file to have 5 replicas and applying it using kubectl apply, you are preserving the intended state of Kubernetes cluster in the YAML file.
Ref: https://kubernetes.io/docs/concepts/cluster-administration/manage-deployment/#in-place-updates-of-resources

---

Q3) You created a Kubernetes deployment by running kubectl run nginx --image=nginx --replicas=1. After a few days, you decided you no longer want this deployment. You identified the pod and deleted it by running kubectl delete pod. You noticed the pod got recreated.
What should you do to delete the deployment and avoid pod getting recreated?

- kubectl delete nginx
- ✅ kubectl delete deployment nginx

Explanation:-This command correctly deletes the deployment. Pods are managed by kubernetes workloads (deployments). When a pod is deleted, the deployment detects the pod is unavailable and brings up another pod to maintain the replica count. The only way to delete the workload is by deleting the deployment itself using the kubectl delete deployment command.
$ kubectl delete deployment nginx
deployment.apps "nginx" deleted
Ref: https://kubernetes.io/docs/reference/kubectl/cheatsheet/#deleting-resources

- kubectl delete pod nginx-84748895c4-k6bzl --no-restart
- kubectl delete --deployment=nginx

Q4) You are migrating your on-premises workloads to GCP VPC, and you want to use Compute Engine virtual machines. You want to separate the Finance team VMs and the Procurement team VMs into separate subnets. You need all VMs to communicate with each other over their internal IP addresses without adding routes. What should you do?

- Use Deployment Manager to create two VPCs, each with a subnet in a different region. Ensure the subnets use non-overlapping IP range.
- ✅ Use Deployment Manager to create a new VPC with 2 subnets in 2 different regions. Ensure the subnets use non-overlapping IP range.

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

- Use Deployment Manager to create a new VPC with 2 subnets in the same region. Ensure the subnets use the same IP range.
- Use Deployment Manager to create two VPCs, each with a subnet in the same region. Ensure the subnets use overlapping IP range.

Q5) You are in the process of migrating a mission-critical application from your on-premises data centre to Google Kubernetes Engine (GKE). Your operations team do not want to take on the overhead for upgrading the GKE cluster and have asked you to ensure the Kubernetes version is always stable and supported. What should you do?

- ✅ Update your GKE cluster to turn on GKE's node auto-upgrade feature.

Explanation:-Node auto-upgrades help you keep the nodes in your cluster up to date with the cluster master version when your master is updated on your behalf. When you create a new cluster or node pool with Google Cloud Console or the gcloud command, node auto-upgrade is enabled by default.
Ref: https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades

- Update your GKE cluster to turn on GKE's node auto-repair feature.
- When provisioning the GKE cluster, use Container Optimized OS node images.
- When provisioning the GKE cluster, ensure you use the latest stable and supported version.

Q6) You want to migrate an XML parser application from the on-premises data centre to Google Cloud Platform. You created a development project, set up the necessary IAM roles and deployed the application in a compute engine instance. The testing has succeeded, and you are ready to deploy the staging instance. You want to create the same IAM roles in a new staging GCP project. How can you do this efficiently without compromising security?

- Make use of the Create Role from Role feature in GCP console to create IAM roles in the Staging project from the Development IAM roles.
- Make use of Create Role feature in GCP console to create all necessary IAM roles from new in the Staging project.
- ✅ Make use of gcloud iam roles copy command to copy the IAM roles from the Development GCP project to the Staging GCP project.

Explanation:-This option fits all the requirements. You copy the roles into the destination project using gcloud iam roles copy and by specifying the staging project destination project.
$gcloud iam roles copy --source "<>" --destination <> --dest-project <>
Ref: https://cloud.google.com/sdk/gcloud/reference/iam/roles/copy

- Make use of gcloud iam roles copy command to copy the IAM roles from the Development GCP organization to the Staging GCP organization.

Q7) You want to ensure the boot disk of a preemptible instance is persisted for re-use. How should you provision the gcloud compute instance to ensure your requirement is met.

- gcloud compute instances create [INSTANCE_NAME] --preemptible --boot-disk-auto-delete=no
- ✅ gcloud compute instances create [INSTANCE_NAME] --preemptible --no-boot-disk-auto-delete

Explanation:-Use --no-boot-disk-auto-delete to disable automatic deletion of boot disks when the instances are deleted. --boot-disk-auto-delete flag is enabled by default. It enables automatic deletion of boot disks when the instances are deleted. In order to prevent automatic deletion, we have to specify --no-boot-disk-auto-delete flag.
Ref: https://cloud.google.com/sdk/gcloud/reference/compute/instances/create

- gcloud compute instances create [INSTANCE_NAME] --no-auto-delete
- gcloud compute instances create [INSTANCE_NAME] --preemptible. The flag --boot-disk-auto-delete is disabled by default.

Q8) Your company plans to store sensitive PII data in a cloud storage bucket. Your compliance department doesn't like encrypting sensitive PII data with Google-managed keys and has asked you to ensure the new objects uploaded to this bucket are encrypted by customer-managed encryption keys. What should you do? (Select Three)

- ✅ In the bucket advanced settings, select the Customer-managed key and then select a Cloud KMS encryption key.

Explanation:-Our compliance department wants us to use customer-managed encryption keys. We can select Customer-Managed radio and provide a cloud KMS encryption key to encrypt objects with the customer-managed key. This fit our requirements.

- ✅ Use gsutil with -o "GSUtil:encryption_key=[KEY_RESOURCE]" when uploading objects to the bucket.

Explanation:-We can have gsutil use an encryption key by using the -o top-level flag: -o "GSUtil:encryption_key=[KEY_RESOURCE]". Ref: https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key

- Use gsutil with --encryption-key=[ENCRYPTION_KEY] when uploading objects to the bucket.
- ✅ Modify .boto configuration to include encryption_key = [KEY_RESOURCE] when uploading objects to bucket.

Explanation:-As an alternative to the -o top-level flag, gsutil can also use an encryption key if .boto configuration is modified to specify the encryption key. encryption_key = [KEY_RESOURCE] Ref: https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key

---

Q9) Your compliance department has asked you to share a compressed zip of sensitive audit logs with an external auditor. The external auditor does not have a Google account, and you want to remove the access after 4 hours. How can you do this securely with the least number of steps?

- ✅ Use gcloud to generate a signed URL on the object with a four-hour expiry. Securely share the URL with the external auditor.

Explanation:-This option fits all requirements. When we generate a signed URL, we can specify an expiry and only users with the signed URL can view/download the objects, and they don't need a google account.
Ref: https://cloud.google.com/storage/docs/access-control/signed-urls

- Make the zip file public and securely share the URL with the external auditor. Set up a lifecycle policy to delete the object after 4 hours.
- Configure Static Website hosting on the Cloud Storage bucket, make the zip file public and ask the auditor to download the file from the website. Delete the zip file after 4 hours.
- Copy the zip file to a new Cloud Storage bucket, make the bucket public and share the URL securely with the external auditor. Delete the new bucket after 4 hours.

Q10) You deployed a Java application on four Google Cloud Compute Engine VMs in two zones behind a network load balancer. During peak usage, the application has stuck threads. This issue ultimately takes down the whole system and requires a reboot of all VMs. Your operations team have recently heard about self-healing mechanisms in Google Cloud and have asked you to identify if it is possible to automatically recreate the VMs if they remain unresponsive for 3 attempts 10 seconds apart. What should you do?

- Use a global HTTP(s) Load Balancer instead and limit Requests Per Second (RPS) to 10.
- ✅ Enable autohealing and set the autohealing health check to healthy (HTTP).

Explanation:-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 essential to use separate health checks for load balancing and 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 replace failing instances proactively, so this health check should be more conservative than a load balancing health check.

- Enable autoscaling on the Managed Instance Group (MIG).
- Use a global HTTP(s) Load Balancer instead and set the load balancer health check to healthy (HTTP).

Q11) You have two Kubernetes resource configuration files.
1. deployment.yaml - creates a deployment

2. service.yaml - sets up a LoadBalancer service to expose the pods.

You don't have a GKE cluster in the development project and you need to provision one. Which of the commands fail with an error in Cloud Shell when you are attempting to create a GKE cluster and deploy the YAML configuration files to create a deployment and service. (Select Two)

○ gcloud container clusters create cluster-1 --zone=us-central1-a
gcloud container clusters get-credentials cluster-1 --zone=us-central1-a
kubectl apply -f deployment.yaml
kubectl apply -f service.yaml

○ gcloud config set compute/zone us-central1-a
gcloud container clusters create cluster-1
gcloud container clusters get-credentials cluster-1 --zone=us-central1-a
kubectl apply -f deployment.yaml
kubectl apply -f service.yaml

✓ gcloud container clusters create cluster-1 --zone=us-central1-a
gcloud container clusters get-credentials cluster-1 --zone=us-central1-a
kubectl apply -f deployment.yaml&&service.yaml

Explanation:-kubectl apply can apply the configuration from a single file or multiple files or even a complete directory. When applying configuration from multiple files, the file names need to be separated by a comma. In this scenario, the filenames are passed as a list and Kubernetes treats the list as literal so looks for files "[deployment.yaml" and "service.yaml]" which it doesn't find. Ref: https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#apply

✓ gcloud container clusters create cluster-1 --zone=us-central1-a
gcloud container clusters get-credentials cluster-1 --zone=us-central1-a
kubectl apply -f [deployment.yaml,service.yaml]

Explanation:-kubectl apply can apply the configuration from a single file or multiple files or even a complete directory. When applying configuration from multiple files, the file names need to be separated by a comma. In this scenario, the filenames are separated by && and kubernetes treats the && as literal so it looks for the file "deployment.yaml&&service.yaml" which it doesn't find. Ref: https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#apply

---

Q12) Your company recently acquired a startup that lets its developers pay for their projects using their company credit cards. You want to consolidate the billing of all GCP projects into a new billing account. You want to follow Google recommended practices. How should you do this?

○ Ensure you have the Billing Account Creator Role. Create a new Billing account yourself and set up a payment method with company credit card details.

✓ In the GCP Console, move all projects to the root organization in the Resource Manager.

Explanation:-If we move all projects under the root organization hierarchy, they still need to modify to use a billing account within the organization (same as the previous option).
Ref: https://cloud.google.com/resource-manager/docs/migrating-projects-billing#top_of_page

Note: The link between projects and billing accounts is preserved, irrespective of the hierarchy. When you move your existing projects into the organization, they will continue to work and be billed as they used to before the migration, even if the corresponding billing account has not been migrated yet.

But in this option, all projects are in the organization resource hierarchy so the organization can uniformly apply organization policies to all its projects which is a Google recommended practice. So this is the better of the two options.

Ref: https://cloud.google.com/billing/docs/concepts

○ Send an email to billing.support@cloud.google.com and request them to create a new billing account and link all the projects to the billing account.

○ Raise a support request with Google Billing Support and request them to create a new billing account and link all the projects to the billing account.

Q13) Your company is migrating a mission-critical application from the on-premises data centre to Google Cloud Platform. The application requires 12 Compute Engine VMs to handle traffic at peak usage times. Your operations team have asked you to ensure the VMs restart automatically (i.e. without manual intervention) if/when they crash, and the processing capacity of the application does not reduce down during system maintenance. What should you do?

✓ Create an instance template with availability policy that turns on the automatic restart behaviour and sets on-host maintenance to live migrate instances during maintenance events. Deploy the application on a Managed Instance Group (MIG) based on this template.

Explanation:-Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash. And Enabling "Migrate VM Instance" enables live migrates, i.e. compute instances are migrated during system maintenance and remain running during the migration.
Automatic Restart - If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the automaticRestart field to true. This setting does not apply if the instance is taken offline through a user action, such as calling sudo shutdown, or during a zone outage.
Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart
Enabling the Migrate VM Instance option migrates your instance away from an infrastructure maintenance event, and your instance remains running during the migration. Your instance might experience a short period of decreased performance, although generally, most instances should not notice any difference. Live migration is ideal for instances that require constant uptime and can tolerate a short period of decreased performance.
Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#live_migrate

○ Create an instance template with availability policy that turns off the automatic restart behaviour and sets on-host maintenance to terminate instances during maintenance events. Deploy the application on a Managed Instance Group (MIG) based on this template.

○ Deploy the application on a Managed Instance Group (MIG) with autohealing health check set to healthy (HTTP).

○ Deploy the application on a Managed Instance Group (MIG) that disables the creation retry mode by setting the --nocreation-retries flag.

Q14) Your manager asked you to write a script to upload objects to a Cloud Storage bucket. How should you set up the IAM access to enable the script running in a Google Compute VM upload objects to Cloud Storage?

○ Create a new IAM service account with the access scope devstorage.write_only and configure the script to use this service account.

- Grant roles/storage.objectAdmin IAM role to the service account used by the VM.
- ✅ Grant roles/storage.objectCreator IAM role to the service account used by the VM.

Explanation:-You need to provide Compute Engine instances permissions to write data into a particular Cloud Storage bucket. Storage Object Creator (roles/storage.objectCreator) allows users to create objects. Does not permit to view, delete, or overwrite objects. This permission is what the script needs to write data to the bucket. So we create a service account, add this IAM role and let the compute engine instances use this service account to write objects to the bucket.
Ref: https://cloud.google.com/storage/docs/access-control/iam-roles

- Create a new IAM service account with the access scope cloud-platform and configure the script to use this service account.

Q15) Your company, which runs highly rated mobile games, has chosen to migrate its analytics backend to BigQuery. The analytics team of 7 analysts need access to perform queries against the data in BigQuery. The analytics team members change frequently. How should you grant them access?

- Create a Cloud Identity account for each analyst and add them all to a group. Grant roles/bigquery.jobUser role to the group.
- Create a Cloud Identity account for each analyst and grant roles/bigquery.jobUser role to each account.
- ✅ Create a Cloud Identity account for each analyst and add them all to a group. Grant roles/bigquery.dataViewer role to the group.

Explanation:-dataViewer provides permissions to Read data (i.e. query) and metadata from the table or view, so this is the right role, and this option also rightly uses groups instead of assigning permissions at the user level.
Ref: https://cloud.google.com/bigquery/docs/access-control-examples

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

- Create a Cloud Identity account for each analyst and grant roles/bigquery.dataViewer role to each account.

Q16) Your company's auditors carry out an annual audit every year and have asked you to provide them with all the IAM policy changes in Google Cloud since the last audit. You want to streamline and expedite the analysis for audit. How should you share the information requested by auditors?

- ✅ Export all audit logs to BigQuery dataset. Make use of ACLs and views to restrict the data shared with the auditors. Have the auditors query the required information quickly.

Explanation:-One option exports to Google Cloud Storage (GCS) bucket whereas other exports to BigQuery. Querying information out of files in a bucket is much harder compared to querying information from BigQuery Dataset where it is as simple as running a job or set of jobs to extract just the required information and in the format required. By enabling the auditor to run jobs in Big Queries, you streamline the log extraction process, and the auditor can review the extracted logs much quicker. While as good as the other option (bucket) is, Export all audit logs to BigQuery dataset. Make use of ACLs and views to restrict the data shared with the auditors. Have the auditors query the required information quickly. is the right answer.
You need to configure log sinks before you can receive any logs, and you can't retroactively export logs that were written before the sink was created.

- Configure alerts in Cloud Monitoring and trigger notifications to the auditors.
- Export all audit logs to Google Cloud Storage bucket and set up the necessary IAM acces to restrict the data shared with auditors.
- Export all audit logs to Cloud Pub/Sub via an export sink. Use a Cloud Function to read the messages and store them in Cloud SQL. Make use of ACLs and views to restrict the data shared with the auditors.

Q17) You created a cluster.YAML file containing
You want to use Cloud Deployment Manager to create this cluster in GKE. What should you do?

- gcloud deployment-manager deployments apply my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- gcloud deployment-manager deployments create my-gcp-ace-cluster --type container.v1.cluster --config cluster.yaml
- gcloud deployment-manager deployments apply my-gcp-ace-cluster --config cluster.yaml
- ✅ gcloud deployment-manager deployments create my-gcp-ace-cluster --config cluster.yaml

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

---

Q18) You have two compute instances in the same VPC but in different regions. You can SSH from one instance to another instance using their external IP address but not their internal IP address. What could be the reason for SSH failing on the internal IP address?

- The compute instances have a static IP for their internal IP.
- The internal IP address is disabled.
- The compute instances are not using the right cross-region SSH IAM permissions
- ✅ The combination of compute instance network tags and VPC firewall rules allow SSH from 0.0.0.0 but denies SSH from the VPC subnets IP range.

Explanation:-The combination of compute instance network tags and VPC firewall rules can certainly result in SSH traffic being allowed on the external IP range but disabled from subnets IP range. The firewall rule can be configured to allow SSH traffic from 0.0.0.0/0 but deny traffic from the VPC range e.g. 10.0.0.0/8. In this case, all SSH traffic from within the VPC is denied but external SSH traffic (i.e. on external IP) is allowed.
Ref: https://cloud.google.com/vpc/docs/using-firewalls

Q19) You are enhancing a production application currently running on an Ubuntu Linux VM on Google Compute Engine. The new enhancements require a connection to Cloud SQL to persist user addresses. Your colleague has created the Cloud SQL instance and an IAM service account with the correct permissions but doesn't know how to configure the VM to use this service account, and has asked for your assistance. What should you do?

- Execute gcloud iam service-accounts keys create to generate a JSON key for the service account. Copy the contents of JSON key to ~/.identity/default-service-account.json overwrite the default service account.
- ✅ Set the service account in the Identity and API access section when provisioning the compute engine VM.

Explanation:-You can set the service account at the time of creating the compute instance. You can also update the service account used by the instance - this requires that you stop the instance first and then update the service account. Setting/Updating the service account can be done either

via the web console or by executing gcloud command or by the REST API. See below an example for updating the service account through gcloud command.
gcloud compute instances set-service-account instance-1 --zone=us-central1-a --service-account=my-new-service-account@gcloud-gcp-ace-lab-266520.iam.gserviceaccount.com
Updated [https://www.googleapis.com/compute/v1/projects/gcloud-gcp-ace-lab-266520/zones/us-central1-a/instances/instance-1].
- Execute gcloud iam service-accounts keys create to generate a JSON key for the service account. Add a metadata tag on the GCP project with key: service-account and value: .
- Execute gcloud iam service-accounts keys create to generate a JSON key for the service account. Add a metadata tag to the compute instance with key: service-account and value: .

Q20) You developed an application that reads objects from a cloud storage bucket. You followed GCP documentation and created a service account with just the permissions to read objects from the cloud storage bucket. However, when your application uses this service account, it fails to read objects from the bucket. You suspect this might be an issue with the permissions assigned to the service account. You would like to authenticate a gsutil session with the service account credentials, reproduce the issue yourself and identify the root cause. How can you authenticate gsutil with service account credentials?

- ✅ Create JSON keys for the service account and execute gcloud auth activate-service-account --key-file [KEY_FILE]

Explanation:-This command correctly authenticates access to Google Cloud Platform with a service account using its JSON key file. To allow gcloud (and other tools in Cloud SDK) to use service account credentials to make requests, use this command to import these credentials from a file that contains a private authorization key, and activate them for use in gcloud
Ref: https://cloud.google.com/sdk/gcloud/reference/auth/activate-service-account

- Create JSON keys for the service account and execute gcloud authenticate service-account --key-file [KEY_FILE]
- Create JSON keys for the service account and execute gcloud auth service-account --key-file [KEY_FILE]
- Create JSON keys for the service account and execute gcloud authenticate activate-service-account --key-file [KEY_FILE]

Q21) You have a number of applications that have bursty workloads and are heavily dependent on topics to decouple publishing systems from consuming systems. Your company would like to go serverless to enable developers to focus on writing code without worrying about infrastructure. Your solution architect has already identified Cloud Pub/Sub as a suitable alternative for decoupling systems. You have been asked to identify a suitable GCP Serverless service that is easy to use with Cloud Pub/Sub. You want the ability to scale down to zero when there is no traffic in order to minimize costs. You want to follow Google recommended practices. What should you suggest?

- App Engine Standard
- ✅ Cloud Functions

Explanation:-Cloud Functions is Google Cloud's event-driven serverless compute platform that lets you run your code locally or in the cloud without having to provision servers. Cloud Functions scales up or down, so you pay only for compute resources you use. Cloud Functions have excellent integration with Cloud Pub/Sub, lets you scale down to zero and is recommended by Google as the ideal serverless platform to use when dependent on Cloud Pub/Sub.
"If you're building a simple API (a small set of functions to be accessed via HTTP or Cloud Pub/Sub), we recommend using Cloud Functions."
Ref: https://cloud.google.com/serverless-options

- Cloud Run for Anthos
- Cloud Run

Q22) You are the Cloud Security Manager at your company, and you want to review IAM users and their assigned roles in the production GCP project. You want to follow Google recommended practices. What should you do?

- Check the output of gcloud iam service-accounts list command.
- ✅ Review the information in the IAM section for the production GCP project in Google Cloud Console.

Explanation:-This option that lets us view roles as well as users (members).
Ref: https://cloud.google.com/iam/docs/overview
See the screenshot below.
A member can be a Google Account (for end-users), a service account (for apps and virtual machines), a Google group, or a G Suite or Cloud Identity domain that can access a resource. The identity of a member is an email address associated with a user, service account, or Google group; or a domain name associated with G Suite or Cloud Identity domains.

- Check the output of gcloud iam roles list command.
- Review the information in the Roles section for the production GCP project in Google Cloud Console.

Q23) An application that you are migrating to Google Cloud relies on overnight batch jobs that take between 2 to 3 hours to complete. You want to do this at a minimal cost. Where should you run these batch jobs?

- Run the batch jobs in a non-preemptible shared core compute engine instance that supports short periods of bursting.
- ✅ Run the batch jobs in a preemptible compute engine instance of appropriate machine type.

Explanation:-We minimize the cost by selecting a preemptible instance of the appropriate type. If the preemptible instance is terminated, the next nightly run picks up the unprocessed volume.

- Run the batch jobs in a GKE cluster on a node pool with a single instance of type e2-small.
- Run the batch jobs in a GKE cluster on a node pool with four instances of type f1-micro.

Q24) A company wants to build an application that stores images in a Cloud Storage bucket and wants to generate thumbnails as well as resize the images. They want to use a google managed service that can scale up and scale down to zero automatically with minimal effort. You have been asked to recommend a service. Which GCP service would you suggest?

- Google Kubernetes Engine
- Google App Engine
- ✅ Cloud Functions

Explanation:-Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.
Ref: https://cloud.google.com/functions

Google Compute Engine

Q25) Your company collects and stores CCTV footage videos in raw format in Google Cloud Storage. Within the first 30 days, the footage is processed regularly for detecting patterns such as threat/object/face detection and suspicious behavior detection. You want to minimize the cost of storing all the data in Google Cloud. How should you store the videos?

Use Google Cloud Nearline Storage for the first 30 days, and use lifecycle rules to transition to Coldline Storage.
✅ Use Google Cloud Regional Storage for the first 30 days, and use lifecycle rules to transition to Coldline Storage.
Explanation:-We save the videos initially in Regional Storage (Standard) which does not have retrieval charges so we do not pay for accessing data within the first 30 days during which the videos are accessed frequently. We only pay for the standard storage costs. After 30 days, we transition the CCTV footage videos to Coldline storage which is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is a better choice than Standard Storage or Nearline Storage in scenarios where slightly lower availability, a 90-day minimum storage duration, and higher costs for data access are acceptable trade-offs for lowered at-rest storage costs. Coldline storage class is cheaper than Nearline storage class.
Ref: https://cloud.google.com/storage/docs/storage-classes#standard
Ref: https://cloud.google.com/storage/docs/storage-classes#coldline
Use Google Cloud Regional Storage for the first 30 days, and use lifecycle rules to transition to Nearline Storage.
Use Google Cloud Regional Storage for the first 30 days, and then move videos to Google Persistent Disk.

Q26) You want to persist logs for 10 years to comply with regulatory requirements. You want to follow Google recommended practices. Which Google Cloud Storage class should you use?

✅ Archive storage class
Explanation:-In April 2019, Google introduced a new storage class "Archive storage class" is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Google previously recommended you use Coldline storage class but the recommendation has since been updated to "Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Note, however, that for data being kept entirely for backup or archiving purposes, Archive Storage is more cost-effective, as it offers the lowest storage costs."
Ref: https://cloud.google.com/storage/docs/storage-classes#archive

Ref: https://cloud.google.com/storage/docs/storage-classes#coldline

So the correct answer is Archive storage class.
Standard storage class
Coldline storage class
Nearline storage class

Q27) You migrated an internal HR system from an on-premises database to Google Cloud Compute Engine Managed Instance Group (MIG). The networks team at your company has asked you to associate the internal DNS records of the VMs with a custom DNS zone. You want to follow Google recommended practices. What should you do?

1. Provision the VMs with custom hostnames.
1. Install a new BIND DNS server on Google Compute Engine, using the BIND name server software (BIND9).
2. Configure a Cloud DNS forwarding zone to direct all requests to the Internal BIND DNS server.
3. When provisioning the VMs, associate the DNS records with the Internal BIND DNS server.
1. Create a new Cloud DNS zone and a new VPC and associate the DNS zone with the VPC.
2. When provisioning the VMs, associate the DNS records with the new DNS zone.
3. Configure firewall rules to block all external (public) traffic.
4. Finally, configure the DNS zone associated with the default VPC to direct all requests to the new DNS zone.
✅ 1. Create a new Cloud DNS zone and set its visibility to private.
2. When provisioning the VMs, associate the DNS records with the new DNS zone.
Explanation:-You should do when you want internal DNS records in a custom zone. Cloud DNS gives you the option of private zones and internal DNS names.
Ref: https://cloud.google.com/dns/docs/overview#concepts

Q28) You have an application deployed in a GKE Cluster as a Kubernetes workload with Daemon Sets. Your application has become very popular and is now struggling to cope up with increased traffic. You want to add more pods to your workload and want to ensure your cluster scales up and scales down automatically based on volume. What should you do?

Perform a rolling update to modify machine type from n1-standard-2 to n1-standard-4.
Create another identical Kubernetes workload and split traffic between the two workloads.
Enable Horizontal Pod Autoscaling for the Kubernetes deployment.
✅ Enable autoscaling on Kubernetes Engine.
Explanation:-GKE's cluster autoscaler automatically resizes the number of nodes in a given node pool, based on the demands of your workloads. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. DaemonSets attempt to adhere to a one-Pod-per-node model.
Ref: https://cloud.google.com/kubernetes-engine/docs/concepts/cluster-autoscaler

Q29) Your Company is planning to migrate all Java web applications to Google App Engine. However, you still want to continue using your on-premise database. How can you set up the app engine to communicate with your on-premise database while minimizing effort?

Setup the application using App Engine Standard environment with Cloud Router to connect to an on-premise database.
Setup the application using App Engine Flexible environment with Cloud Router to connect to an on-premise database.
✅ Setup the application using App Engine Flexible environment with Cloud VPN to connect to an on-premise database.
Explanation:-You need Cloud VPN to connect VPC to an on-premise network.
Ref: https://cloud.google.com/vpn/docs/concepts/overview

Unlike App Engine Standard which is serverless, App Engine Flex instances are already within the VPC, so they can use Cloud VPN to connect to the on-premise network.

- Setup the application using App Engine Standard environment with Cloud VPN to connect to an on-premise database.

Q30) You want to ingest and analyze large volumes of stream data from sensors in real-time, matching the high speeds of IoT data to track normal and abnormal behavior. You want to run it through a data processing pipeline and store the results. Finally, you want to enable customers to build dashboards and drive analytics on their data in real-time. What services should you use for this task?

- Cloud Pub/Sub, Cloud Dataflow, Cloud Dataprep
- Stackdriver, Cloud Dataflow, BigQuery
- Cloud Pub/Sub, Cloud Dataflow, Cloud Dataproc
- ✅ Cloud Pub/Sub, Cloud Dataflow, BigQuery

Explanation:-You want to ingest large volumes of streaming data at high speeds. So you need to use Cloud Pub/Sub. Cloud Pub/Sub provides a simple and reliable staging location for your event data on its journey towards processing, storage, and analysis. Cloud Pub/Sub is serverless and you can ingest events at any scale.
Ref: https://cloud.google.com/pubsub
Next, you want to analyze this data. Cloud Dataflow is a fully managed streaming analytics service that minimizes latency, processing time, and cost through autoscaling and batch processing. Dataflow enables fast, simplified streaming data pipeline development with lower data latency.
Ref: https://cloud.google.com/dataflow

Next, you want to store these results. BigQuery is an ideal place to store these results as BigQuery supports the querying of streaming data in real-time. This assists in real-time predictive analytics.
Ref: https://cloud.google.com/bigquery

Therefore the correct answer is Cloud Pub/Sub, Cloud Dataflow, BigQuery

Here's more information from Google docs about the Stream analytics use case. Google recommends we use Dataflow along with Pub/Sub and BigQuery.
https://cloud.google.com/dataflow#section-6
Google's stream analytics makes data more organized, useful, and accessible from the instant it's generated. Built on Dataflow along with Pub/Sub and BigQuery, our streaming solution provisions the resources you need to ingest, process, and analyze fluctuating volumes of real-time data for real-time business insights. This abstracted provisioning reduces complexity and makes stream analytics accessible to both data analysts and data engineers.

and

https://cloud.google.com/solutions/stream-analytics

Ingest, process, and analyze event streams in real time. Stream analytics from Google Cloud makes data more organized, useful, and accessible from the instant it's generated. Built on the autoscaling infrastructure of Pub/Sub, Dataflow, and BigQuery, our streaming solution provisions the resources you need to ingest, process, and analyze fluctuating volumes of real-time data for real-time business insights.

Q31) You have three gcloud configurations - one for each of development, test and production projects. You want to list all the configurations and switch to a new configuration. With the fewest steps possible, what's the fastest way to switch to the correct configuration?

- To list configurations - gcloud config list
To activate a configuration - gcloud config activate.
- ✅ To list configurations - gcloud config configurations list
To activate a configuration - gcloud config configurations activate.

Explanation:-The two commands together achieve the intended outcome. gcloud config configurations list - lists existing named configurations and gcloud config configurations activate - activates an existing named configuration
Ref: https://cloud.google.com/sdk/gcloud/reference/config/configurations/list

Ref: https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate

See an example below

$ gcloud config configurations list
NAME IS_ACTIVE ACCOUNT PROJECT DEFAULT_ZONE DEFAULT_REGION
dev-configuration False gcp-ace-lab-dev
prod-configuration False gcp-ace-lab-prod
test-configuration True gcp-ace-lab-test

$ gcloud config configurations activate prod-configuration
Activated [prod-configuration].

$ gcloud config configurations list
NAME IS_ACTIVE ACCOUNT PROJECT DEFAULT_ZONE DEFAULT_REGION
dev-configuration False gcp-ace-lab-dev
prod-configuration True gcp-ace-lab-prod
test-configurationFalsegcp-ace-lab-test

- To list configurations - gcloud configurations list
To activate a configuration - gcloud config activate.
- To list configurations - gcloud configurations list
To activate a configuration - gcloud configurations activate

Q32) You want to migrate an application from Google App Engine Standard to Google App Engine Flex. Your application is

currently serving live traffic and you want to ensure everything is working in Google App Engine Flex before migrating all traffic. You want to minimize effort and ensure the availability of service. What should you do?

- 1. Set env: app-engine-flex in app.yaml
2. gcloud app deploy --no-promote --version=[NEW_VERSION]
3. Validate [NEW_VERSION] in App Engine Flex
4. gcloud app versions start [NEW_VERSION]
- ✅ 1. Set env: flex in app.yaml
2. gcloud app deploy --no-promote --version=[NEW_VERSION]
3. Validate [NEW_VERSION] in App Engine Flex
4. gcloud app versions migrate [NEW_VERSION]

Explanation:-These commands together achieve the end goal while satisfying our requirements. Setting env: flex in app.yaml and executing gcloud app deploy --no-promote --version=[NEW_VERSION] results in a new version deployed to flex engine. but the new version is not configured to serve traffic. We take the opportunity to review this version before migrating it to serve live traffic by running gcloud app versions migrate [NEW_VERSION]
Ref: https://cloud.google.com/sdk/gcloud/reference/app/versions/migrate
Ref: https://cloud.google.com/sdk/gcloud/reference/app/deploy

- 1. Set env: flex in app.yaml
2. gcloud app deploy --version=[NEW_VERSION]
3. Validate [NEW_VERSION] in App Engine Flex
4. gcloud app versions migrate [NEW_VERSION]
- 1. Set env: app-engine-flex in app.yaml
2. gcloud app deploy --version=[NEW_VERSION]
3. Validate [NEW_VERSION] in App Engine Flex
4. gcloud app versions start [NEW_VERSION]

Q33) You've created a Kubernetes engine cluster named "my-gcp-ace-proj-1", which has a cluster pool named my-gcp-ace-primary-node-pool. You want to increase the number of nodes within your cluster pool from 10 to 20 to meet capacity demands. What is the command to change the number of nodes in your pool?

- ✅ gcloud container clusters resize my-gcp-ace-proj-1 --node-pool my-gcp-ace-primary-node-pool --num-nodes 20

Explanation:-gcloud container clusters resize can be used to specify the number of nodes using the --num-nodes parameter which is the target number of nodes in the cluster.
Ref: https://cloud.google.com/sdk/gcloud/reference/container/clusters/resize

- kubectl container clusters update my-gcp-ace-proj-1 --node-pool my-gcp-ace-primary-node-pool --num-nodes 20
- gcloud container clusters resize my-gcp-ace-proj-1 --node-pool my-gcp-ace-primary-node-pool --new-size 20
- gcloud container clusters update my-gcp-ace-proj-1 --node-pool my-gcp-ace-primary-node-pool --num-nodes 20

Q34) You work at a large organization where each team has a distinct role. The development team can create Google Cloud projects but can't link them to a billing account – this role is reserved for the finance team, and the development team do not want finance team to make changes to their project resources. How should you configure IAM access controls to enable this?

- Grant the finance team Billing Account User (roles/billing.user) role on the billing account.
- ✅ Grant the finance team Billing Account User (roles/billing.user) role on the billing account and Project Billing Manager (roles/billing.projectManager) on the GCP organization.

Explanation:-To link a project to a billing account, you need the necessary roles at the project level as well as at the billing account level. In this scenario, we are assigning the finance team the Billing Account User role on the billing account, which allows them to create new projects linked to the billing account on which the role is granted. We are also assigning them the Project Billing Manager role on the organization (trickles down to the project as well) which lets them attach the project to the billing account, but does not grant any rights over resources.

- Grant the development team Billing Account User (roles/billing.user) role on the billing account and Project Billing Manager (roles/billing.projectManager) on the GCP organization.
- Grant the development team Billing Account User (roles/billing.user) role on the billing account.

Q35) You are exploring the possibility of migrating a mission-critical application from your on-premises data centre to Google Cloud Platform. You want to host this on a GKE cluster with autoscaling enabled, and you need to ensure each node can run a pod to push the application logs to a third-party logging platform. How should you deploy the pod?

- Initialize the logging pod during the GKE Cluster creation.
- Deploy the logging pod in a StatefulSet Kubernetes object.
- Add the logging pod in the Deployment YAML file.
- ✅ Deploy the logging pod in a DaemonSet Kubernetes object.

Explanation:-In GKE, DaemonSets manage groups of replicated Pods and adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. So, this is a perfect fit for our monitoring pod.
https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset
DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluentd, and node monitoring daemons like collectd. For example, you could have DaemonSets for each type of daemon run on all of your nodes. Alternatively, you could run multiple DaemonSets for a single type of daemon, but have them use different configurations for different hardware types and resource needs.

Q36) Your team is responsible for the migration of all legacy on-premises applications to Google Cloud. Your team is a big admirer of serverless and has chosen App Engine Standard as the preferred choice for compute workloads. Your manager asked you to migrate a legacy accounting application built in C++, but you realized App Engine Standard doesn't support C++. What GCP compute services should you use instead to maintain the serverless aspect? (Choose two answers)

- ✅ Deploy the containerized version of the application in Cloud Run on GKE.

Explanation:-Cloud Run implements the Knative serving API, an open-source project to run serverless workloads on top of Kubernetes. That means you can deploy Cloud Run services anywhere Kubernetes runs. And suppose you need more control over your services (like access to GPU or more memory). In that case, you can also deploy these serverless containers in your GKE cluster instead of using the fully managed environment.

When using the fully managed environment, Cloud Run on GKE is serverless. Ref: https://github.com/knative/serving/blob/master/docs/spec/spec.md Ref: https://cloud.google.com/blog/products/serverless/cloud-run-bringing-serverless-to-containers

- Deploy the containerized version of the application in App Engine Flex.
- Deploy the containerized version of the application in Google Kubernetes Engine (GKE).
- ✅ Deploy the containerized version of the application in Cloud Run.

Explanation:-Cloud Run is a fully managed compute platform that automatically scales your stateless containers. Cloud Run is serverless: it abstracts away all infrastructure management, so you can focus on what matters most—building great applications. Run your containers in fully managed Cloud Run or on Anthos, which supports both Google Cloud and on-premises environments. Cloud Run is built upon an open standard, Knative, enabling the portability of your applications. Ref: https://cloud.google.com/run

---

Q37) 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?

- Create a key with the ssh-keygen command. Upload the key to the instance. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.
- Create a key with the ssh-keygen command. Then use the gcloud compute ssh command.
- Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.
- ✅ Use the gcloud compute ssh command.

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/sdk/gcloud/reference/compute/ssh

---

Q38) Your company runs several internal applications on bare metal Kubernetes servers in your on-premises data centre. One of the applications deployed in the Kubernetes cluster uses a NAS share to save files. In preparation for the upcoming migration to Google Cloud, you want to update the application to use Google Cloud Storage instead; however, security policies prevent virtual machines from having public IP addresses. What should you do?

- Make an exception and assign public IP addresses to the servers. Configure firewall rules to allow traffic from the VM public IP addresses to the IP range of Cloud Storage.
- Create a new VPC in GCP and deploy a proxy server like HAProxy/Squid to forward requests to Cloud Storage. Configure a VPN tunnel between the on-premises data centre and the GCP VPC. Have the servers access Cloud Storage through the proxy.
- ✅ Configure a VPN tunnel between the on-premises data centre and the GCP VPC. Create a custom route in the VPC for Google Restricted APIs IP range (199.36.153.4/30) and propagate the route over VPN. Resolve *.googleapis.com as a CNAME record to restricted.googleapis.com in your on-premises DNS server.

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

Configure a VPN tunnel between the on-premises data centre and the GCP VPC. Create a custom route in the VPC for Google Restricted APIs IP range (199.36.153.4/30) and propagate the route over VPN. Resolve *.googleapis.com as a CNAME record to restricted.googleapis.com in your on-premises DNS server. is the right answer, 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 you've 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."

- Migrate all VMs from the data centre to Google Compute Engine. Set up a Load Balancer on the GCP bucket and have the servers access Cloud Storage through the load balancer.

Q39) A GKE cluster (test environment) in your test GCP project is experiencing issues with a sidecar container connecting to Cloud SQL. This issue has resulted in a massive amount of log entries in Cloud Logging and shot up your bill by 25%. Your manager has asked you to disable these logs as quickly as possible and using the least number of steps. You want to follow Google recommended practices. What should you do?

- ✅ In Cloud Logging, disable the log source for GKE container resource in the Logs ingestion window.

Explanation:-We want to disable logs from a specific GKE container, and this is the only option that does that.
More information about logs exclusions: https://cloud.google.com/logging/docs/exclusions.

- Recreate the GKE cluster and disable Cloud Monitoring.
- Recreate the GKE cluster and disable Cloud Logging.
- In Cloud Logging, disable the log source for GKE Cluster Operations resource in the Logs ingestion window.

Q40) Your company has a number of GCP projects that are managed by the respective project teams. Your expenditure of all GCP projects combined has exceeded your operational expenditure budget. At a review meeting, it has been agreed that your finance team should be able to set budgets and view the current charges for all projects in the organization but not view the project resources; and your developers should be able to see the Google Cloud Platform billing charges for only their own

projects as well as view resources within the project. You want to follow Google recommended practices to set up IAM roles and permissions. What should you do?

- Add the developers and finance managers to the Viewer role for the Project.
- ✅ Add the finance team to the Billing Account Administrator role for each of the billing accounts that they need to manage. Add the developers to the Viewer role for the Project.

Explanation:-Billing Account Administrator role is an owner role for a billing account. It provides permissions to manage payment instruments, configure billing exports, view cost information, set budgets, link and unlink projects and manage other user roles on the billing account.
Ref: https://cloud.google.com/billing/docs/how-to/billing-access
Project viewer role provides permissions for read-only actions that do not affect the state, such as viewing (but not modifying) existing resources or data; including viewing the billing charges for the project.
https://cloud.google.com/iam/docs/understanding-roles#primitive_roles

- Add the finance team to the Viewer role for the Project. Add the developers to the Security Reviewer role for each of the billing accounts.
- Add the finance team to the default IAM Owner role. Add the developers to a custom role that allows them to see their own spend only.

Q41) You are enhancing a production application currently running on an Ubuntu Linux VM on Google Compute Engine. The new enhancements require a connection to SQL Server instance to persist user appointments. Your colleague has provisioned an SQL Server instance in a Google Compute Engine VM in US-Central region and has asked for your assistance to RDP to the VM in the least number of steps. What should you suggest?

- Add a firewall rule to allow TCP traffic on port 3389. In the GCP console, add a username and password for the Windows VM instance. Install Chrome RDP for Google Cloud Platform extension and click the RDP button in the console to connect to the instance with the credentials.
- Add a firewall rule to allow TCP traffic on port 22. In the GCP console, add a password for the Windows VM instance. Install Chrome RDP for Google Cloud Platform extension and click the RDP button in the console to connect to the instance with the credentials.
- Add a firewall rule to allow TCP traffic on port 3389. Install an RDP client and connect to the instance.
- ✅ In the GCP console, add a username and password for the Windows VM instance. Install an RDP client and connect to the instance with username and password.

Explanation:-This option correctly sets the username/password, which is essential. Also, the default VPC comes with port 3389 open to the public. The question doesn't explicitly state the compute engine is in a custom VPC, so it is safe to assume we are using default VPC which has default RDP access open to the public. Finally, you install an RDP client on the desktop and use the credentials set up earlier to RDP to the server.

Q42) Your compliance team requested all audit logs are stored for 10 years and to allow access for external auditors to view. You want to follow Google recommended practices. What should you do? (Choose two)

- Create an account for auditors to have view access to Stackdriver Logging.
- Export audit logs to Splunk via a Pub/Sub export sink.
- ✅ Export audit logs to Cloud Storage via an export sink.

Explanation:-Among all the storage solutions offered by Google Cloud Platform, Cloud storage offers the best pricing for long term storage of logs. 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 the storage options covered by the above options above. Ref: https://cloud.google.com/storage/pricing

- ✅ Generate a signed URL to the Stackdriver export destination for auditors to access.

Explanation:-In Google Cloud Storage, you can generate a signed URL to provide limited permission and time to make a request. Anyone who possesses it can use the signed URL to perform specified actions, such as reading an object, within a specified period of time. In our scenario, we do not need to create accounts for our auditors to provide access to logs in Cloud Storage. Instead, we can generate them signed URLs which are time-bound and lets them access/download log files. Ref: https://cloud.google.com/storage/docs/access-control/signed-urls

---

Q43) Your company plans to migrate all applications from its on-premises data centre to Google Cloud Platform. The DevOps team currently use Jenkins extensively to automate configuration updates in applications. How should you provision Jenkins in Google Cloud with the least number of steps?

- Download Jenkins binary from https://www.jenkins.io/download/ and deploy in Google App Engine Standard Service.
- ✅ Provision Jenkins from GCP marketplace.

Explanation:-The simplest way to launch a Jenkins server is from GCP Market place. GCP market place has several builds available for Jenkins: https://console.cloud.google.com/marketplace/browse?q=jenkins. All you need to do is spin up an instance from a suitable market place build, and you have a Jenkins server in a few minutes with just a few clicks.

- Create a Kubernetes Deployment YAML file referencing the Jenkins docker image and deploy to a new GKE cluster.
- Download Jenkins binary from https://www.jenkins.io/download/ and deploy in a new Google Compute Engine instance.

Q44) You developed an application to serve production users and you plan to use Cloud SQL to host user state data which is very critical for the application flow. You want to protect your user state data from zone failures. What should you do?

- ✅ Configure High Availability (HA) for Cloud SQL and Create a Failover replica in the same region but in a different zone.

Explanation:-If a HA-configured instance becomes unresponsive, Cloud SQL automatically switches to serving data from the standby instance. The HA configuration provides data redundancy. A Cloud SQL instance configured for HA has instances in the primary zone (Master node) and secondary zone (standby/failover node) within the configured region. Through synchronous replication to each zone's persistent disk, all writes made to the primary instance are also made to the standby instance. If the primary goes down, the standby/failover node takes over and your data continues to be available to client applications.
Ref: https://cloud.google.com/sql/docs/mysql/high-availability

- Create a Read replica in a different region.
- Create a Read replica in the same region but in a different zone.
- Configure High Availability (HA) for Cloud SQL and Create a Failover replica in a different region

---

Q45) Your company runs a very successful web platform and has accumulated 3 petabytes of customer activity data in sharded MySQL database located in your datacenter. Due to storage limitations in your on-premise data center, your company has decided to move this data to GCP. The data must be available all through the day. Your business analysts, who have experience of using a SQL Interface, have asked for a seamless transition. How should you store the data so that availability is ensured while optimizing the ease of analysis for the business analysts?

- Import data into Google Cloud Datastore.
- Import data into Google Cloud SQL.
- ✅ Import data into Google BigQuery.

Explanation:-Bigquery is a petabyte-scale serverless, highly scalable, and cost-effective cloud data warehouse that offers blazing-fast speeds, and with zero operational overhead. BigQuery supports a standard SQL dialect that is ANSI:2011 compliant, which reduces the impact and enables a seamless transition for your business analysts.
Ref: https://cloud.google.com/bigquery

- Import flat files into Google Cloud Storage.

Q46) Your company is migrating an application from its on-premises data centre to Google Cloud. One of the applications uses a custom Linux distribution that is not available on Google Cloud. Your solution architect has suggested using VMWare tools to exporting the image and store it in a Cloud Storage bucket. The VM Image is a single compressed 64 GB tar file. You started copying this file using gsutil over a dedicated 1Gbps network, but the transfer is taking a very long time to complete. Your solution architect has suggested using all of the 1Gbps Network to transfer the file quickly. What should you do?

- Upload the file Multi-Regional instead and move the file to Nearline Storage Class.
- ✅ Use parallel composite uploads to speed up the transfer.

Explanation:-With cloud storage, Object composition can be used for uploading an object in parallel: you can divide your data into multiple chunks, upload each chunk to a distinct object in parallel, compose your final object, and delete any temporary source objects. This option helps maximize your bandwidth usage and ensures the file is uploaded as fast as possible.
Ref: https://cloud.google.com/storage/docs/composite-objects#uploads

- Increase the transfer speed by decreasing the TCP window size.
- Restart the transfer from GCP console.

Q47) You are developing a simple application in App Engine Standard service. Unit testing and user acceptance testing has succeeded, and you want to build a new App Engine application to serve as your performance testing environment. What should you do?

- Configure a Deployment Manager YAML template to copy the application from the development GCP project into the performance testing GCP project.
- ✅ Create a new GCP project for the performance testing environment using gcloud and deploy your App Engine application to the new GCP project.

Explanation:-You can deploy to a different project by using --project flag.
By default, the service is deployed the current project configured via:

$ gcloud config set core/project PROJECT
To override this value for a single deployment, use the --project flag:

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

- Create a new GCP project for the performance testing environment using gcloud and copy the application from the development GCP project into the performance testing GCP project.
- Use gcloud to deploy the application to a new performance testing GCP project by specifying the --project parameter. Select Yes when prompted for confirmation on creating a new project.

Q48) Your networks team has set up Google compute network as shown below. In addition, firewall rules in the VPC network have been configured to allow egress to 0.0.0.0/0
Which instances have access to Google APIs and Services such as Google Cloud Storage?

- VM A1, VM A2
- ✅ VM A1, VM A2, VM B2

Explanation:-VM A1 can access Google APIs and services, including Cloud Storage because its network interface is located in subnet-a, which has Private Google Access enabled. Private Google Access applies to the instance because it only has an internal IP address.
VM B1 cannot access Google APIs and services because it only has an internal IP address and Private Google Access is disabled for subnet-b.

VM A2 and VM B2 can both access Google APIs and services, including Cloud Storage, because they each have external IP addresses. Private Google Access has no effect on whether or not these instances can access Google APIs and services because both have external IP addresses.

So the correct answer is VM A1, VM A2, VM B2

Ref: https://cloud.google.com/vpc/docs/private-access-options#example

- VM A1, VM A2, VM B1, VM B2
- VM A1, VM A2, VM B1

Q49) You created a compute instance by running gcloud compute instances create instance1. You intended to create the instance in project gcp-ace-proj-266520 but the instance got created in a different project. Your cloud shell gcloud configuration is as shown.
What should you do to delete the instance that was created in the wrong project and recreate it in gcp-ace-proj-266520 project?

- ✅ gcloud compute instances delete instance1
gcloud config set project gcp-ace-proj-266520
gcloud compute instances create instance1

Explanation:-This sequence of commands correctly deletes the instance from gcp-ace-lab-266520 which is the default project in the active gcloud configuration, then modifies the current configuration to set the default project to gcp-ace-proj-266520, and finally creates the instance in the project gcp-ace-proj-266520 which is the default project in active gcloud configuration at the time of running the command. This produces the intended outcome of deleting the instance from gcp-ace-lab-266520 project and recreating it in gcp-ace-prod-266520
Ref: https://cloud.google.com/sdk/gcloud/reference/config/set

Ref: https://cloud.google.com/sdk/gcloud/reference/compute/instances/create

Ref: https://cloud.google.com/sdk/gcloud/reference/compute/instances/delete
- gcloud compute instances delete instance1

gcloud compute instances create instance1
- gcloud compute instances delete instance1

gcloud config set compute/project gcp-ace-proj-266520
gcloud compute instances create instance1
- gcloud config set project gcp-ace-proj-266520

gcloud compute instances recreate instance1 --previous-project gcp-ace-lab-266520

Q50) Your team manages the game backend for a popular game with users all over the world. The game backend APIs runs on a fleet of VMs behind a Managed Instance Group (MIG) with autoscaling enabled. You have configured the scaling policy on the MIG to add more instances if the CPU utilization is consistently over 85%, and to scale down when the CPU utilization is consistently lower than 65%. You noticed the autoscaler adds more VMs than is necessary during the scale-up, and you suspect this might be down to an incorrect configuration in the health check – the initial delay on the health check is 30 seconds. Each VM takes just under 3 minutes before it is ready to process the requests from the web application and mobile app. What should you do to fix the scaling issue?

- ✅ Update the autoscaling health check to increase the initial delay to 200 seconds.

Explanation:-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 issue 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 the 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
- Update the Managed Instances template to set the maximum instances to 1.
- Update the Managed Instances template to set the maximum instances to 5.
- Update the autoscaling health check from HTTP to TCP.
