ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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

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

Answer: C

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

Question: 2

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

A. .00.0.0/0
B. 10.0.0.0/8
C. 172.16.0.0/12
D. 192.168.0.0/16

Answer: B

Explanation:

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

Question: 3

You want to select and configure a cost-effective solution for relational data on Google Cloud Platform. You are working with a small set of operational data in one geographic location. You need to support point-in-time recovery. What should you do?

A. Select Cloud SQL (MySQL). Verify that the enable binary logging option is selected.
B. Select Cloud SQL (MySQL). Select the create failover replicas option.
C. Select Cloud Spanner. Set up your instance with 2 nodes.
D. Select Cloud Spanner. Set up your instance as multi-regional.

Answer: A

Explanation:
Reference: https://cloud.google.com/sql/docs/mysql/backup-recovery/restore

https://cloud.google.com/sql/docs/mysql/backup-recovery/pitr#disk-usage

Question: 4

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 recreation 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. Set the Autohealing health check to healthy (HTTP).
D. Create a managed instance group. 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 autohealer determines that an application isnt

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.

Question: 5

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

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

Answer: D

Explanation:
Reference: https://medium.com/google-cloud/kubernetes-engine-kubectl-config-b6270d2b656c

kubectl config view -o jsonpath='{.users[].name}' # display the first user
kubectl config view -o jsonpath='{.users[*].name}' # get a list of users
kubectl config get-contexts # display list of contexts
kubectl config current-context # display the current-context
kubectl config use-context my-cluster-name # set the default context to my-cluster-name

https://kubernetes.io/docs/reference/kubectl/cheatsheet/

Question: 6

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

Explanation:
Reference: https://cloud.google.com/storage/docs/storage-classes#nearline

https://cloud.google.com/blog/products/gcp/introducing-coldline-and-a-unified-platform-for-data- storage

Cloud Storage Coldline: a low-latency storage class for long-term archiving Coldline is a new Cloud Storage class designed for long-term archival and disaster recovery. Coldline is perfect for the archival needs of big data or multimedia content, allowing businesses to archive years of data. Coldline provides fast and instant (millisecond) access to data and changes the way that companies think about storing and accessing their cold data.

Question: 7

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

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

Answer: D

Explanation:
(https://cloud.google.com/resource-manager/docs/project-migration#change_billing_account)

https://cloud.google.com/billing/docs/concepts
https://cloud.google.com/resource-manager/docs/project-migration

Question: 8

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.

Question: 9

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

Explanation:
Reference: https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed
https://cloud.google.com/appengine/docs/standard/go/config/appref

"App Engine calculates the number of instances necessary to serve your current application traffic based on scaling settings such as target_cpu_utilization and target_throughput_utilization. Setting min_idle_instances specifies the number of instances to run in addition to this calculated number. For example, if App Engine calculates that 5 instances are necessary to serve traffic, and min_idle_instances is set to 2, App Engine will run 7 instances (5, calculated based on traffic, plus 2 additional per min_idle_instances)."

Automatic scaling creates dynamic instances based on request rate, response latencies, and other application metrics. However, if you specify the number of minimum idle instances, that specified number of instances run as resident instances while any additional instances are dynamic.
Ref: https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed

Question: 10

You have a development project with appropriate IAM roles defined. You are creating a production project and want to have the same IAM roles on the new project, using the fewest possible steps. What should you do?

A. Use gcloud iam roles copy and specify the production project as the destination project.
B. Use gcloud iam roles copy and specify your organization as the destination organization.
C. In the Google Cloud Platform Console, use the 'create role from role' functionality.
D. In the Google Cloud Platform Console, use the 'create role' functionality and select all applicable

permissions.

Answer: A

Explanation:
Reference: https://cloud.google.com/sdk/gcloud/reference/iam/roles/copy

To create a copy of an existing role spanner.databaseAdmin into a project with PROJECT_ID, run: gcloud iam roles copy --source="roles/spanner.databaseAdmin" -- destination=CustomSpannerDbAdmin --dest-project=PROJECT_ID

Question: 11

You need a dynamic way of provisioning VMs on Compute Engine. The exact specifications will be in a dedicated configuration file. You want to follow Google's recommended practices. Which method should you use?

A. Deployment Manager
B. Cloud Composer
C. Managed Instance Group
D. Unmanaged Instance Group

Answer: A

Explanation:
https://cloud.google.com/deployment-manager/docs/configuration/create-basic-configuration

Question: 12

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

A. Use kubectl app deploy <dockerfilename>.
B. Use gcloud app deploy <dockerfilename>.
C. Create a docker image from the Dockerfile and upload it to Container Registry. Create a Deployment YAML file to point to that image. Use kubectl to create the deployment with that file.
D. Create a docker image from the Dockerfile and upload it to Cloud Storage. Create a Deployment YAML file to point to that image. Use kubectl to create the deployment with that file.

Answer: C

Explanation:

Reference: https://cloud.google.com/kubernetes-engine/docs/tutorials/hello-app

Question: 13

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

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

Answer: D

Explanation:

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

Question: 14

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

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

Answer: B

Explanation:
Reference: https://cloud.google.com/sdk/gcloud/reference/deployment- manager/deployments/update

Question: 15

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

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

Answer: B

Explanation:
Reference: https://cloud.google.com/bigquery/docs/estimate-costs

On-demand pricing Under on-demand pricing, BigQuery charges for queries by using one metric: the number of bytes processed (also referred to as bytes read). You are charged for the number of bytes processed whether the data is stored in BigQuery or in an external data source such as Cloud Storage, Drive, or Cloud Bigtable. On-demand pricing is based solely on usage.
https://cloud.google.com/bigquery/pricing#on_demand_pricing

Question: 16

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

Question: 17

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

Question: 18

You need to set up a policy so that videos stored in a specific Cloud Storage Regional bucket are moved to Coldline after 90 days, and then deleted after one year from their creation. How should you set up the policy?

A. Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete actions. Set the SetStorageClass action to 90 days and the Delete action to 275 days (365 – 90) B. Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete actions. Set the SetStorageClass action to 90 days and the Delete action to 365 days.
C. Use gsutil rewrite and set the Delete action to 275 days (365-90).
D. Use gsutil rewrite and set the Delete action to 365 days.

Answer: A

Explanation:

https://cloud.google.com/storage/docs/lifecycle#setstorageclass-cost
The object's time spent set at the original storage class counts towards any minimum storage duration that applies for the new storage class.

Question: 19

You have a Linux VM that must connect to Cloud SQL. You created a service account with the appropriate access rights. You want to make sure that the VM uses this service account instead of the default Compute Engine service account. What should you do?

A. When creating the VM via the web console, specify the service account under the 'Identity and API Access' section.
B. Download a JSON Private Key for the service account. On the Project Metadata, add that JSON as the value for the key compute-engine-service-account.
C. Download a JSON Private Key for the service account. On the Custom Metadata of the VM, add that JSON as the value for the key compute-engine-service-account.
D. Download a JSON Private Key for the service account. After creating the VM, ssh into the VM and save the JSON under ~/.gcloud/compute-engine-service-account.json.

Answer: A

Explanation:

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

https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances Changing the service account and access scopes for an instance If you want to run the VM as a different identity, or you determine that the instance needs a different set of scopes to call the required APIs, you can change the service account and the access scopes of an existing instance. For example, you can change access scopes to grant access to a new API, or change an instance so that it runs as a service account that you created, instead of the Compute Engine default service account. However, Google recommends that you use the fine-grained IAM policies instead of relying on access scopes to control resource access for the service account. To change an instance's service account and access scopes, the instance must be temporarily stopped. To stop your instance, read the documentation for Stopping an instance. After changing the service account or access scopes, remember to restart the instance. Use one of the following methods to the change service account or access scopes of the stopped instance.

Question: 20

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

A. Install a RDP client on your desktop. Verify that a firewall rule for port 3389 exists.
B. Install a RDP client in your desktop. Set a Windows username and password in the GCP Console. Use the credentials to log in to the instance.
C. Set a Windows password in the GCP Console. Verify that a firewall rule for port 22 exists. Click the RDP button in the GCP Console and supply the credentials to log in.
D. Set a Windows username and password in the GCP Console. Verify that a firewall rule for port 3389 exists. Click the RDP button in the GCP Console, and supply the credentials to log in.

Answer: D

Explanation:

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

https://cloud.google.com/compute/docs/instances/connecting-to-windows#before-you-begin

Question: 21

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

A. Create two configurations using gcloud config configurations create [NAME]. Run gcloud config configurations activate [NAME] to switch between accounts when running the commands to start the Compute Engine instances.

B. Create two configurations using gcloud config configurations create [NAME]. Run gcloud configurations list to start the Compute Engine instances.

C. Activate two configurations using gcloud configurations activate [NAME]. Run gcloud config list to start the Compute Engine instances.

D. Activate two configurations using gcloud configurations activate [NAME]. Run gcloud configurations list to start the Compute Engine instances.

Answer: A

Explanation:
"Run gcloud configurations list to start the Compute Engine instances". How the heck are you expecting to "start" GCE instances doing "configuration list".

Each gcloud configuration has a 1 to 1 relationship with the region (if a region is defined). Since we have two different regions, we would need to create two separate configurations using gcloud config configurations create
Ref: https://cloud.google.com/sdk/gcloud/reference/config/configurations/create

Secondly, you can activate each configuration independently by running gcloud config configurations activate [NAME]
Ref: https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate

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

Question: 22

You significantly changed a complex Deployment Manager template and want to confirm that the dependencies of all defined resources are properly met before committing it to the project. You want the most rapid feedback on your changes. What should you do?

A. Use granular logging statements within a Deployment Manager template authored in Python. B. Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.

C. Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.

D. Execute the Deployment Manager template using the –-preview option in the same project, and observe the state of interdependent resources.

Answer: D

Explanation:

Reference: https://cloud.google.com/deployment-manager/docs/deployments/updating- deployments

Question: 23

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

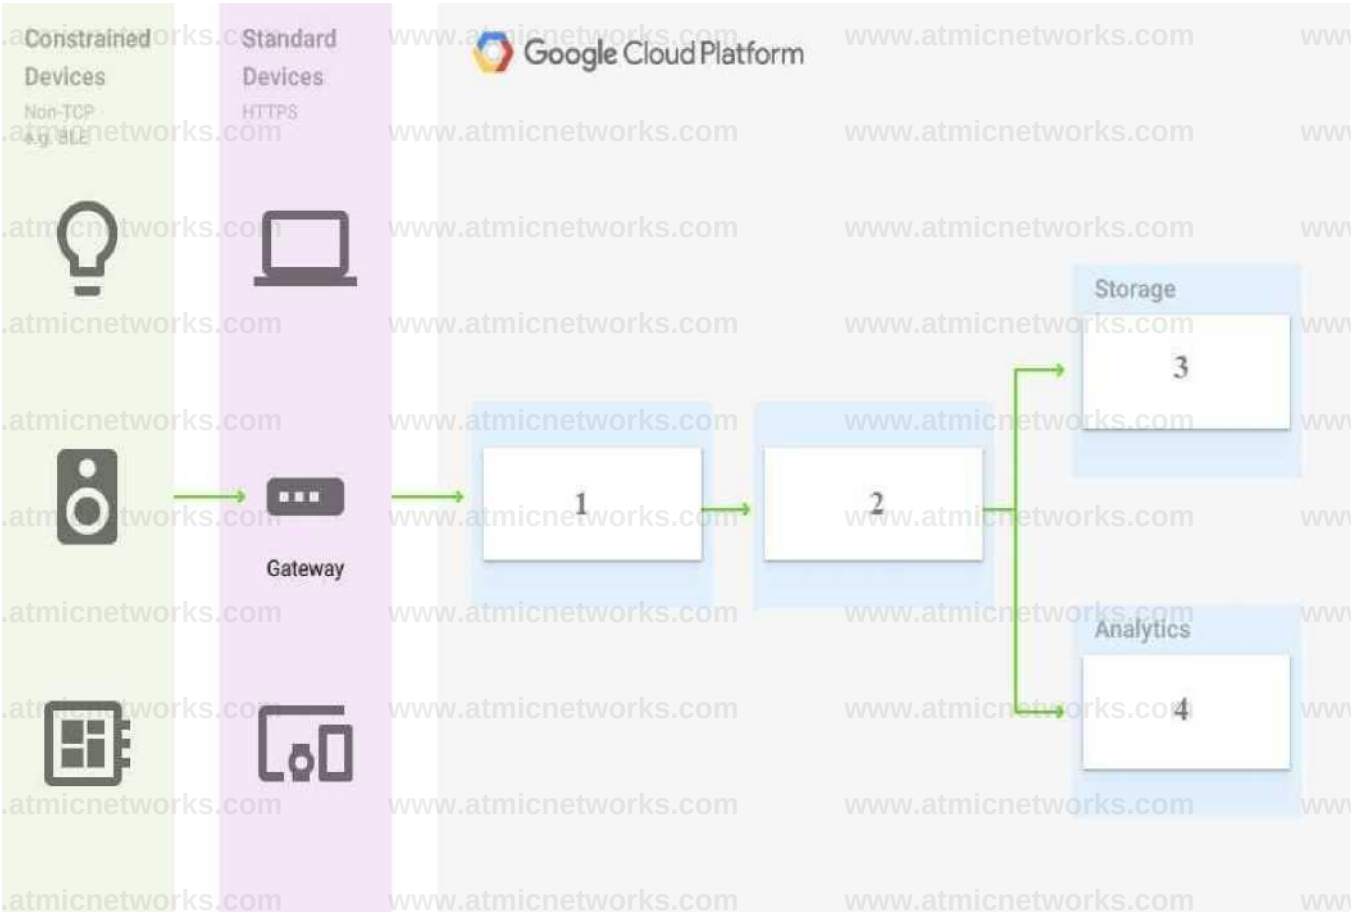

A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery

C. Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable
D. Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

Answer: D

Explanation:

Reference: https://cloud.google.com/solutions/correlating-time-series-dataflow

https://cloud.google.com/blog/products/data-analytics/handling-duplicate-data-in-streaming- pipeline-using-

pubsub-dataflow
https://cloud.google.com/bigtable/docs/schema-design-time-series

Question: 24

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

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

Answer: A

Explanation:

You can deploy to a different project by using –project flag.
By default, the service is deployed the current project configured via:
$ gcloud config set core/project PROJECT
To override this value for a single deployment, use the –project flag:
$ gcloud app deploy ~/my_app/app.yaml –project=PROJECT

Ref: https://cloud.google.com/sdk/gcloud/reference/app/deploy

Question: 25

You need to configure IAM access audit logging in BigQuery for external auditors. You want to follow Google-recommended practices. What should you do?

A. Add the auditors group to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
B. Add the auditors group to two new custom IAM roles.
C. Add the auditor user accounts to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.
D. Add the auditor user accounts to two new custom IAM roles.

Answer: A

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

Because if you directly add users to the IAM roles, then if any users left the organization then you have to remove the users from multiple places and need to revoke his/her access from multiple places. But, if you put a user into a group then its very easy to manage these type of situations. Now, if any user left then you just need to remove the user from the group and all the access got revoked

The organization creates a Google group for these external auditors and adds the current auditor to the group. This group is monitored and is typically granted access to the dashboard application. During normal access, the auditors' Google group is only granted access to view the historic logs stored in BigQuery. If any anomalies are discovered, the group is granted permission to view the actual Cloud Logging Admin Activity logs via the dashboard's elevated access mode. At the end of each audit period, the group's access is then revoked. Data is redacted using Cloud DLP before being made accessible for viewing via the dashboard application. The table below explains IAM logging roles that an Organization Administrator can grant to the service account used by the dashboard, as well as the resource level at which the role is granted.

Question: 26

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

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

Question: 27

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

A. Using the GCP Console, filter the Activity log to view the information.

B. Using the GCP Console, filter the Stackdriver log to view the information.
C. View the bucket in the Storage section of the GCP Console.
D. Create a trace in Stackdriver to view the information.

Answer: A

Explanation:

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

Question: 28

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

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

Answer: B

Explanation:

Storage Admin (roles/storage.admin) Grants full control of buckets and objects.
When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.

firebase.projects.get
resourcemanager.projects.get
resourcemanager.projects.list
storage.buckets.*
storage.objects.*
https://cloud.google.com/storage/docs/access-control/iam-roles

This role grants full control of buckets and objects. When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.
Ref: https://cloud.google.com/iam/docs/understanding-roles#storage-roles

Question: 29

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

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

Answer: A

Explanation:

Signed URLs are used to give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account. https://cloud.google.com/storage/docs/access- control/signed-urls

Question: 30

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

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

Answer: B

Explanation:

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

https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns

DaemonSets attempt to adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed.

In GKE, DaemonSets manage groups of replicated Pods and adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. So, this is a perfect fit for our monitoring pod. Ref: https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset

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

Question: 31

You want to send and consume Cloud Pub/Sub messages from your App Engine application. The Cloud Pub/Sub API is currently disabled. You will use a service account to authenticate your application to the API. You want to make sure your application can use Cloud Pub/Sub. What should you do?

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

Answer: A

Explanation:

Quickstart: using the Google Cloud Console
This page shows you how to perform basic tasks in Pub/Sub using the Google Cloud Console.

Note: If you are new to Pub/Sub, we recommend that you start with the interactive tutorial.

Before you begin
Set up a Cloud Console project.

Set up a project

Click to:

Create or select a project.
Enable the Pub/Sub API for that project.
You can view and manage these resources at any time in the Cloud Console.

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

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

Question: 32

You need to monitor resources that are distributed over different projects in Google Cloud Platform.
You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?

A. Use Shared VPC to connect all projects, and link Stackdriver to one of the projects.
B. For each project, create a Stackdriver account. In each project, create a service account for that project and grant it the role of Stackdriver Account Editor in all other projects.
C. Configure a single Stackdriver account, and link all projects to the same account.
D. Configure a single Stackdriver account for one of the projects. In Stackdriver, create a Group and add the other project names as criteria for that Group.

Answer: C

Explanation:

When you intially click on Monitoring(Stackdriver Monitoring) it creates a workspac(a stackdriver account) linked to the ACTIVE(CURRENT) Project from which it was clicked.

Now if you change the project and again click onto Monitoring it would create an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don't want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).

If you have accidently created two diff workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.

If we have only one workspace and two projects we can simply add other GCP Project under Monitoring > Settings > GCP Projects > Add GCP Projects.

https://cloud.google.com/monitoring/settings/multiple-projects
Nothing about groups https://cloud.google.com/monitoring/settings?hl=en

Question: 33

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

A. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.

B. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 1.

C. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

D. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

Answer: A

Explanation:

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

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

Since we need the application running at all times, we need a minimum 1 instance.

Only a single instance of the VM should run, we need a maximum 1 instance.

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

number of instances to 1.

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

Question: 34

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. Review the output section.

B. Run gcloud iam service-accounts list. Review the output section.

C. Navigate to the project and then to the IAM section in the GCP Console. Review the members and roles.

D. Navigate to the project and then to the Roles section in the GCP Console. 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.

Question: 35

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

A. Verify that you are Project Billing Manager for the GCP project. Update the existing project to link it to the existing billing account.

B. Verify that you are Project Billing Manager for the GCP project. Create a new billing account and link the new billing account to the existing project.

C. Verify that you are Billing Administrator for the billing account. Create a new project and link the new project to the existing billing account.

D. Verify that you are Billing Administrator for the billing account. Update the existing project to link it to the existing billing account.

Answer: B

Explanation:
Billing Administrators can not create a new billing account, and the project is presumably already created. Project Billing Manager allows you to link the created billing account to the project. It is

vague on how the billing account gets created but by process of elimination

Question: 36

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

Question: 37

You created a Google Cloud Platform project with an App Engine application inside the project. You initially configured the application to be served from the us-central region. Now you want the application to be served from the asia-northeast1 region. What should you do?

A. Change the default region property setting in the existing GCP project to asia-northeast1.
B. Change the region property setting in the existing App Engine application from us-central to asia-northeast1.
C. Create a second App Engine application in the existing GCP project and specify asia-northeast1 as the region to serve your application.
D. Create a new GCP project and create an App Engine application inside this new project. Specify asia-northeast1 as the region to serve your application.

Answer: D

Explanation:
https://cloud.google.com/appengine/docs/flexible/managing-projects-apps-billing#:~:text=Each%20Cloud%20project%20can%20contain%20only%20a%20single%20App%20Engine%20application%2C%20and%20once%20created%20you%20cannot%20change%20the%20location%20of%20your%20App%20Engine%20application.

Two App engine can't be running on the same project: you can check this easy diagram for more info:
https://cloud.google.com/appengine/docs/standard/an-overview-of-app-engine#components_of_an_application
And you can't change location after setting it for your app Engine.
https://cloud.google.com/appengine/docs/standard/locations

App Engine is regional and you cannot change an apps region after you set it. Therefore, the only way to have an app run in another region is by creating a new project and targeting the app engine to run in the required region (asia-northeast1 in our case).

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

Question: 38

You need to grant access for three users so that they can view and edit table data on a Cloud Spanner instance. What should you do?

A. Run gcloud iam roles describe roles/spanner.databaseUser. Add the users to the role.

B. Run gcloud iam roles describe roles/spanner.databaseUser. Add the users to a new group. Add the group to the role.

C. Run gcloud iam roles describe roles/spanner.viewer --project my-project. Add the users to the role.

D. Run gcloud iam roles describe roles/spanner.viewer --project my-project. Add the users to a new group. Add the group to the role.

Answer: B

Explanation:

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

Using the gcloud tool, execute the gcloud iam roles describe roles/spanner.databaseUser command on Cloud Shell. Attach the users to a newly created Google group and add the group to the role.

Question: 39

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

A. Enable the Node Auto-Repair feature for your GKE cluster.

B. Enable the Node Auto-Upgrades feature for your GKE cluster.

C. Select the latest available cluster version for your GKE cluster.

D. Select "Container-Optimized OS (cos)" as a node image for your GKE cluster.

Answer: B

Explanation:

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

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

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

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

Question: 40

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

Explanation:
Reference: https://cloud.google.com/load-balancing/docs/https/
According to this guide for setting up an HTTP (S) load balancer in GCP: The client SSL session terminates at the load balancer. Sessions between the load balancer and the instance can either be HTTPS (recommended) or HTTP.
https://cloud.google.com/load-balancing/docs/ssl

Question: 41

You have 32 GB of data in a single file that you need to upload to a Nearline Storage bucket. The WAN connection you are using is rated at 1 Gbps, and you are the only one on the connection. You want to use as much of the rated 1 Gbps as possible to transfer the file rapidly. How should you upload the file?

A. Use the GCP Console to transfer the file instead of gsutil.
B. Enable parallel composite uploads using gsutil on the file transfer.
C. Decrease the TCP window size on the machine initiating the transfer.
D. Change the storage class of the bucket from Nearline to Multi-Regional.

Answer: B

Explanation:

https://cloud.google.com/storage/docs/parallel-composite-uploads
https://cloud.google.com/storage/docs/uploads-downloads#parallel-composite-uploads

Question: 42

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

```
apiVersion: apps/vl
kind: Deployment metadata:
  name: myappl-deployment spec: selector:
    matchLabels: app: myappl replicas: 2 template:
  metadata: labels:
      app: myappl
    spec: containers: - name: main-container
      image: ger.io/my-company-repo/myappl:1.4
      env:
      - name: DB_PASSWORD value: "t0ugh2qu&ssl
      ports: - containerport: 3080
```

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

A. Store the database password inside the Docker image of the container, not in the YAML file.

B. Store the database password inside a Secret object. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.

C. Store the database password inside a ConfigMap object. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.

D. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

Answer: B

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

Question: 43

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 then 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 isnt up and isnt 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

Question: 44

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. Use a single-node cluster with a small instance type.
B. Select Google Kubernetes Engine. Use a three-node cluster with micro instance types.
C. Select Compute Engine. Use preemptible VM instances of the appropriate standard machine type.
D. Select Compute Engine. 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

Question: 45

You recently deployed a new version of an application to App Engine and then discovered a bug in the release. You need to immediately revert to the prior version of the application. What should you do?

A. Run gcloud app restore.

B. On the App Engine page of the GCP Console, select the application that needs to be reverted and click Revert.

C. On the App Engine Versions page of the GCP Console, route 100% of the traffic to the previous version.

D. Deploy the original version as a separate application. Then go to App Engine settings and split traffic between applications so that the original version serves 100% of the requests.

Answer: C

Explanation:
Reference: https://medium.com/google-cloud/app-engine-project-cleanup-9647296e796a

Question: 46

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

A. Check the app.yaml file for your application and check project settings.

B. Check the web-application.xml file for your application and check project settings.

C. Go to Deployment Manager and review settings for deployment of applications.

D. Go to Cloud Shell and run gcloud config list to review the Google Cloud configuration used for deployment.

Answer: D

Explanation:
C:\GCP\appeng>gcloud config list
[core]
account = xxx@gmail.com
disable_usage_reporting = False
project = my-first-demo-xxxx

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

Question: 47

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

A. Create an instance template for the instances. Set the 'Automatic Restart' to on. Set the 'On-host maintenance' to Migrate VM instance. Add the instance template to an instance group.
B. Create an instance template for the instances. Set 'Automatic Restart' to off. Set 'On-host maintenance' to Terminate VM instances. Add the instance template to an instance group. C. Create an instance group for the instances. Set the 'Autohealing' health check to healthy (HTTP). D. Create an instance group for the instance. Verify that the 'Advanced creation options' setting for 'do not retry machine creation' is set to off.

Answer: A

Explanation:

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

highly available Add the instance template to an instance group so instances can be managed.

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

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

Question: 48

You host a static website on Cloud Storage. Recently, you began to include links to PDF files on this site. Currently, when users click on the links to these PDF files, their browsers prompt them to save the file onto their local system. Instead, you want the clicked PDF files to be displayed within the browser window directly, without prompting the user to save the file locally. What should you do?

A. Enable Cloud CDN on the website frontend.
B. Enable 'Share publicly' on the PDF file objects.
C. Set Content-Type metadata to application/pdf on the PDF file objects.
D. Add a label to the storage bucket with a key of Content-Type and value of application/pdf.

Answer: C

Explanation:

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

Question: 49

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

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

Answer: D

Explanation:

In Google compute engine, if predefined machine types don't meet your needs, you can create an instance with custom virtualized hardware settings. Specifically, you can create an instance with a custom number of vCPUs and custom memory, effectively using a custom machine type. Custom machine types are ideal for the following scenarios:

1. Workloads that aren't a good fit for the predefined machine types that are available to you.
2. Workloads that require more processing power or more memory but don't need all of the upgrades that are provided by the next machine type level.

In our scenario, we only need a memory upgrade. Moving to a bigger instance would also bump up the CPU which we don't need so we have to use a custom machine type. It is not possible to change memory while the instance is running so you need to first stop the instance, change the memory and then start it again. See below a screenshot that shows how CPU/Memory can be customized for an instance that has been stopped.

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

Question: 50

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

requirements?

A. Create a single custom VPC with 2 subnets. Create each subnet in a different region and with a different CIDR range.

B. Create a single custom VPC with 2 subnets. Create each subnet in the same region and with the same CIDR range.

C. Create 2 custom VPCs, each with a single subnet. Create each subnet is a different region and with a different CIDR range.

D. Create 2 custom VPCs, each with a single subnet. Create each subnet in the same region and with the same CIDR range.

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

Question: 51

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

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

B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.

C. In the Instance Template, add the label 'health-check'.

D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Answer: A

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

Question: 52

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

A. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery jobUser role to the group.

B. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery dataViewer user role to the group.

C. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery jobUser role to the group.

D. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery dataViewer user role to the group.

Answer: C

Explanation:

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

BigQuery Data Viewer
(roles/bigquery.dataViewer)

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

Read data and metadata from the table or view.
This role cannot be applied to individual models or routines.

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

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

Lowest-level resources where you can grant this role:

Table
View

BigQuery Job User
(roles/bigquery.jobUser)

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

Lowest-level resources where you can grant this role:

Project

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

Question: 53

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

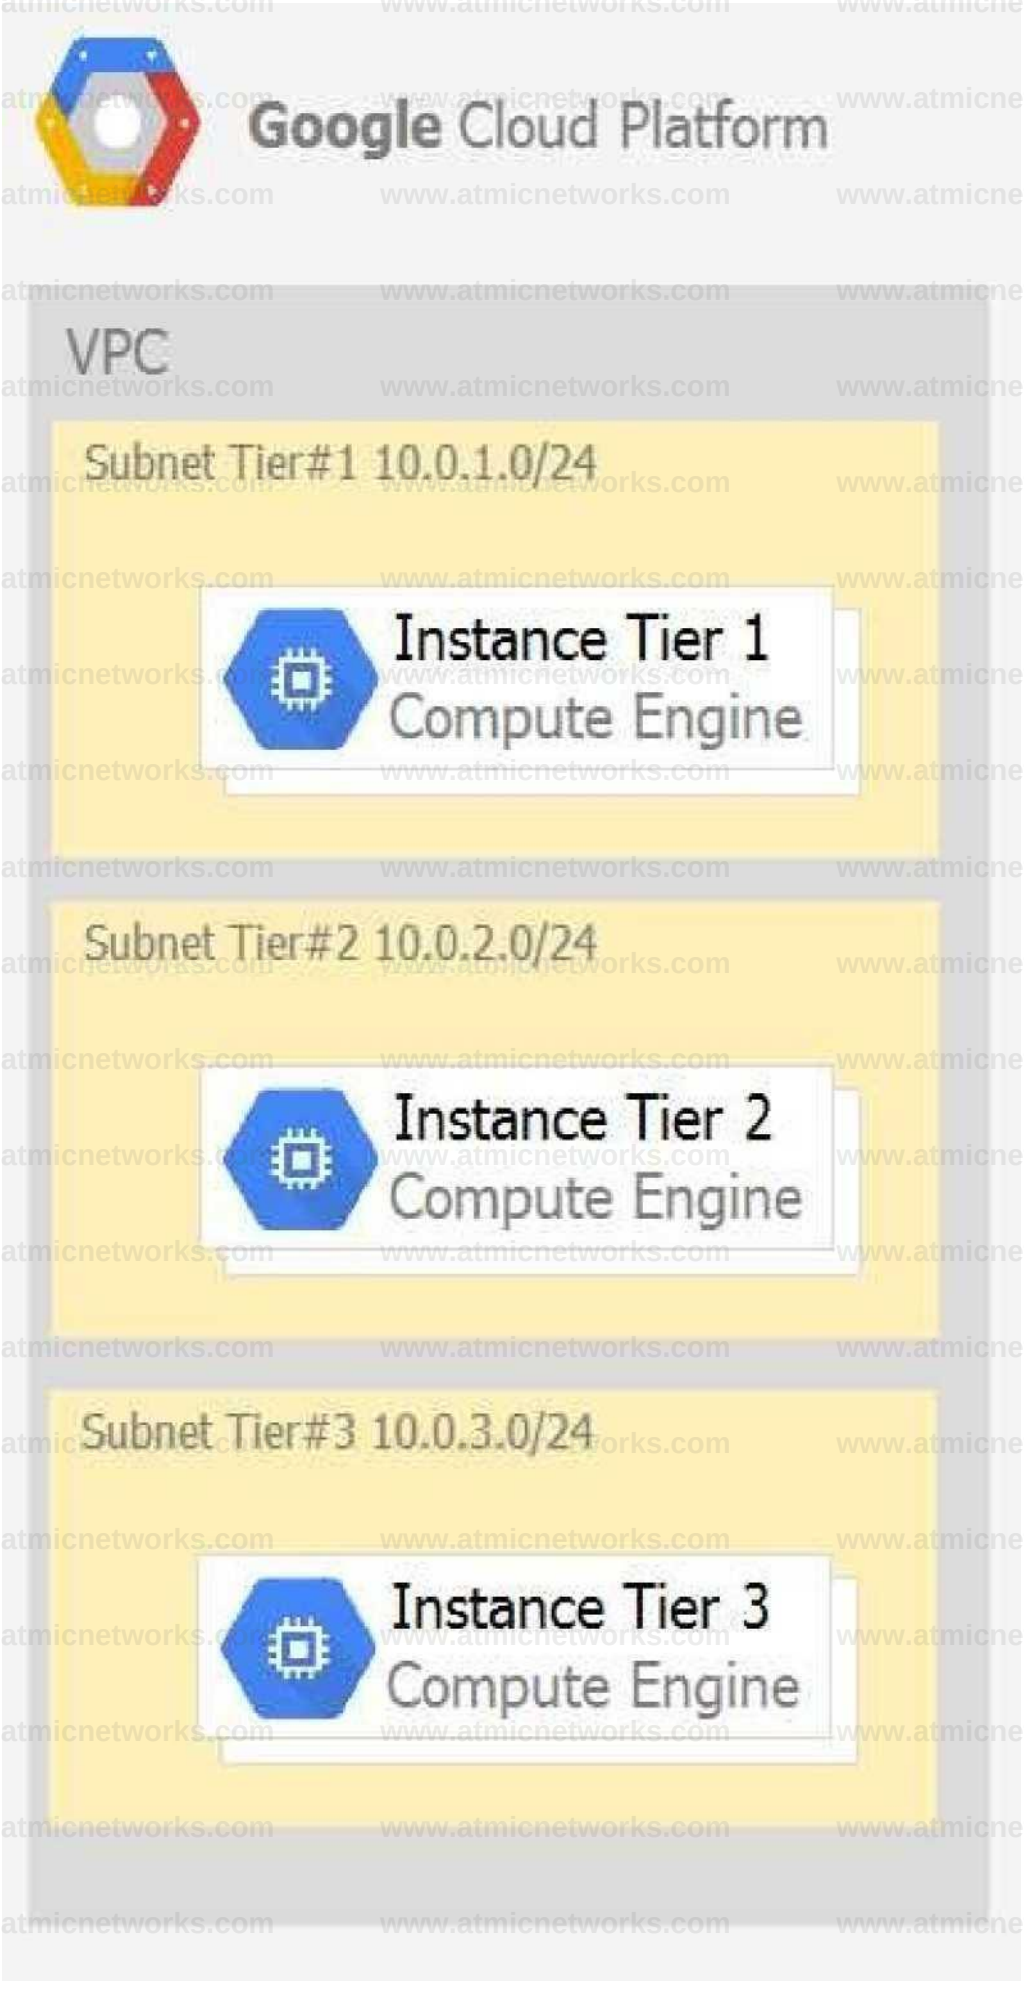

Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:

- 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 all2. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow all

D. 1. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow TCP: 80802. Create an egress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)• Protocols: allow TCP: 8080

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

Question: 54

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

A. 1. Create a subnetwork in the same VPC, in europe-west1.2. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

B. 1. Create a VPC and a subnetwork in europe-west1.2. Expose the application with an internal load balancer.3. Create the new instance in the new subnetwork and use the load balancer's address as the endpoint.

C. 1. Create a subnetwork in the same VPC, in europe-west1.2. Use Cloud VPN to connect the two subnetworks.3.

Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

D. 1. Create a VPC and a subnetwork in europe-west1.2. Peer the 2 VPCs.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

Answer: C

Explanation:

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

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

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

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

Question: 55

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

A. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE container resource.

B. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE Cluster Operations resource.

C. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Logging.

D. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Monitoring.

Answer: A

Explanation:

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

GKE Containers have more log than GKE Cluster Operations:

.-GKE Containe:

cluster_name: An immutable name for the cluster the container is running in.
namespace_id: Immutable ID of the cluster namespace the container is running in.
instance_id: Immutable ID of the GCE instance the container is running in.
pod_id: Immutable ID of the pod the container is running in.
container_name: Immutable name of the container.
zone: The GCE zone in which the instance is running.

VS

.-GKE Cluster Operations
project_id: The identifier of the GCP project associated with this resource, such as "my-project".
cluster_name: The name of the GKE Cluster.
location: The location in which the GKE Cluster is running.

Question: 56

You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?

A. Deploy the new version in the same application and use the --migrate option.
B. Deploy the new version in the same application and use the --splits option to give a weight of 99 to the current version and a weight of 1 to the new version.
C. Create a new App Engine application in the same project. Deploy the new version in that application. Use the App Engine library to proxy 1% of the requests to the new version. D. Create a new App Engine application in the same project. Deploy the new version in that application. Configure your network load balancer to send 1% of the traffic to that new application.

Answer: B

Explanation:

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

Question: 57

You have a web application deployed as a managed instance group. You have a new version of the application to gradually deploy. Your web application is currently receiving live web traffic. You want to ensure that the available capacity does not decrease during the deployment. What should you do?

A. Perform a rolling-action start-update with maxSurge set to 0 and maxUnavailable set to 1.
B. Perform a rolling-action start-update with maxSurge set to 1 and maxUnavailable set to 0.
C. Create a new managed instance group with an updated instance template. Add the group to the backend service for the load balancer. When all instances in the new managed instance group are healthy, delete the old managed instance group.
D. Create a new instance template with the new application version. Update the existing managed instance

group with the new instance template. Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

Answer: B

Explanation:

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

Question: 58

You are building an application that stores relational data from users. Users across the globe will use this application. Your CTO is concerned about the scaling requirements because the size of the user base is unknown. You need to implement a database solution that can scale with your user growth with minimum configuration changes. Which storage solution should you use?

A. Cloud SQL
B. Cloud Spanner
C. Cloud Firestore
D. Cloud Datastore

Answer: B

Explanation:

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

enterprise-grade security

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

	CLOUD SPANNER	TRADITIONAL RELATIONAL	TRADITIONAL NON-RELATIONAL
Schema	xZ Yes	xZ Yes	X No
SQL	xZ Yes	xZ Yes	X No
Consistency	xZ Strong	^Z Strong	X Eventual
Availability	xZ High	X Failover	xZ High

Scalability	xZ Horizontal	X Vertical	xZ Horizontal
Replication	xZ Automatic	Configurable	Configurable

Question: 59

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."

Question: 60

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

A. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Set the service's externalTrafficPolicy to Cluster.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.
B. 1. In GKE, create a Service of type NodePort that uses the application's Pods as backend.2. Create a Compute Engine instance called proxy with 2 network interfaces, one in each VPC.3. Use iptables on this instance to forward traffic from gce-network to the GKE nodes.4. Configure the Compute Engine instance to use the address of

proxy in gce-network as endpoint.

C. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add an annotation to this service: cloud.google.com/load-balancer-type: Internal3. Peer the two VPCs together.4. Configure the Compute Engine instance to use the address of the load balancer that has been created.

D. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG's instances.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.

Answer: A

Explanation:

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

https://medium.com/pablo-perez/k8s-externaltrafficpolicy-local-or-cluster-40b259a19404
https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing

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

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

Question: 61

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 costefficient as well. Coldline Storage is a very low-cost, highly durable storage service for storing infrequently accessed data.

Question: 62

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.

Question: 63

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.

Question: 64

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.

Question: 65

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

Question: 66

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

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

Answer: A

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

Question: 67

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

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

Answer: A

Explanation:

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

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

Question: 68

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

A. 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.2. Update your instances' metadata to add the following value: logs-destination: bq://platform-logs.
B. 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.2. Create a Cloud Function that is triggered by messages in the logs topic.3. Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.
C. 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.2. Click Create Export.3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.
D. 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.2. Configure this Cloud Function to create a BigQuery Job that executes this query:INSERT INTO dataset.platform-logs (timestamp,

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

Answer: C

Explanation:

I. In Stackdriver Logging, create a filter to view only Compute Engine logs. 2. Click Create Export. 3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

Question: 69

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

Question: 70

You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML. You need to

enable authentication to the APIs from your on-premises environment. What should you do?

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

permissions, and use this user account for authentication from your data center.

Answer: B

Explanation:
Reference: https://cloud.google.com/vision/automl/docs/before-you-begin

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

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

Question: 71

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

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

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

Question: 72

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

```
apiVersion: apps/vl kind: Deployment met adat a:
  name: myapp-deployment
spec:
  selector:
    matchLabels:
      app: myapp replicas: 2 template:
    metadata:
      labels:
        app: myapp
    spec:
      containers:
      - name: myapp image: myapp: 1.1 port s:
        - containerPort: 80
```

```
apiVersion: vl kind: Service metadata:
  name: myapp-service
spec:
  ports:
  - port: 8000
    target Port: 8 0 protocol: TCP
  selector:
    app: myapp
```

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

```
kubectl get pods -1 app=myapp
NAME                                READY     STATUS    RESTART   AGE
myapp-deployment-58ddbbb995-lp8 6m  0/1       Pending   0         9m
myapp-deployment-58ddbbb995-qjpkg   1/1       Running   0         9m
```

You want to find out why the pod is stuck in pending status. What should you do?

A. Review details of the myapp-service Service object and check for error messages.
B. Review details of the myapp-deployment Deployment object and check for error messages.
C. Review details of myapp-deployment-58ddbbb995-lp86m Pod and check for warning messages.
D. View logs of the container in myapp-deployment-58ddbbb995-lp86m pod and check for warning messages.

Answer: C

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

Question: 73

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

A. After the VM has been created, use your Google Account credentials to log in into the VM.

B. After the VM has been created, use gcloud compute reset-windows-password to retrieve the login credentials for the VM.

C. When creating the VM, add metadata to the instance using 'windows-password' as the key and a password as the value.

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

Answer: B

Explanation:

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

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

Question: 74

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

A. Set metadata to enable-oslogin=true for the instance. Grant the dev1 group the compute.osLogin role. Direct them to use the Cloud Shell to ssh to that instance.

B. Set metadata to enable-oslogin=true for the instance. Set the service account to no service account for that instance. Direct them to use the Cloud Shell to ssh to that instance.

C. Enable block project wide keys for the instance. Generate an SSH key for each user in the dev1 group. Distribute the keys to dev1 users and direct them to use their third-party tools to connect. D. Enable block project wide keys for the instance. Generate an SSH key and associate the key with that instance. Distribute the key to dev1 users and direct them to use their third-party tools to connect.

Answer: A

Explanation:
Reference: https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys
After you enable OS Login on one or more instances in your project, those VMs accept connections only from user accounts that have the necessary IAM roles in your project or organization. In this case, we are granting the group compute.osLogin which lets them log in as non-administrator account. And since we are directing them to use Cloud Shell to ssh, we dont need to add their SSH keys to the instance metadata.
Ref: https://cloud.google.com/compute/docs/instances/managing-instance-access#configure_users Ref: https://cloud.google.com/compute/docs/instances/managing-instance- access#add_oslogin_keys

Question: 75

You need to produce a list of the enabled Google Cloud Platform APIs for a GCP project using the gcloud command line in the Cloud Shell. The project name is my-project. What should you do?

A. Run gcloud projects list to get the project ID, and then run gcloud services list --project <project ID>.
B. Run gcloud init to set the current project to my-project, and then run gcloud services list -available.
C. Run gcloud info to view the account value, and then run gcloud services list --account <Account>. D. Run gcloud projects describe <project ID> to verify the project value, and then run gcloud services list --available.

Answer: A

Explanation:

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

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

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

gcloud services list

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

gcloud services list –available

https://cloud.google.com/sdk/gcloud/reference/services/list#--available
--available
Return the services available to the project to enable. This list will include any services that the project has already enabled.

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

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

Question: 76

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

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

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

Question: 77

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

A. Fill in local SSD. Fill in persistent disk storage and snapshot storage.
B. Fill in local SSD. Add estimated cost for cluster management.
C. Select Add GPUs. Fill in persistent disk storage and snapshot storage.
D. Select Add GPUs. Add estimated cost for cluster management.

Answer: A

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

Question: 78

You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?

A. Create a Kubernetes Service of type NodePort for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.

B. Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service.

C. Create a Kubernetes Service of type NodePort to expose the application on port 443 of each node of the Kubernetes cluster. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.

D. Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application. Forward the public traffic to HAProxy with an iptable rule. Configure the DNS name of your application using the public IP of the node HAProxy is running on.

Answer: A

Explanation:

Reference: https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer

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

Kubernetes Service of type ClusterIP exposes the Service on a cluster-internal IP. Choosing this value makes the Service only reachable from within the cluster so you can not route external traffic to this IP.

Ref: https://kubernetes.io/docs/concepts/services-networking/service/

Question: 79

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

A. Verify that both projects are in a GCP Organization. Create a new VPC and add all instances.

B. Verify that both projects are in a GCP Organization. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.

C. Verify that you are the Project Administrator of both projects. Create two new VPCs and add all instances.

D. Verify that you are the Project Administrator of both projects. Create a new VPC and add all instances.

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."

Question: 80

You want to add a new auditor to a Google Cloud Platform project. The auditor should be allowed to read, but not modify, all project items.

How should you configure the auditor's permissions?

A. Create a custom role with view-only project permissions. Add the user's account to the custom role.
B. Create a custom role with view-only service permissions. Add the user's account to the custom role.
C. Select the built-in IAM project Viewer role. Add the user's account to this role.
D. Select the built-in IAM service Viewer role. Add the user's account to this role.

Answer: C

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/access-control-proj

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

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

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

Question: 81

You are operating a Google Kubernetes Engine (GKE) cluster for your company where different teams can run non-production workloads. Your Machine Learning (ML) team needs access to Nvidia Tesla P100 GPUs to train their models. You want to minimize effort and cost. What should you do?

A. Ask your ML team to add the "accelerator: gpu" annotation to their pod specification.
B. Recreate all the nodes of the GKE cluster to enable GPUs on all of them.
C. Create your own Kubernetes cluster on top of Compute Engine with nodes that have GPUs. Dedicate this cluster to your ML team.
D. Add a new, GPU-enabled, node pool to the GKE cluster. Ask your ML team to add the cloud.google.com/gke -accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Answer: D

Explanation:

This is the most optimal solution. Rather than recreating all nodes, you create a new node pool with GPU enabled. You then modify the pod specification to target particular GPU types by adding node

selector to your workloads Pod specification. YOu still have a single cluster so you pay Kubernetes cluster management fee for just one cluster thus minimizing the cost.
Ref: https://cloud.google.com/kubernetes-engine/docs/how-to/gpus
Ref: https://cloud.google.com/kubernetes-engine/pricing

Example:

```
apiVersion: v1
kind: Pod
metadata:
name: my-gpu-pod
spec:
containers:
name: my-gpu-container
image: nvidia/cuda:10.0-runtime-ubuntu18.04
command: [/bin/bash]
resources:
limits:
nvidia.com/gpu: 2
nodeSelector:
cloud.google.com/gke-accelerator: nvidia-tesla-k80 # or nvidia-tesla-p100 or nvidia-tesla-p4 or nvidia-tesla-v100 or nvidia-tesla-t4
```

Question: 82

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

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

Answer: A

Explanation:

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

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

Question: 83

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

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

Answer: B

Explanation:
Reference: https://cloud.google.com/resource-manager/docs/creating-managing-organization Default behavior does not grant access to the "your GCP Project" Default behavior allow only create billing account and project - When the organization is created, all users in your domain are automatically granted Project Creator and Billing Account Creator IAM roles at the organization level. This enables users in your domain to continue creating projects with no disruption.

Question: 84

You have a Google Cloud Platform account with access to both production and development projects. You need to create an

automated process to list all compute instances in development and production projects on a daily basis. What should you do?

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

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

Question: 85

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

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

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

Question: 86

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

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

Answer: A

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

Question: 87

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

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

Answer: C

Explanation:

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

Question: 88

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

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

Answer: A

Explanation:

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

Question: 89

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

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

Answer: A

Explanation:

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

Question: 90

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

Explanation:
Reference: https://cloud.google.com/storage/docs/managing-lifecycles

Nearline Storage is ideal for data you plan to read or modify on average once per month or less. And this option archives just the noncurrent versions which is what we want to do.

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

Question: 91

Your company's infrastructure is on-premises, but all machines are running at maximum capacity.
You want to burst to Google Cloud. The workloads on Google Cloud must be able to directly communicate to the workloads on-premises using a private IP range. What should you do?

A. In Google Cloud, configure the VPC as a host for Shared VPC.
B. In Google Cloud, configure the VPC for VPC Network Peering.
C. Create bastion hosts both in your on-premises environment and on Google Cloud. Configure both as proxy servers using their public IP addresses.
D. Set up Cloud VPN between the infrastructure on-premises and Google Cloud.

Answer: D

Explanation:

"Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same

organization."
https://cloud.google.com/vpc/docs/vpc-peering

while
"Cloud Interconnect provides low latency, high availability connections that enable you to reliably transfer data between your on-premises and Google Cloud Virtual Private Cloud (VPC) networks." https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview

and
"HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your onpremises network to your VPC

network through an IPsec VPN connection in a single region." https://cloud.google.com/network-connectivity/docs/vpn/concepts/overview

Question: 92

You want to select and configure a solution for storing and archiving data on Google Cloud Platform.
You need to support compliance objectives for data from one geographic location. This data is archived after 30 days and needs to be accessed annually. What should you do?

A. Select Multi-Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.
B. Select Multi-Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
C. Select Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Nearline Storage.
D. Select Regional Storage. Add a bucket lifecycle rule that archives data after 30 days to Coldline Storage.

Answer: D

Explanation:

Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

The Real description is about Coldline storage Class:

Coldline Storage
Coldline Storage 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 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 costeffective, as it offers the lowest storage costs.

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

Question: 93

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

Question: 94

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:

```
S kubectl get pods -l app=myapp
NAME                                READY     STATUS    RESTART   AGE
myapp-deployment-58ddbbb995-lp86m   0/1       Pending   0           9m
myapp-deployment-58ddbbb995-qjpkg   1/1       Running   0           9m
```

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. 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. Heres 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.

Question: 95

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

Question: 96

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. Apply the appropriate number of slots to each Project.

Answer: B,E

Explanation:

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

Question: 97

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. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.

D. Use the cos_containerd image for your GKE nodes. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Answer: C

Explanation:

Reference: https://cloud.google.com/kubernetes-engine/sandbox/

GKE Sandbox provides an extra layer of security to prevent untrusted code from affecting the host kernel on your cluster nodes when containers in the Pod execute unknown or untrusted code. Multitenant clusters and clusters whose containers run untrusted workloads are more exposed to security vulnerabilities than other clusters. Examples include SaaS providers, web-hosting providers, or other organizations that allow their users to upload and run code. When you enable GKE Sandbox on a node pool, a sandbox is created for each Pod running on a node in that node pool. In addition, nodes running sandboxed Pods are prevented from accessing other Google Cloud services or cluster metadata. Each sandbox uses its own userspace kernel. With this in mind, you can make decisions

about how to group your containers into Pods, based on the level of isolation you require and the characteristics of your applications.

Ref: https://cloud.google.com/kubernetes-engine/docs/concepts/sandbox-pods

Question: 98

Your customer has implemented a solution that uses Cloud Spanner and notices some read latency- related performance issues on one table. This table is accessed only by their users using a primary key. The table schema is shown below.

```
CREATE TABLE Persons { person_id   // sequential number based on number of registration
INT64 NOT NULL,                                    // system date
account_creation_date DATE,        // customer birthdate
birthdate DATE, firstname STRING       // first name
(255), lastname STRING (255),          // last name
profile_picture BYTES (255) )          // profile picture
PRIMARY KEY (person_id)
```

You want to resolve the issue. What should you do?

A. Remove the RMlk^Kt^ field from the table.
B. Add a secondary mdex on the Mud column.
C. Change the primary key to not have monotonically increasing values.
D. Create a secondary index using the following Data Definition Language (DDL) `CREATE INDEX person_id_ix ON Persons ( person_id, firstname, lastname ) STORING ( profile_picture )`

A. Option A B. Option B C. Option C D. Option D

Answer: C

Explanation:

As mentioned in Schema and data model, you should be careful when choosing a primary key to not accidentally create hotspots in your database. One cause of hotspots is having a column whose value monotonically increases as the first key part, because this results in all inserts occurring at the end of your key space. This pattern is undesirable because Cloud Spanner divides data among servers by key ranges, which means all your inserts will be directed at a single server that will end up doing all the work. https://cloud.google.com/spanner/docs/schema-design#primary-key-prevent-hotspots

Question: 99

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 thegroup for the finance team to roles/billing user role.
B. Add thegroup for the finance team to roles/billingadmin role.
C. Add thegroup for the finance team to roles/billingviewer role.
D. Add thegroup 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

Question: 100

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

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

Answer: D

Explanation:

Question: 101

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

Question: 102

Your organization needs to grant users access to query datasets in BigQuery but prevent them from accidentally deleting the datasets. You want a solution that follows Google-recommended practices. What should you do?

A. Add users to roles/bigquery user role only, instead of roles/bigquery dataOwner.
B. Add users to roles/bigquery dataEditor role only, instead of roles/bigquery dataOwner.
C. Create a custom role by removing delete permissions, and add users to that role only.
D. Create a custom role by removing delete permissions. Add users to the group, and then add the group to the custom role.

Answer: D

Explanation:
https://cloud.google.com/bigquery/docs/access-control#custom_roles
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.

Question: 103

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

Question: 104

Your company set up a complex organizational structure on Google Could Platform. The structure includes hundreds of folders and projects. Only a few team members should be able to view the hierarchical structure. You need to assign minimum permissions to these team members and you want to follow Google-recommended practices. What should you do?

A. Add the users to roles/browser role.
B. Add the users to roles/iam.roleViewer role.
C. Add the users to a group, and add this group to roles/browser role.
D. Add the users to a group, and add this group to roles/iam.roleViewer role.

Answer: C

Explanation:

We need to apply the GCP Best practices. roles/browser Browser Read access to browse the

hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project. https://cloud.google.com/iam/docs/understanding-roles

Question: 105

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

Question: 106

Your organization has a dedicated person who creates and manages all service accounts for Google Cloud projects. You need to assign this person the minimum role for projects. What should you do?

A. Add the user to roles/iam.roleAdmin role.
B. Add the user to roles/iam.securityAdmin role.
C. Add the user to roles/iam.serviceAccountUser role.
D. Add the user to roles/iam.serviceAccountAdmin role.

Answer: D

Explanation:

Reference: https://cloud.google.com/iam/docs/creating-managing-service-accounts

Service Account User (roles/iam.serviceAccountUser): Includes permissions to list service accounts, get details about a service account, and impersonate a service account. Service Account Admin (roles/iam.serviceAccountAdmin): Includes permissions to list service accounts and get details about a service account. Also includes permissions to create, update, and delete service accounts, and to view or change the IAM policy on a service account.

Question: 107

You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your dat a. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

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

Answer: A

Explanation:

Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Since we have a requirement to access data once a quarter and want to go with the most cost-efficient option, we should select Coldline Storage.

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

S

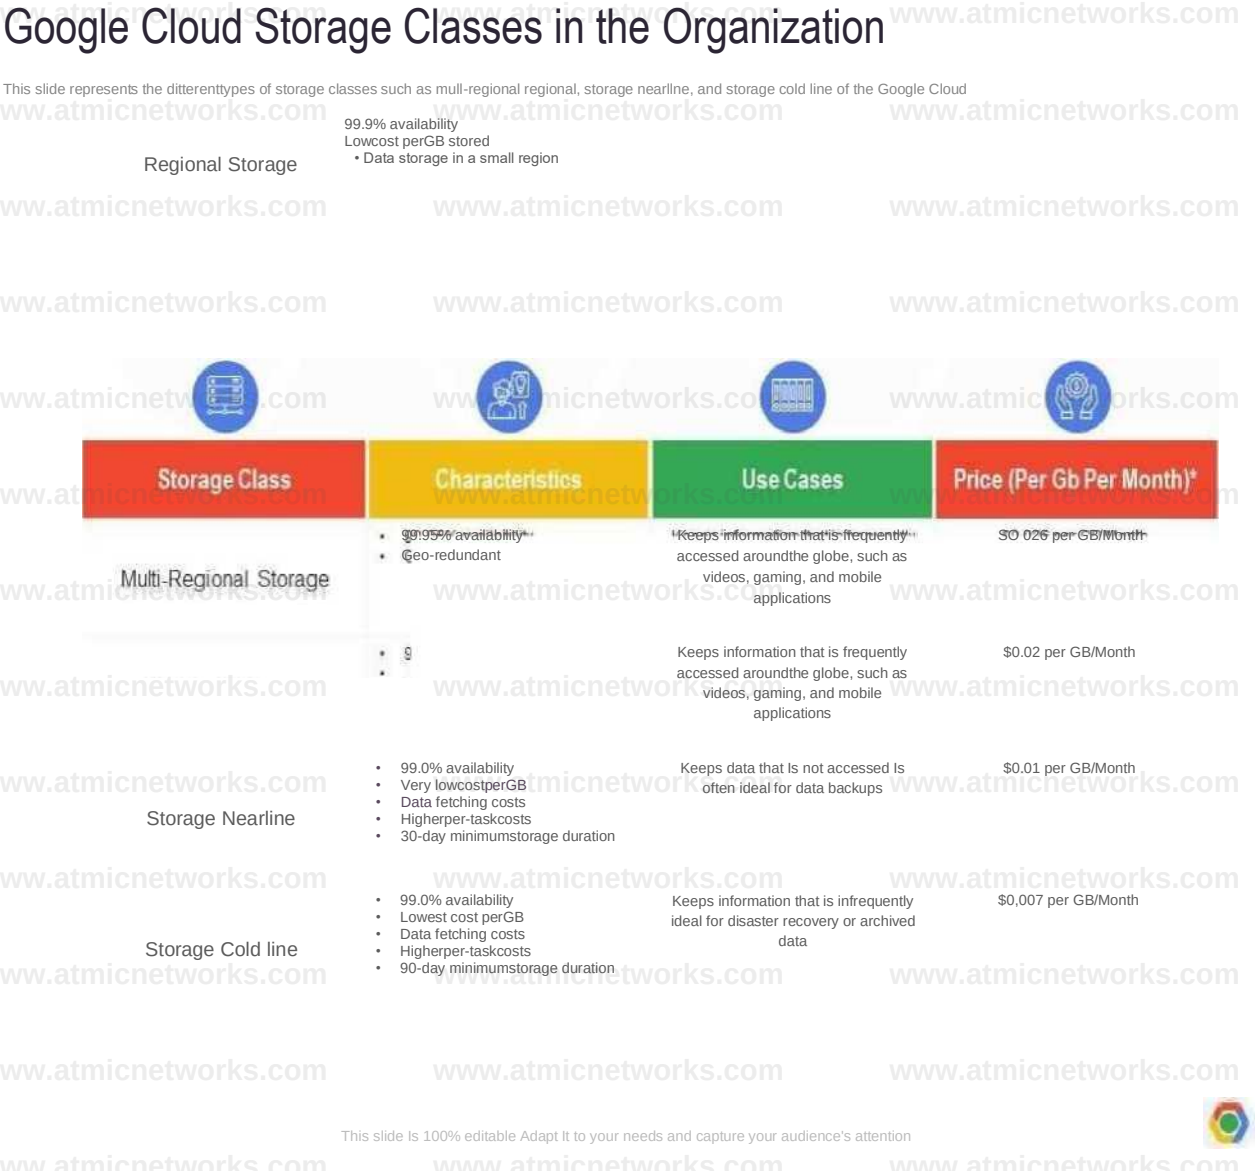

Question: 108

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

Question: 109

Your organization has user identities in Active Directory. Your organization wants to use Active Directory as their source of truth for identities. Your organization wants to have full control over the Google accounts used by employees for all Google services, including your Google Cloud Platform (GCP) organization. What should you do?

A. Use Google Cloud Directory Sync (GCDS) to synchronize users into Cloud Identity.
B. Use the cloud Identity APIs and write a script to synchronize users to Cloud Identity.
C. Export users from Active Directory as a CSV and import them to Cloud Identity via the Admin Console.
D. Ask each employee to create a Google account using self signup. Require that each employee use their company email address and password.

Answer: A

Explanation:

Reference: https://cloud.google.com/solutions/federating-gcp-with-active-directory-introduction

Directory Sync Google Cloud Directory Sync enables administrators to synchronize users, groups and other data from an Active Directory/LDAP service to their Google Cloud domain directory

https://tools.google.com/dlpage/dirsync/

Question: 110

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. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A

Explanation:
This aligns with Googles 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

Question: 111

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

Question: 112

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

Question: 113

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

A. Assign the auditor the IAM role roles/logging.privateLogViewer. Perform the export of logs to Cloud Storage.

B. Assign the auditor the IAM role roles/logging.privateLogViewer. Direct the auditor to also review the logs for changes to Cloud IAM policy.

C. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permission. Perform the export of logs to Cloud Storage.

D. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permission. Direct the auditor to also review the logs for changes to Cloud IAM policy.

Answer: B

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

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

Question: 114

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

A. Navigate to Stackdriver Logging and select resource.labels.project_id="*"

B. Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days.

C. Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule

to delete objects after 60 days.

D. Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days.

Answer: B

Explanation:
Navigate to Stackdriver Logging and select resource.labels.project_id=*. is not right.
Log entries are held in Stackdriver Logging for a limited time known as the retention period which is 30 days (default configuration). After that, the entries are deleted. To keep log entries longer, you need to export them outside of Stackdriver Logging by configuring log sinks.

Ref: https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud- audit-logging

Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days. is not right.
While this works, it makes no sense to use Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery when Google provides a feature (export sinks) that does exactly the same thing and works out of the box.

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

Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days. is not right.
You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub.
Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organication, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud organization.
Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks

Either way, we now have the data in Cloud Storage, but querying logs information from Cloud Storage is harder than Querying information from BigQuery dataset. For this reason, we should prefer Big Query over Cloud Storage.
Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days. is the right answer.

You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub.
Ref: https://cloud.google.com/logging/docs/export/configure_export_v2

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

Either way, we now have the data in a BigQuery Dataset. Querying information from a Big Query dataset is easier and quicker than analyzing contents in Cloud Storage bucket. As our requirement is to Quickly analyze the log contents, we should prefer Big Query over Cloud Storage.
Also, You can control storage costs and optimize storage usage by setting the default table expiration for newly created tables in a dataset. If you set the property when the dataset is created, any table created in the dataset is deleted after the expiration period. If you set the property after the dataset is created, only new tables are deleted after the expiration period.
For example, if you set the default table expiration to 7 days, older data is automatically deleted after 1 week.
Ref: https://cloud.google.com/bigquery/docs/best-practices-storage

Reference: https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google- cloud-audit- logging

Question: 115

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.

Question: 116

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 webapplications.
D. Give bigquery.dataViewer role to crm-databases-proj and appropriate roles to web-applications.

Answer: C

Explanation:

Reference: https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google- cloud-audit- logging
bigquery.dataViewer role provides permissions to read the datasets metadata and list tables in the dataset as well as Read data and metadata from the datasets tables. This is exactly what we need to fulfil this requirement and follows the least privilege principle.
Ref: https://cloud.google.com/iam/docs/understanding-roles#bigquery-roles

Question: 117

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 Stackdriver. Search for the user's email as the principal.
B. View System Event Logs in Stackdriver. Search for the service account associated with the user.
C. View Data Access audit logs in Stackdriver. Search for the user's email as the principal.
D. View the Admin Activity log in Stackdriver. Search for the service account associated with theuser.

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

Question: 118

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

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

Answer: A

Explanation:
When setting support levels for permissions in custom roles, you can set to one of SUPPORTED, TESTING or NOT_SUPPORTED.
Ref: https://cloud.google.com/iam/docs/custom-roles-permissions-support

Question: 119

Your company has a large quantity of unstructured data in different file formats. You want to perform ETL transformations on the dat
a. You need to make the data accessible on Google Cloud so it can be processed by a Dataflow job.
What should you do?

A. Upload the data to BigQuery using the bq command line tool.
B. Upload the data to Cloud Storage using the gsutil command line tool.
C. Upload the data into Cloud SQL using the import function in the console.
D. Upload the data into Cloud Spanner using the import function in the console.

Answer: B

Explanation:
"large quantity" : Cloud Storage or BigQuery "files" a file is nothing but an Object

Reference: https://cloud.google.com/solutions/performing-etl-from-relational-database-into- bigquery

Question: 120

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?

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

Question: 121

Your managed instance group raised an alert stating that new instance creation has failed to create new instances. You need to maintain the number of running instances specified by the template to be able to process expected application traffic. What should you do?

A. Create an instance template that contains valid syntax which will be used by the instance group. Delete any persistent disks with the same name as instance names.

B. Create an instance template that contains valid syntax that will be used by the instance group. Verify that the instance name and persistent disk name values are not the same in the template. C. Verify that the instance template being used by the instance group contains valid syntax. Delete any persistent disks with the same name as instance names. Set the disks.autoDelete property to true in the instance template.

D. Delete the current instance template and replace it with a new instance template. Verify that the instance name and persistent disk name values are not the same in the template. Set the disks.autoDelete property to true in the instance

template.

Answer: A

Explanation:
https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-migs
https://cloud.google.com/compute/docs/instance-templates#how_to_update_instance_templates

Question: 122

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

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

Answer: B

Explanation:
https://medium.com/google-cloud/how-to-deploy-cassandra-and-connect-on-google-cloud- platform-with-a-few-clicks-11ee3d7001d1
https://cloud.google.com/blog/products/databases/open-source-cassandra-now-managed-on- google-cloud
https://cloud.google.com/marketplace
You can deploy Cassandra as a Service, called Astra, on the Google Cloud Marketplace. Not only do you get a unified bill for all GCP services, your Developers can now create Cassandra clusters on Google Cloud in minutes and build applications with Cassandra as a database as a service without the operational overhead of managing Cassandra

Question: 123

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.
3. Configure your email address in the notification channel.

C. 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.
3. Create an uptime check for the instance in Stackdriver.

D. 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.
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

Question: 124

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. SREs would scale resources up or down accordingly.

C. Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshold. Google support would scale resources up or down accordingly. D. Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshold. 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

Question: 125

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. 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 dollars. Schedule the Cloud Function using Cloud Scheduler to run hourly.

D. 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 dollars. 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

Question: 126

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

A. For each GCP product in the solution, review the pricing details on the products pricing page. Use the pricing calculator to total the monthly costs for each GCP product.

B. For each GCP product in the solution, review the pricing details on the products pricing page. Create a Google Sheet that summarizes the expected monthly costs for each product.

C. Provision the solution on GCP. Leave the solution provisioned for 1 week. Navigate to the Billing Report page in the Google Cloud Platform Console. Multiply the 1 week cost to determine the monthly costs.

D. Provision the solution on GCP. Leave the solution provisioned for 1 week. Use Stackdriver to determine the provisioned and used resource amounts. Multiply the 1 week cost to determine the monthly costs.

Answer: A

Explanation:

You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You dont incur any charges for doing so.

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

Question: 127

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 Balancer. Add a firewall rule allowing ingress traffic from 0.0.0.0/0 on the target instances.

Answer: C

Explanation:

Reference: https://cloud.google.com/load-balancing/docs/ssl

SSL Proxy Load Balancing support for the following ports: 25, 43, 110, 143, 195, 443, 465, 587, 700, 993, 995, 1883, 3389, 5222, 5432, 5671, 5672, 5900, 5901, 6379, 8085, 8099, 9092, 9200, and 9300. When you use Google- managed SSL certificates with SSL Proxy Load Balancing, the frontend port for traffic must be 443 to enable the Google-managed SSL certificates to be provisioned and renewed. https://cloud.google.com/load-balancing/images/choose-lb.svg

Question: 128

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.
Reference: https://cloud.google.com/compute/docs/disks/performance

Local SSDs
Local SSDs are physically attached to the server that hosts your VM instance. Local SSDs have higher throughput and lower latency than standard persistent disks or SSD persistent disks. The data that you store on a local SSD persists only until the instance is stopped or deleted. Each local SSD is 375 GB in size, but you can attach a maximum of 24 local SSD partitions for a total of 9 TB per instance.

Performance
Local SSDs are designed to offer very high IOPS and low latency. Unlike persistent disks, you must manage the striping on local SSDs yourself. Combine multiple local SSD partitions into a single logical volume to achieve the best local SSD performance per instance, or format local SSD partitions individually.

Local SSD performance depends on which interface you select. Local SSDs are available through both SCSI and NVMe interfaces.

Question: 129

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

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

Answer: A

Explanation:
/25:
CIDR to IP Range
Result
CIDR Range 172.16.20.128/25
Netmask 255.255.255.128
Wildcard Bits 0.0.0.127
First IP 172.16.20.128
First IP (Decimal) 2886734976
Last IP 172.16.20.255
Last IP (Decimal) 2886735103
Total Host 128
CIDR
172.16.20.128/25

/24:
CIDR to IP Range
Result
CIDR Range 172.16.20.128/24
Netmask 255.255.255.0

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

Question: 130

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

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

Answer: B

Explanation:

The resource that you need to get access is in the other project.

roles/bigquery.dataViewer BigQuery Data Viewer
When applied to a table or view, this role provides permissions to:

Read data and metadata from the table or view.
This role cannot be applied to individual models or routines.

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

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

Question: 131

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

A. Create a Compute Engine snapshot of your base VM. Create your images from that snapshot.
B. Create a Compute Engine snapshot of your base VM. Create your instances from that snapshot.
C. Create a custom Compute Engine image from a snapshot. Create your images from that image.
D. Create a custom Compute Engine image from a snapshot. Create your instances from that image.

Answer: D

Explanation:

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

Reference: https://cloud.google.com/compute/docs/instances/create-start-instance

Preparing your instance for an image
You can create an image from a disk even while it is attached to a running VM instance. However, your image will be more reliable if you put the instance in a state that is easier for the image to

capture. Use one of the following processes to prepare your boot disk for the image:

Stop the instance so that it can shut down and stop writing any data to the persistent disk.

If you can't stop your instance before you create the image, minimize the amount of writes to the disk and sync your file system.

Pause apps or operating system processes that write data to that persistent disk.
Run an app flush to disk if necessary. For example, MySQL has a FLUSH statement. Other apps might have similar processes.
Stop your apps from writing to your persistent disk.
Run sudo sync.
After you prepare the instance, create the image.

https://cloud.google.com/compute/docs/images/create-delete-deprecate-private-images#prepare_instance_for_image

Question: 132

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

Explanation:

Reference: https://cloud.google.com/error-reporting/docs/setup/compute-engine

Cloud Loging knows nothing about applications installed on the system without an agent collecting logs. Using the serial console is not a best-practice and is impractical on a large scale.
The VM images for Compute Engine and Amazon Elastic Compute Cloud (EC2) don't include the Logging agent, so you must complete these steps to install it on those instances. The agent runs under both Linux and Windows. Source: https://cloud.google.com/logging/docs/agent/logging/installation

Question: 133

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. Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
D. 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

Reference: https://cloud.google.com/billing/docs/onboarding-checklist

Question: 134

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

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

Question: 135

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

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

Answer: A

Explanation:

https://cloud.google.com/sdk/gcloud/reference/projects/create

Quickstart: Creating a New Instance Using the Command Line

Before you begin

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

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

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

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

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

Question: 136

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 ofline, 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. Start and stop the instance as needed.
D. Create an Instance Template with Preemptible VMs On. Create a Managed Instance Group from the template and adjust Target CPU Utilization. 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 ofline 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.

Question: 137

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

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

Answer: A

Explanation:

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

Go to the Cloud Marketplace solution for Jenkins.

Click Launch on Compute Engine.

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

Machine type selection for Jenkins deployment.

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

Jenkins has been deployed.

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

Question: 138

You have downloaded and installed the gcloud command line interface (CLI) and have authenticated with your Google Account. Most of your Compute Engine instances in your project run in the europe- west1-d zone. You want to avoid having to specify this zone with each CLI command when managing these instances. What should you do?

A. Set the europe-west1-d zone as the default zone using the gcloud config subcommand.
B. In the Settings page for Compute Engine under Default location, set the zone to europe–west1-d.
C. In the CLI installation directory, create a file called default.conf containing zone=europe–west1–d.
D. Create a Metadata entry on the Compute Engine page with key compute/zone and value europe– west1–d.

Answer: A

Explanation:

Change your default zone and region in the metadata server Note: This only applies to the default configuration. You can change the default zone and region in your metadata server by making a

request to the metadata server. For example: gcloud compute project-info add-metadata \ -metadata google-compute-default-region=europe-west1,google-compute-default-zone=europe- west1-b The gcloud command-line tool only picks up on new default zone and region changes after you rerun the gcloud init command. After updating your default metadata, run gcloud init to reinitialize your default configuration. https://cloud.google.com/compute/docs/gcloud-

compute#change_your_default_zone_and_region_in_the_metadata_server

Question: 139

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

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

Answer: D

Explanation:

Keyword need to look for

- "High Throughput",
- "Consistent",
- "Property based data insert/fetch like ngine status, distance traveled, fuel level, and more." which can be designed in column,
- "Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds" This will go for pera bytes situation,
- Export data based on the time of the event.
- Atomic

o BigTable will fit all requirement.
o DataStore is not fully Atomic
o CloudStorage is not a option where we can export data based on time of event. We need another solution to do that
o FireStore can be used with MobileSDK.

Question: 140

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

Question: 141

You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

A. Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
B. Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
C. Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
D. Use the Reports view in the Cloud Billing Console to view the desired cost information.

Answer: A

Explanation:

https://cloud.google.com/billing/docs/how-to/export-data-bigquery "Cloud Billing export to BigQuery enables you to export detailed Google Cloud billing data (such as usage, cost estimates, and pricing data) automatically throughout the day to a BigQuery dataset that you specify."

Reference: https://cloud.google.com/billing/docs/how-to/visualize-data

Question: 142

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

Question: 143

Your team maintains the infrastructure for your organization. The current infrastructure requires changes. You need to share your proposed changes with the rest of the team. You want to follow Google's recommended best practices. What should you do?

A. Use Deployment Manager templates to describe the proposed changes and store them in a Cloud Storage bucket.
B. Use Deployment Manager templates to describe the proposed changes and store them in Cloud Source Repositories.
C. Apply the change in a development environment, run gcloud compute instances list, and then save the output in a shared Storage bucket.
D. Apply the change in a development environment, run gcloud compute instances list, and then save the output in Cloud Source Repositories.

Answer: B

Explanation:

Showing Deployment Manager templates to your team will allow you to define the changes you want to implement in your cloud infrastructure. You can use Cloud Source Repositories to store Deployment Manager templates and collaborate with your team. Cloud Source Repositories are fully-featured, scalable, and private Git repositories you can use to store, manage and track changes to your code.
https://cloud.google.com/source-repositories/docs/features

Question: 144

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

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

Answer: B

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

Question: 145

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

A. Use gcloud container clusters upgrade. Deploy the new services.
B. Create a new Node Pool and specify machine type n2–highmem–16. Deploy the new pods.
C. Create a new cluster with n2–highmem–16 nodes. Redeploy the pods and delete the old cluster.
D. Create a new cluster with both n1–standard–2 and n2–highmem–16 nodes. Redeploy the pods and delete the old cluster.

Answer: B

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

Question: 146

You have an application that uses Cloud Spanner as a database backend to keep current state information about users. Cloud Bigtable logs all events triggered by users. You export Cloud Spanner data to Cloud Storage during daily backups. One of your analysts asks you to join data from Cloud Spanner and Cloud Bigtable for specific users. You want to complete this ad hoc request as efficiently as possible. What should you do?

A. Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
B. Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
C. Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
D. Create two separate BigQuery external tables on Cloud Storage and Cloud Bigtable. Use the BigQuery console to join these tables through user fields, and apply appropriate filters.

Answer: D

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

text files."
https://cloud.google.com/dataflow/docs/guides/templates/provided-batch#cloudspannertogcstext

"The Dataflow connector for Cloud Spanner lets you read data from and write data to Cloud Spanner in a Dataflow pipeline"
https://cloud.google.com/spanner/docs/dataflow-connector

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

Question: 147

You are hosting an application from Compute Engine virtual machines (VMs) in us–central1–a. You want to adjust your design to support the failure of a single Compute Engine zone, eliminate downtime, and minimize cost. What should you do?

A. – Create Compute Engine resources in us–central1–b.
– Balance the load across both us–central1–a and us–central1–b.
B. – Create a Managed Instance Group and specify us–central1–a as the zone.
– Configure the Health Check with a short Health Interval.
C. – Create an HTTP(S) Load Balancer.
– Create one or more global forwarding rules to direct traffic to your VMs.
D. – Perform regular backups of your application.
– Create a Cloud Monitoring Alert and be notified if your application becomes unavailable.
– Restore from backups when notified.

Answer: A

Explanation:

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

Distribute your resources across multiple zones and regions to tolerate outages. Google designs zones to be independent from each other: a zone usually has power, cooling, networking, and control planes that are isolated from other zones, and most single failure events will affect only a single zone. Thus, if a zone becomes unavailable, you can transfer traffic to another zone in the same region to keep your services running. Similarly, if a region experiences any disturbances, you should have backup services running in a different region. For more information about distributing your resources and designing a robust system, see Designing Robust Systems. Decreased network latency To decrease network latency, you might want to choose a region or zone that is close to your point of service. https://cloud.google.com/compute/docs/regions-zones#choosing_a_region_and_zone

Question: 148

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

Question: 149

You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

A. Create a new subnet in the same region as the subnet being used.
B. Add an alias IP range to the subnet used by the GKE clusters.
C. Create a new VPC, and set up VPC peering with the existing VPC.
D. Expand the CIDR range of the relevant subnet for the cluster.

Answer: A

Explanation:

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

Question: 150

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

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

Answer: A

Explanation:

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

Question: 151

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

Question: 152

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

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

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

Question: 153

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

Question: 154

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

A. Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.

B. Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.

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

D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

Answer: B

Explanation:

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

Standard Storage

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

Archive Storage

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

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

Question: 155

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 dat
a. 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 bucker 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

Question: 156

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

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

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

Reference: https://cloud.google.com/billing/docs/how-to/budgets

Question: 157

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

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

Answer: D

Explanation:

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

Question: 158

Your company uses a large number of Google Cloud services centralized in a single project. All teams have specific projects for testing and development. The DevOps team needs access to all of the

production services in order to perform their job. You want to prevent Google Cloud product changes from broadening their permissions in the future. You want to follow Google-recommended practices. What should you do?

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

D. Create a custom role that combines the required permissions. Grant the DevOps team the custom role on the organization level.

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

Question: 159

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

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

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

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

Answer: AB

Explanation:

(A) Object Lifecycle Management

Delete

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

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

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

(B) Signed URLs

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

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

Question: 160

Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

A. Assign the appropriate permissions, and then use Cloud Monitoring to review metrics
B. Use the export logs API to provide the Admin Activity Audit Logs in the format they want
C. Turn on Data Access Logs for the buckets they want to audit, and Then build a query in the log viewer that filters on Cloud Storage
D. Assign the appropriate permissions, and then create a Data Studio report on Admin Activity Audit Logs

Answer: C

Explanation:

Types of audit logs Cloud Audit Logs provides the following audit logs for each Cloud project, folder, and organization: Admin Activity audit logs Data Access audit logs System Event audit logs Policy Denied audit logs ***Data Access audit logs contain API calls that read the configuration or metadata

of resources, as well as user-driven API calls that create, modify, or read user-provided resource data.

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

https://cloud.google.com/logging/docs/audit#data-access Cloud Storage: When Cloud Storage usage logs are enabled, Cloud Storage writes usage data to the Cloud Storage bucket, which generates Data Access audit logs for the bucket. The generated Data Access audit log has its caller identity redacted.

Question: 161

You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their application on Google Cloud. They manage the resources in their own project, but they need access to the BigQuery dataset in your project. You want to provide the partner company with access to the dataset What should you do?

A. Create a Service Account in your own project, and grant this Service Account access to BigGuery in your project
B. Create a Service Account in your own project, and ask the partner to grant this Service Account access to BigQuery in

their project

C. Ask the partner to create a Service Account in their project, and have them give the Service Account access to BigQuery in their project

D. Ask the partner to create a Service Account in their project, and grant their Service Account access to the BigQuery dataset in your project

Answer: D

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

Question: 162

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 lo Cloud Storage" Schedule the batch job on the desired interval

B. In the Cloud Console, go to Cloud Storage Upload the relevant images to the appropriate bucket C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage Schedule the script as a cron job

D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topic. Create an application that sends all medical images to the Pub/Sub lope

Answer: C

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

Question: 163

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

A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.

C. Deploy the container on App Engine flexible environment with autoscaling. and set the value min_instances to zero in the app yaml

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

Answer: B

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

Question: 164

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

Question: 165

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

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

Answer: D

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

Question: 166

You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database

for three years for audit purposes. What should you do?

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

Answer: C

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

Question: 167

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

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

Answer: D

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

Question: 168

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

A.
1. Verify that you ace assigned the Billing Administrator IAM role tor your organization's Google Cloud Project for the Marketing department
2. Link the new project to a Marketing Billing Account

B.
1. Verify that you are assigned the Billing Administrator IAM role for your organization's Google Cloud account

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

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

Answer: A

Explanation:

Question: 169

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

A. Use a custom mode VPC network, configure static routes, and use active/passive routing
B. Use an automatic mode VPC network, configure static routes, and use active/active routing
C. Use a custom mode VPC network use Cloud Router border gateway protocol (86P) routes, and use active/passive routing
D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Answer: C

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

Question: 170

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

Question: 171

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

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

Answer: B

Explanation:
https://cloud.google.com/sql/docs/postgres

Question: 172

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 learn 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 protect with the ID acme-marketing-data-digest for the Marketing team and deploy the resources there
D. Create a new protect named Meeting Data Digest and use the ID acme-data-digest Grant the Project Editor role to the Marketing team.

Answer: C

Explanation:

Question: 173

You have created an application that is packaged into a Docker image. You want to deploy the Docker image as a workload on Google Kubernetes Engine. What should you do?

A. Upload the image to Cloud Storage and create a Kubernetes Service referencing the image.
B. Upload the image to Cloud Storage and create a Kubernetes Deployment referencing the image.
C. Upload the image to Container Registry and create a Kubernetes Service referencing the image.
D. Upload the image to Container Registry and create a Kubernetes Deployment referencing the image.

Answer: D

Explanation:
A deployment is responsible for keeping a set of pods running. A service is responsible for enabling network access to a set of pods.

Question: 174

You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images. The microservice responsible for the image rendering requires a large amount of CPU time compared to the memory it requires. The other microservices are workloads that are optimized for n1-standard machine types. You need to optimize your cluster so that all workloads are using resources as efficiently as possible. What should you do?

A. Assign the pods of the image rendering microservice a higher pod priority than the older microservices
B. Create a node pool with compute-optimized machine type nodes for the image rendering microservice Use the node pool with general-purpose machine type nodes for the other microservices
C. Use the node pool with general-purpose machine type nodes for lite mage rendering microservice Create a nodepool with compute-optimized machine type nodes for the other microservices D. Configure the required amount of CPU and memory in the resource requests specification of the image rendering microservice deployment Keep the resource requests for the other microservices at the default

Answer: B

Explanation:

Question: 175

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

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

Answer: B

Explanation:

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

Question: 176

You need to track and verity modifications to a set of Google Compute Engine instances in your Google Cloud project. In particular, you want to verify OS system patching events on your virtual machines (VMs). What should you do?

A. Review the Compute Engine activity logs Select and review the Admin Event logs
B. Review the Compute Engine activity logs Select and review the System Event logs
C. Install the Cloud Logging Agent In Cloud Logging review the Compute Engine syslog logs
D. Install the Cloud Logging Agent In Cloud Logging, review the Compute Engine operation logs

Answer: A

Explanation:

Question: 177

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

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

Answer: B

Explanation:
https://cloud.google.com/iam/docs/best-practices-for-using-and-managing-service-accounts
If an application uses third-party or custom identities and needs to access a resource, such as a BigQuery dataset or a Cloud Storage bucket, it must perform a transition between principals. Because Google Cloud APIs don't recognize third-party or custom identities, the application can't propagate the end-user's identity to BigQuery or Cloud Storage. Instead, the application has to perform the access by using a different Google identity.

Question: 178

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

Explanation:

Question: 179

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

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

Answer: A

Explanation:

Question: 180

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

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

Answer: B

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

Question: 181

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

Explanation:

Question: 182

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

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

Answer: B

Explanation:
Folders are nodes in the Cloud Platform Resource Hierarchy. A folder can contain projects, other folders, or a combination of both. Organizations can use folders to group projects under the organization node in a hierarchy. For example, your organization might contain multiple departments, each with its own set of Google Cloud resources. Folders allow you to group these resources on a per-department basis. Folders are used to group resources that share common IAM policies. While a folder can contain

multiple folders or resources, a given folder or resource can have exactly one parent. https://cloud.google.com/resource-manager/docs/creating-managing-folders

Question: 183

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 SL A. 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

Question: 184

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

A. Invite the user to transfer their existing account
B. Invite the user to use an email alias to resolve the conflict
C. Tell the user that they must delete their existing account
D. Tell the user to remove all personal email from the existing account

Answer: A

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

Question: 185

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: C

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

Question: 186

You are assigned to maintain a Google Kubernetes Engine (GKE) cluster named dev that was deployed on Google Cloud. You want to manage the GKE configuration using the command line interface (CLI). You have just downloaded and installed the Cloud SDK. You want to ensure that future CLI commands by default address this specific cluster. What should you do?

A. Use the command gcloud config set container/cluster dev.
B. Use the command gcloud container clusters update dev.
C. Create a file called gke.default in the ~/.gcloud aname.
D. Create a file called defaults.json in the ~/.gcloud folder that contains the cluster name.

Answer: A

Explanation:
To set a default cluster for gcloud commands, run the following command: gcloud config set container/cluster CLUSTER_NAME
https://cloud.google.com/kubernetes-engine/docs/how- to/managing-clusters?hl=en

Question: 187

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

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

Answer: A

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

Question: 188

You are assisting a new Google Cloud user who just installed the Google Cloud SDK on their VM. The server needs access to Cloud Storage. The user wants your help to create a new storage bucket. You need to make this change in multiple environments. What should you do?

A. Use a Deployment Manager script to automate creating storage buckets in an appropriate region
B. Use a local SSD to improve performance of the VM for the targeted workload

C. Use the gsutii command to create a storage bucket in the same region as the VM
D. Use a Persistent Disk SSD in the same zone as the VM to improve performance of the VM

Answer: A

Explanation:

Question: 189

You received a JSON file that contained a private key of a Service Account in order to get access to several resources in a Google Cloud project. You downloaded and installed the Cloud SDK and want to use this private key for authentication and authorization when performing gcloud commands. What should you do?

A. Use the command gcloud auth login and point it to the private key
B. Use the command gcloud auth activate-service-account and point it to the private key
C. Place the private key file in the installation directory of the Cloud SDK and rename it to "credentials ison"
D. Place the private key file in your home directory and rename it to ''GOOGLE_APPUCATION_CREDENTiALS".

Answer: B

Explanation:
Authorizing with a service account
gcloud auth activate-service-account authorizes access using a service account. As with gcloud init and gcloud auth login, this command saves the service account credentials to the local system on successful completion and sets the specified account as the active account in your Cloud SDK configuration.

https://cloud.google.com/sdk/docs/authorizing#authorizing_with_a_service_account

Question: 190

You are managing a Data Warehouse on BigQuery. An external auditor will review your company's processes, and multiple external consultants will need view access to the dat
a. You need to provide them with view access while following Google-recommended practices. What should you do?
A. Grant each individual external consultant the role of BigQuery Editor
B. Grant each individual external consultant the role of BigQuery Viewer
C. Create a Google Group that contains the consultants and grant the group the role of BigQuery Editor
D. Create a Google Group that contains the consultants, and grant the group the role of BigQuery Viewer

Answer: D

Explanation:

Question: 191

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

Question: 192

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

A. Configure Cloud NAT for all subnets of your VPC to be used when egressing from the VM instances.

B. Create a private zone on Cloud DNS, and configure the applications with the DNS name.
C. Configure the IP of the database as custom metadata for each instance, and query the metadata server.
D. Query the Compute Engine internal DNS from the applications to retrieve the IP of the database.

Answer: B

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

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

Question: 193

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 λ€Billing Account Userλ€ 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

Explanation:

Question: 194

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

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

Answer: B

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

Question: 195

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

Question: 196

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

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

Answer: C

Explanation:

Question: 197

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 dat
a. 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.

Question: 198

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

Question: 199

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

Explanation:

Question: 200

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

Question: 201

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

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

Answer: C

Explanation:
https://cloud.google.com/vpc/docs/vpc-peering

Question: 202

You have created a new project in Google Cloud through the gcloud command line interface (CLI) and linked a billing account. You need to create a new Compute
Engine instance using the CLI. You need to perform the prerequisite steps. What should you do?

A. Create a Cloud Monitoring Workspace.
B. Create a VPC network in the project.
C. Enable the compute googleapis.com API.
D. Grant yourself the IAM role of Compute Admin.

Answer: D

Explanation:

Question: 203

Your company has developed a new application that consists of multiple microservices. You want to

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

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

Answer: C

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

Question: 204

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

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

Answer: C

Explanation:
as gcloud config configurations list can help check for the existing configurations and activate can help switch to the configuration.
gcloud config configurations list lists existing named configurations
gcloud config configurations activate activates an existing named configuration

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

Question: 205

The storage costs for your application logs have far exceeded the project budget. The logs are currently being retained indefinitely in the Cloud Storage bucket myapp-gcp-ace-logs. You have been

asked to remove logs older than 90 days from your Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

A. Write a script that runs gsutil Is -| – gs://myapp-gcp-ace-logs/** to find and remove items older than 90 days. Schedule the script with cron.
B. Write a lifecycle management rule in JSON and push it to the bucket with gsutil lifecycle set config- json-file.
C. Write a lifecycle management rule in XML and push it to the bucket with gsutil lifecycle set configxml-file.
D. Write a script that runs gsutil Is -lr gs://myapp-gcp-ace-logs/** to find and remove items older than 90 days. Repeat this process every morning.

Answer: B

Explanation:

You write a lifecycle management rule in XML and push it to the bucket with gsutil lifecycle set config-xml-file. is not right.
gsutil lifecycle set enables you to set the lifecycle configuration on one or more buckets based on the configuration file provided. However, XML is not a valid supported type for the configuration file.

Ref: https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

Write a script that runs gsutil ls -lr gs://myapp-gcp-ace-logs/** to find and remove items older than 90 days. Repeat this process every morning. is not right.
This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

Write a script that runs gsutil ls -l gs://myapp-gcp-ace-logs/** to find and remove items older than 90 days. Schedule the script with cron. is not right.
This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

Write a lifecycle management rule in JSON and push it to the bucket with gsutil lifecycle set config- json-file. is the right answer.
You can assign a lifecycle management configuration to a bucket. The configuration contains a set of rules which apply to current and future objects in the bucket. When an object meets the criteria of one of the rules, Cloud Storage automatically performs a specified action on the object. One of the supported actions is to Delete objects. You can set up a lifecycle management to delete objects older than 90 days. gsutil lifecycle set enables you to set the lifecycle configuration on the bucket based on the configuration file. JSON is the only supported type for the configuration file. The config-json-file specified on the command line should be a path to a local file containing the lifecycle configuration JSON document.

Ref: https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

Ref: https://cloud.google.com/storage/docs/lifecycle

Question: 206

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

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

Answer: A

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

Question: 207

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

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

Answer: C

Explanation:

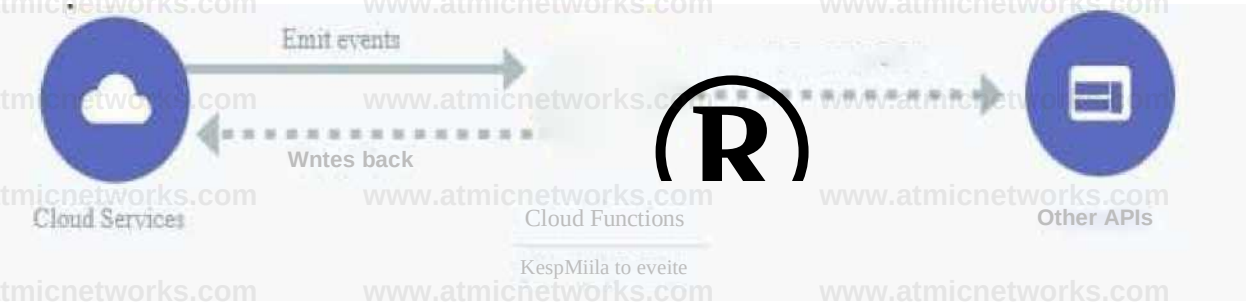

Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.

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

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

Question: 208

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

Regional *

Your data is stored in a specific region with replication across availability zones in that region.

Multi-regional

Your data is distributed redundantly across US, EU or Asia.

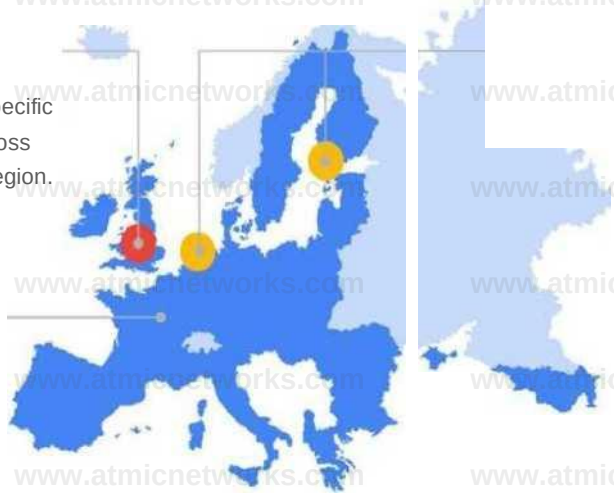

Dual-regional

Your data is replicated across a specific pair of regions.

Question: 209

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

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

Answer: A

Explanation:

Google HTTP(S) Load Balancing has native support for the WebSocket protocol when you use HTTP or HTTPS, not HTTP/2, as the protocol to the backend.

Ref: https://cloud.google.com/load-balancing/docs/https#websocket_proxy_support

So the next possible step is to Meet with the cloud enablement team to discuss load balancer options.
We dont need to convert WebSocket code to use HTTP streaming or Redesign the application, as WebSocket support is offered by Google HTTP(S) Load Balancing. Reviewing the encryption requirements is a good idea but it has nothing to do with WebSockets.

Question: 210

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. Then use the gcloud compute ssh command.
D. 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.

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/sdk/gcloud/reference/compute/ssh

Question: 211

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

Question: 212

You created a Kubernetes deployment by running kubectl run nginx image=nginx labels=app=prod. Your Kubernetes cluster is also used by a number of other deployments. How can you find the identifier of the pods for this nginx deployment?

A. kubectl get deployments –output=pods
B. gcloud get pods –selector=”app=prod”
C. kubectl get pods -l “app=prod”
D. gcloud list gke-deployments -filter={pod }

Answer: C

Explanation:
This command correctly lists pods that have the label app=prod. When creating the deployment, we used the label app=prod so listing pods that have this label retrieve the pods belonging to nginx deployments. You can list pods by using Kubernetes CLI kubectl get pods.

Ref: https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/ Ref: https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container- images/#list-containers-filtering-by-pod-label

Question: 213

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

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

NAME READY STATUS RESTARTS AGE
nginx-84748895c4-k6bzl 1/1 Running 0 25s
```

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

A. kubectl delete deployment nginx
B. kubectl delete –deployment=nginx
C. kubectl delete pod nginx-84748895c4-k6bzl –no-restart 2
D. kubectl delete inginx

Answer: A

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

$ kubectl delete deployment nginx deployment.apps nginx deleted Ref: https://kubernetes.io/docs/reference/kubectl/cheatsheet/#deleting-resources

Question: 214

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?

A. Cloud Run for Anthos
B. Cloud Run
C. App Engine Standard
D. Cloud Functions.

Answer: D

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

Question: 215

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

A. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace

B. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace:v1

C. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace

D. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace:v1

Answer: B

Explanation:

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

Question: 216

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. Create a JSON key for the service account. Execute the command gsutil signurl -m 1h gs:///*.

B. Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command gsutil signurl -d 1h gs:///**.

C. Create a service account with just the permissions to access files in the bucket. Create a JSON key for the service account. Execute the command gsutil signurl -p 60m gs:///.

D. Create a JSON key for the Default Compute Engine Service Account. 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

Question: 217

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

A. Your zonal capacity is limited, causing all preemptible VM's to be shutdown torecover capacity. Try deploying your group to another zone.

B. You have hit your instance quota for the region.

C. Your managed instance group's VM's are toggled to only last 1 minute inpreemptible settings.

D. Your managed instance group's health check is repeatedly failing, either to amisconfigured health check or misconfigured firewall rules not allowing the healthcheck to access the instance

Answer: D

Explanation:

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

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

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

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

the criteria that define a successful probe.

Question: 218

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

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

Answer: B

Explanation:

Question: 219

You are building a multi-player gaming application that will store game information in a database. As the popularity of the application increases, you are concerned about delivering consistent performance. You need to ensure an optimal gaming performance for global users, without increasing the management complexity. What should you do?

A. Use Cloud SQL database with cross-region replication to store game statistics in the EU, US, and APAC regions.
B. Use Cloud Spanner to store user data mapped to the game statistics.
C. Use BigQuery to store game statistics with a Redis on Memorystore instance in the front to provide global consistency.
D. Store game statistics in a Bigtable database partitioned by username.

Answer: B

Explanation:

Question: 220

Your company has multiple projects linked to a single billing account in Google Cloud. You need to visualize the costs with specific metrics that should be dynamically calculated based on companyspecific criteri
a. You want to automate the process. What should you do?
A. In the Google Cloud console, visualize the costs related to the projects in the Reports section.
B. In the Google Cloud console, visualize the costs related to the projects in the Cost breakdown section.
C. In the Google Cloud console, use the export functionality of the Cost table. Create a Looker Studiodashboard on top of the CSV export.
D. Configure Cloud Billing data export to BigOuery for the billing account. Create a Looker Studio dashboard on top of the BigQuery export.

Answer: D

Explanation:
Cloud Billing export to BigQuery enables you to export detailed Google Cloud billing data (such as usage, cost estimates, and pricing data) automatically throughout the day to a BigQuery dataset that you specify. Then you can access your Cloud Billing data from BigQuery for detailed analysis, or use a tool like Looker Studio to visualize your data.

Question: 221

You have an application that runs on Compute Engine VM instances in a custom Virtual Private Cloud (VPC). Your company's security policies only allow the use to internal IP addresses on VM instances and do not let VM instances connect to the internet. You need to ensure that the application can access a file hosted in a Cloud Storage bucket within your project. What should you do?

A. Enable Private Service Access on the Cloud Storage Bucket.
B. Add slorage.googleapis.com to the list of restricted services in a VPC Service Controls perimeter and add your project to the list to protected projects.
C. Enable Private Google Access on the subnet within the custom VPC.
D. Deploy a Cloud NAT instance and route the traffic to the dedicated IP address of the Cloud Storage bucket.

Answer: D

Explanation:

Question: 222

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

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

Answer: B

Explanation:

Question: 223

All development (dev) teams in your organization are located in the United States. Each dev team has its own Google Cloud project. You want to restrict access so that each dev team can only create cloud resources in the United States (US). What should you do?

A. Create a folder to contain all the dev projects Create an organization policy to limit resources in US locations.
B. Create an organization to contain all the dev projects. Create an Identity and Access Management (IAM) policy to limit the resources in US regions.
C. Create an Identity and Access Management <IAM) policy to restrict the resources locations in the US. Apply the policy to all dev projects.
D. Create an Identity and Access Management (IAM)policy to restrict the resources locations in all dev projects. Apply the policy to all dev roles.

Answer: A

Explanation:

Question: 224

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 lo 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 lo the load balancer, and create two NS records to point WWW and HOME to mydomain.com respectively.

Answer: C

Explanation:

Question: 225

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

A.

* Create service accounts sa-app and sa-db.
- Associate service account: sa-app with the application servers and the service account sa-db with the database servers.
- Create an ingress firewall rule to allow network traffic from source service account sa-app to target service account sa-db.

B.

- Create network tags app-server and db-server.
- Add the app-server lag lo the application servers and the db-server lag to the database servers.
- Create an egress firewall rule to allow network traffic from source network tag app-server to target network tag db-server.

C.

- Create a service account sa-app and a network tag db-server.

* Associate the service account sa-app with the application servers and the network tag db-server with the database servers.

- Create an ingress firewall rule to allow network traffic from source VPC IP addresses and target the subnet-a IP addresses.

D.

- Create a network lag app-server and service account sa-db.
- Add the tag to the application servers and associate the service account with the database servers.
- Create an egress firewall rule to allow network traffic from source network tag app-server to target service account sa-db.

Answer: C

Explanation:

Question: 226

Your learn wants to deploy a specific content management system (CMS) solution lo Google Cloud.

You need a quick and easy way to deploy and install the solution. What should you do?

A. Search for the CMS solution in Google Cloud Marketplace. Use gcloud CLI to deploy the solution.

B. Search for the CMS solution in Google Cloud Marketplace. Deploy the solution directly from Cloud Marketplace.

C. Search for the CMS solution in Google Cloud Marketplace. Use Terraform and the Cloud Marketplace ID to deploy the solution with the appropriate parameters.

D. Use the installation guide of the CMS provider. Perform the installation through your configuration management system.

Answer: B

Explanation:

Question: 227

You are working for a startup that was officially registered as a business 6 months ago. As your customer base grows, your use of Google Cloud increases. You want to allow all engineers to create new projects without asking them for their credit card information. What should you do?

A. Create a Billing account, associate a payment method with it, and provide all project creators with permission to associate that billing account with their projects.
B. Grant all engineer's permission to create their own billing accounts for each new project.
C. Apply for monthly invoiced billing, and have a single invoice tor the project paid by the finance team.
D. Create a billing account, associate it with a monthly purchase order (PO), and send the PO to Google Cloud.

Answer: A

Explanation:

Question: 228

You recently received a new Google Cloud project with an attached billing account where you will work. You need to create instances, set firewalls, and store data in Cloud Storage. You want to follow Google-recommended practices. What should you do?

A. Use the gcloud CLI services enable cloudresourcemanager.googleapis.com command to enable all resources.
B. Use the gcloud services enable compute.googleapis.com command to enable Compute Engine and the gcloud services enable storage-api.googleapis.com command to enable the Cloud Storage APIs.
C. Open the Google Cloud console and enable all Google Cloud APIs from the API dashboard.
D. Open the Google Cloud console and run gcloud init --project <project-id> in a Cloud Shell.

Answer: B

Explanation:

Question: 229

Your company is using Google Workspace to manage employee accounts. Anticipated growth will increase the number of personnel from 100 employees to 1.000 employees within 2 years. Most employees will need access to your company's Google Cloud account. The systems and processes will

need to support 10x growth without performance degradation, unnecessary complexity, or security issues. What should you do?

A. Migrate the users to Active Directory. Connect the Human Resources system to Active Directory. Turn on Google Cloud Directory Sync (GCDS) for Cloud Identity. Turn on Identity Federation from Cloud Identity to Active Directory.

B. Organize the users in Cloud Identity into groups. Enforce multi-factor authentication in Cloud Identity.

C. Turn on identity federation between Cloud Identity and Google Workspace. Enforce multi-factor authentication for domain wide delegation.

D. Use a third-party identity provider service through federation. Synchronize the users from Google Workplace to the third-party provider in real time.

Answer: B

Explanation:

Question: 230

Your application development team has created Docker images for an application that will be deployed on Google Cloud. Your team does not want to manage the infrastructure associated with this application. You need to ensure that the application can scale automatically as it gains popularity. What should you do?

A. Create an Instance template with the container image, and deploy a Managed Instance Group with Autoscaling.

B. Upload Docker images to Artifact Registry, and deploy the application on Google Kubernetes Engine using Standard mode.

C. Upload Docker images to the Cloud Storage, and deploy the application on Google Kubernetes Engine using Standard mode.

D. Upload Docker images to Artifact Registry, and deploy the application on Cloud Run.

Answer: D

Explanation:

Question: 231

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. Allow ingress traffic

from the IP range 35.235.240.0/20 on port 22.

D. Create a bastion host with public internet access. Create the SSH tunnel to the instance through the bastion host.

Answer: D

Explanation:

Question: 232

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

A. Store the application data on a zonal persistent disk. Create a snapshot schedule for the disk. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.

B. Store the application data on a zonal persistent disk. If an outage occurs, create an instance in another zone with this disk attached.

C. Store the application data on a regional persistent disk. Create a snapshot schedule for the disk. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.

D. Store the application data on a regional persistent disk If an outage occurs, create an instance in another zone with this disk attached.

Answer: D

Explanation:

Question: 233

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. instanceAdmln." permissions to the policy Attach the policy to the DevOps group.

C. Create a custom role at the folder level and grant all compute. instanceAdmln. * permissions to the role Grant the custom role to the DevOps group.

D. Grant the basic role roles/editor to the DevOps group.

Answer: D

Explanation:

Question: 234

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 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. Use the generated Docker images and deploy them to Cloud Function. Use the same configuration as on-premises.

C. Use the existing codebase and deploy each service as a separate Cloud Function Update the configurations and the required endpoints.

D. Use your existing codebase and deploy each service as a separate Cloud Run Use the same configurations as on-premises.

Answer: A

Explanation:

Question: 235

You are running a web application on Cloud Run for a few hundred users. Some of your users complain that the initial web page of the application takes much longer to load than the following pages. You want to follow Google's recommendations to mitigate the issue. What should you do?

A. Update your web application to use the protocol HTTP/2 instead of HTTP/1.1

B. Set the concurrency number to 1 for your Cloud Run service.

C. Set the maximum number of instances for your Cloud Run service to 100.

D. Set the minimum number of instances for your Cloud Run service to 3.

Answer: D

Explanation:

Question: 236

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

A. Use kubect1 to delete the topic resource.

B. Use gcloud CLI to delete the topic.

C. Use kubect1 to create the label deleted-by-cnrm and to change its value to true for the topic resource.

D. Use gcloud CLI to update the topic label managed-by-cnrm to false.

Answer: C

Explanation:

Question: 237

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: A

Explanation:

Question: 238

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

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

Answer: B

Explanation:

Question: 239

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

- A Flask web application
- AbackendAPI
- A scheduled long-running background job for ETL and reporting.

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

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

Answer: B

Explanation:

Question: 240

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

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

Answer: A

Explanation:

Question: 241

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

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

environment variables in your command line tool.

Answer: D

Explanation:

Question: 242

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 Goorj e Kubernfl:es Engine duster while optimizing for cost. What should you do?

A. Create a cluster with a single node-pool by using standard VMs. Label the fault-tolerant Deployments as spot-true.
B. Create a cluster with a single node-pool by using Spot VMs. Label the critical Deployments as spot- false.
C. Create a cluster with both a Spot W node pool and a rode pool by using standard VMs Deploy the critical.

deployments on the Spot VM node pool and the fault; tolerant deployments on the node pool by using standard VMs.

D. Create a cluster with both a Spot VM node pool and by using standard VMs. Deploy the critical deployments on the mode pool by using standard VMs and the fault-tolerant deployments on the Spot VM node pool.

Answer: C

Explanation:

Question: 243

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

A. Deploy the application lo Cloud. Run. Use gradual rollouts for traffic spelling.
B. Deploy the application lo Google Kubemetes Engine. Use Anthos Service Mesh for traffic splitting.
C. Deploy the application to Cloud functions. Saucily the version number in the functions name.
D. Deploy the application to App Engine. For each new version, create a new service.

Answer: A

Explanation:

Question: 244

Your company's security vulnerability management policy wonts 3 member of the security team to have visibility into vulnerabilities and other OS metadata for a specific Compute Engine instance This Compute Engine instance hosts a critical application in your Goggle Cloud project. You need to implement your company's security vulnerability management policy. What should you dc?

A.
- Ensure that the Ops Agent Is Installed on the Compute Engine instance.
- Create a custom metric in the Cloud Monitoring dashboard.
- Provide the security team member with access to this dashboard.

B.
- Ensure that the Ops Agent is installed on tie Compute Engine instance.
- Provide the security team member roles/configure.inventoryViewer permission.

C.
- Ensure that the OS Config agent Is Installed on the Compute Engine instance.
- Provide the security team member roles/configure.vulnerabilityViewer permission.

D.
- Ensure that the OS Config agent is installed on the Compute Engine instance
- Create a log sink Co a BigQuery dataset.
- Provide the security team member with access to this dataset.

Answer: C

Explanation:

Question: 245

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

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

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

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

Answer: C

Explanation:

Question: 246

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

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

Answer: A

Explanation:

Question: 247

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

A. Create a GKE Autopilot cluster. Enroll the cluster in the rapid release channel.
B. Create a GKE Autopilot cluster. Enroll the cluster in the stable release channel.
C. Create a zonal GKE standard cluster. Enroll the cluster in the stable release channel.
D. Create a regional GKE standard cluster. Enroll the cluster in the rapid release channel.

Answer: B

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

Question: 248

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%20organization%20level

Question: 249

You used the gcloud container clusters command to create two Google Cloud Kubernetes (GKE) clusters prod-cluster and dev-cluster.
- prod-cluster is a standard cluster.
- dev-cluster is an auto-pilot duster.

When you run the Kubect1 get nodes command, you only see the nodes from prod-cluster Which commands should you run to check the node status for dev-cluster?

A.

```
JClOlld CO^talnex dUSte-l get-CreientliilO dev-llustei
kubect1 qet nodes
```

B.

```
gcloud container clusters update -genexate-password dev-ciuster kub^ctl get nMes
```

C.

```
kubectl config set-context dev-cluster
kubectl cluster-info
```

D.

kubectl config set-credential./ dev-cluste kribectl du^tar-lnfo

A. Option A B. Option B C. Option C D. Option D

Answer: C

Explanation:

Question: 250

You have a Bigtable instance that consists of three nodes that store personally identifiable information (Pll) dat a. You need to log all read or write operations, including any metadata or configuration reads of this database table, in your company's Security Information and Event Management (SIEM) system. What should you do?

A. • Navigate to Cloud Mentioning in the Google Cloud console, and create a custom monitoring job for the Bigtable instance to track all changes.
• Create an alert by using webhook endpoints. with the SIEM endpoint as a receiver

B. • Navigate to the Audit Logs page in the Google Cloud console, and enable Data Read. Data Write and Admin Read logs for the Bigtable instance
• Create a Pub/Sub topic as a Cloud Logging sink destination, and add your SIEM as a subscriber to the topic.

C. • Install the Ops Agent on the Bigtable instance during configuration. K
• Create a service account with read permissions for the Bigtable instance.
• Create a custom Dataflow job with this service account to export logs to the company's SIEM system.

D. • Navigate to the Audit Logs page in the Google Cloud console, and enable Admin Write logs for the Biglable instance.
• Create a Cloud Functions instance to export logs from Cloud Logging to your SIEM.

Answer: D

Explanation:

Question: 251

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 Functions. Use Cloud Scheduler to start the application.

B. Create a container for the set of binaries. Use Cloud Scheduler to start a Cloud Run job for the container.

C. 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.

D. Lift and shift to a VM on Compute Engine. Use an instance schedule to start and stop the instance.

Answer: B

Explanation:

Question: 252

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: D

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.

Question: 253

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

A. Use Cloud Logging filters to create log-based metrics for firewall and instance actions. Monitor the changes and set up reasonable alerts.

B. Install Kibana on a compute Instance. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Pub/Sub. Target the Pub/Sub topic to push messages to the Kibana instance. Analyze the logs on Kibana in real time.

C. Turn on Google Cloud firewall rules logging, and set up alerts for any insert, update, or delete events.

D. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Cloud Storage. Use BigQuery to periodically analyze log events in the storage bucket.

Answer: A

Explanation:

Question: 254

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.

Question: 255

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

A. • Add a step for human approval to the CI/CD pipeline before the execution of the infrastructure provisioning.
• Use the human approvals IAM account for the provisioning.

B. • Attach a single service account to the compute instances.
• Add minimal rights to the service account.
• Allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources.

C. • Attach a single service account to the compute instances.
• Add all required Identity and Access Management (IAM) permissions to this service account to create, update, or delete resources

D. • Create multiple service accounts, one for each pipeline with the appropriate minimal Identity and Access Management (IAM) permissions.
• Use a secret manager service to store the key files of the service accounts.
• Allow the CI/CD pipeline to request the appropriate secrets during the execution of the pipeline.

Answer: B

Explanation:

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

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

Reference:

1: https://cloud.google.com/iam/docs/impersonating-service-accounts

2: https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys

3: https://cloud.google.com/iam/docs/understanding-service-accounts

Question: 256

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-sa. Enforce an org policy to deny service account key creation with an exception to pj-sa.

B. Implement a Kubernetes Cronjob to rotate all service account keys periodically. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

C. Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hours. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.

D. Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hours. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

Answer: A

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.

Reference:

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

Question: 257

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

A. Instruct the external consultant to use the gcloud compute ssh command line tool by using Identity-Aware Proxy to access the instance.

B. Instruct the external consultant to use the gcloud compute ssh command line tool by using the public IP address of the instance to access it.

C. Instruct the external consultant to generate an SSH key pair, and request the public key from the consultant. Add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key.

D. Instruct the external consultant to generate an SSH key pair, and request the private key from the consultant. Add the private key to the instance yourself, and have the consultant access the instance through SSH with their public key.

Answer: C

Explanation:

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

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

Reference:

1: https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys

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

3: https://cloud.google.com/compute/docs/instances/connecting-advanced#sshbetweeninstances

Question: 258

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

Choose 2 answers

A. Run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI.

B. Create a Google Cloud service account, and download the service account key. Place the key file in a folder on your machine where gcloud CLI can find it.

C. Download your Cloud Identity user account key. Place the key file in a folder on your machine where gcloud CLI can find it.

D. Run gcloud config set compute/zone $my_zone to set the default zone for gcloud CLI.

E. Run gcloud config set project $my_project to set the default project for gcloud CLI.

Answer: AE

Explanation:

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

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

To set the default project for gcloud CLI, you need to run gcloud config set project $my_project, where $my_project is the ID of the project that contains the instances you want to list. This will save you from having to specify the project flag for every gcloud command2.

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

Reference:

1: https://cloud.google.com/sdk/docs/authorizing

2: https://cloud.google.com/sdk/gcloud/reference/config/set

3: https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys

4: https://cloud.google.com/identity/docs/overview

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

Question: 259

During a recent audit of your existing Google Cloud resources, you discovered several users with email addresses outside of your Google Workspace domain.

You want to ensure that your resources are only shared with users whose email addresses match your domain. You need to remove any mismatched users, and you want to avoid having to audit your resources to identify mismatched users. What should you do?

A. Create a Cloud Scheduler task to regularly scan your projects and delete mismatched users.

B. Create a Cloud Scheduler task to regularly scan your resources and delete mismatched users.

C. Set an organizational policy constraint to limit identities by domain to automatically remove mismatched users.

D. Set an organizational policy constraint to limit identities by domain, and then retroactively remove the existing mismatched users.

Answer: D

Explanation:

https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints This list constraint defines the set of domains that email addresses added to Essential Contacts can have. By default, email addresses with any domain can be added to Essential Contacts. The allowed/denied list must specify one or more domains of the form @example.com. If this constraint is active and configured with allowed values, only email addresses with a suffix matching one of the entries from the list of allowed domains can be added in Essential Contacts. This constraint has no effect on updating or removing existing contacts.
constraints/essentialcontacts.allowedContactDomains

Question: 260

You are responsible for a web application on Compute Engine. You want your support team to be notified automatically if users experience high latency for at least 5 minutes. You need a Google- recommended solution with no development cost. What should you do?

A. Create an alert policy to send a notification when the HTTP response latency exceeds the specified threshold.

B. Implement an App Engine service which invokes the Cloud Monitoring API and sends a notification in case of anomalies.

C. Use the Cloud Monitoring dashboard to observe latency and take the necessary actions when the response latency exceeds the specified threshold.
D. Export Cloud Monitoring metrics to BigQuery and use a Looker Studio dashboard to monitor your web applications latency.

Answer: A

Explanation:
https://cloud.google.com/monitoring/alerts#alerting-example

Question: 261

Your team is building a website that handles votes from a large user population. The incoming votes will arrive at various rates. You want to optimize the storage and processing of the votes. What should you do?

A. Save the incoming votes to Firestore. Use Cloud Scheduler to trigger a Cloud Functions instance to periodically process the votes.
B. Use a dedicated instance to process the incoming votes. Send the votes directly to this instance.
C. Save the incoming votes to a JSON file on Cloud Storage. Process the votes in a batch at the end of the day.
D. Save the incoming votes to Pub/Sub. Use the Pub/Sub topic to trigger a Cloud Functions instance to process the votes.

Answer: B

Explanation:
Pub/Sub is a scalable and reliable messaging service that can handle large volumes of data from different sources at different rates. It allows you to decouple the producers and consumers of the data, and provides a durable and persistent storage for the messages until they are delivered. Cloud Functions is a serverless platform that can execute code in response to events, such as messages published to a Pub/Sub topic. It can scale automatically based on the load, and you only pay for the resources you use. By using Pub/Sub and Cloud Functions, you can optimize the storage and processing of the votes, as you can handle the variable rates of incoming votes, process them in real time or near real time, and avoid managing servers or VMs. Reference:
Pub/Sub documentation
Cloud Functions documentation
Choosing a messaging service for Google Cloud

Question: 262

Your team has developed a stateless application which requires it to be run directly on virtual machines. The application is expected to receive a fluctuating amount of traffic and needs to scale automatically. You need to deploy the application. What should you do?

A. Deploy the application on a managed instance group and configure autoscaling.
B. Deploy the application on a Kubernetes Engine cluster and configure node pool autoscaling.
C. Deploy the application on Cloud Functions and configure the maximum number instances.
D. Deploy the application on Cloud Run and configure autoscaling.

Answer: A

Explanation:
A managed instance group (MIG) is a group of identical virtual machines (VMs) that you can manage as a single entity. You can use a

MIG to deploy and maintain a stateless application that runs directly on VMs. A MIG can automatically scale the number of VMs based on the load or a schedule. A MIG

can also automatically heal the VMs if they become unhealthy or unavailable. A MIG is suitable for applications that need to run on VMs rather than containers or serverless platforms.

B is incorrect because Kubernetes Engine is a managed service for running containerized applications on a cluster of nodes. It is not necessary to use Kubernetes Engine if the application does not use containers and can run directly on VMs.

C is incorrect because Cloud Functions is a serverless platform for running event-driven code in response to triggers. It is not suitable for applications that need to run continuously and handle HTTP requests.

D is incorrect because Cloud Run is a serverless platform for running stateless containerized applications. It is not suitable for applications that do not use containers and can run directly on VMs. Reference:

Managed instance groups documentation

Choosing a compute option for Google Cloud

Question: 263

A colleague handed over a Google Cloud 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 Google Cloud 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 gcloud projects get-iam-policy command to view the current role assignments.

Answer: D

Explanation:

The gcloud projects get-iam-policy command displays the IAM policy for a project, which includes the roles and members assigned to those roles. The Project Owner role grants full access to all resources and actions in the project. By using this command, you can review who has been granted this role and make any necessary changes. Reference:

1: Associate Cloud Engineer Certification Exam Guide | Learn - Google Cloud

2: gcloud projects get-iam-policy | Cloud SDK Documentation | Google Cloud

3: Understanding roles | Cloud IAM Documentation | Google Cloud

Question: 264

You are deploying a web application using Compute Engine. You created a managed instance group (MIG) to host the application. You want to follow Google-recommended practices to implement a secure and highly available solution. What should you do?

A. Use SSL proxy load balancing for the MIG and an A record in your DNS private zone with the load balancer's IP address.
B. Use SSL proxy load balancing for the MIG and a CNAME record in your DNS public zone with the load balancer's IP address.
C. Use HTTP(S) load balancing for the MIG and a CNAME record in your DNS private zone with the load balancer's IP address.
D. Use HTTP(S) load balancing for the MIG and an A record in your DNS public zone with the load balancer's IP address.

Answer: D

Explanation:

HTTP(S) load balancing is a Google-recommended practice for distributing web traffic across multiple regions and zones, and providing high availability, scalability, and security for web applications. It supports both IPv4 and IPv6 addresses, and can handle SSL/TLS termination and encryption. It also integrates with Cloud CDN, Cloud Armor, and Cloud Identity-Aware Proxy for enhanced performance and protection. A MIG can be used as a backend service for HTTP(S) load balancing, and can automatically scale and heal the VM instances that host the web application.
To configure DNS for HTTP(S) load balancing, you need to create an A record in your DNS public zone with the load balancer's IP address. This will map your domain name to the load balancer's IP address, and allow users to access your web application using the domain name. A CNAME record is not recommended, as it can cause latency and DNS resolution issues. A private zone is not suitable, as it is only visible within your VPC network, and not to the public internet.
Reference:
HTTP(S) Load Balancing documentation
Setting up DNS records for HTTP(S) load balancing
Choosing a load balancer

Question: 265

You want to host your video encoding software on Compute Engine. Your user base is growing rapidly, and users need to be able 3 to encode their videos at any time without interruption or CPU limitations. You must ensure that your encoding solution is highly available, and you want to follow Google-recommended practices to automate operations. What should you do?

A. Deploy your solution on multiple standalone Compute Engine instances, and increase the number of existing instances wnen CPU utilization on Cloud Monitoring reaches a certain threshold.
B. Deploy your solution on multiple standalone Compute Engine instances, and replace existing instances with high-CPU instances when CPU utilization on Cloud Monitoring reaches a certain threshold.
C. Deploy your solution to an instance group, and increase the number of available instances whenever you see high CPU utilization in Cloud Monitoring.
D. Deploy your solution to an instance group, and set the autoscaling based on CPU utilization.

Answer: D

Explanation:
Instance groups are collections of virtual machine (VM) instances that you can manage as a single entity. Instance groups can help you simplify the management of multiple instances, reduce operational costs, and improve the availability and performance of your applications. Instance groups support autoscaling, which automatically adds or removes instances from the group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower. You can set the autoscaling policy
based on CPU utilization, load balancing capacity, Cloud Monitoring metrics, or a queue-based workload. In this case, since the video encoding software is CPU-intensive, setting the autoscaling based on CPU utilization is the best option to ensure high availability and optimal performance. Reference:
Instance groups
Autoscaling groups of instances

Question: 266

Your customer wants you to create a secure website with autoscaling based on the compute instance CPU load. You want to enhance performance by storing static content in Cloud Storage. Which resources are needed to distribute the user traffic?

A. An internal HTTP(S) load balancer together with Identity-Aware Proxy to allow only HTTPS traffic. B. An external HTTP(S) load

balancer to distribute the load and a URL map to target the requests for the static content to the Cloud Storage backend. Install the HTTPS certificates on the instance.

C. An external HTTP(S) load balancer with a managed SSL certificate to distribute the load and a URL map to target the requests for the static content to the Cloud Storage backend.

D. An external network load balancer pointing to the backend instances to distribute the load evenly. The web servers will forward the request to the Cloud Storage as needed.

Answer: C

Explanation:

An external HTTP(S) load balancer is a Google-recommended solution for distributing web traffic across multiple regions and zones, and providing high availability, scalability, and security for web applications. It supports both IPv4 and IPv6 addresses, and can handle SSL/TLS termination and encryption. It also integrates with Cloud CDN, Cloud Armor, and Cloud Identity-Aware Proxy for enhanced performance and protection. A managed instance group (MIG) can be used as a backend service for the HTTP(S) load balancer, and can automatically scale the number of VM instances based on the CPU load. A Cloud Storage bucket can also be used as a backend service for the HTTP(S) load balancer, and can serve static content such as images, videos, or HTML files. A URL map can be used to route requests to different backend services based on the path or host of the request. For example, a URL map can send requests for /static/* to the Cloud Storage bucket, and requests for /dynamic/* to the MIG. A managed SSL certificate can be used to secure the connection between the clients and the load balancer, and can be automatically provisioned and renewed by Google. A is incorrect because an internal HTTP(S) load balancer is only visible within a VPC network, and not to the public internet. It is used for internal applications that need to communicate with other internal services. Identity-Aware Proxy is a service that provides secure access to web applications without using a VPN. It is not a load balancer, and it does not distribute user traffic.

B is incorrect because installing HTTPS certificates on the instance is not necessary, as the HTTP(S) load balancer can handle SSL/TLS termination and encryption. It is also more complex and less secure to manage the certificates on the instance level, as they need to be updated and synchronized across multiple instances.

D is incorrect because an external network load balancer is a TCP/UDP load balancer that operates at the network layer. It is not suitable for web applications that use HTTP(S) protocols, as it does not support SSL/TLS termination and encryption, URL maps, or Cloud Storage backends. It is also less efficient and scalable to forward the requests to the Cloud Storage from the web servers, as it adds

an extra hop and latency.

Reference:

HTTP(S) Load Balancing documentation

Setting up HTTP(S) Load Balancing with Cloud Storage

Creating and using SSL certificates

Choosing a load balancer

Question: 267

Your manager asks you to deploy a workload to a Kubernetes cluster. You are not sure of the workloads resource requirements or how the requirements might vary depending on usage patterns, external dependencies, or other factors. You need a solution that makes cost-effective recommendations regarding CPU and memory requirements, and allows the workload to function consistently in any situation. You want to follow Google-recommended practices. What should you do?

A. Configure the Horizontal Pod Autoscaler for availability, and configure the cluster autoscaler for suggestions.

B. Configure the Horizontal Pod Autoscaler for availability, and configure the Vertical Pod Autoscaler recommendations for suggestions.

C. Configure the Vertical Pod Autoscaler recommendations for availability, and configure the Cluster autoscaler for

suggestions.

D. Configure the Vertical Pod Autoscaler recommendations for availability, and configure the Horizontal Pod Autoscaler for suggestions.

Answer: D

Explanation:

Question: 268

You need to extract text from audio files by using the Speech-to-Text API. The audio files are pushed to a Cloud Storage bucket. You need to implement a fully managed, serverless compute solution that requires authentication and aligns with Google-recommended practices. You want to automate the call to the API by submitting each file to the API as the audio file arrives in the bucket. What should you do?

A. Run a Kubernetes job to scan the bucket regularly for incoming files, and call the Speech-to-Text API for each unprocessed file.

B. Create an App Engine standard environment triggered by Cloud Storage bucket events to submit the file URI to the Google Speech-to-Text API.

C. Run a Python script by using a Linux cron job in Compute Engine to scan the bucket regularly for incoming files, and call the Speech-to-Text API for each unprocessed file.

D. Create a Cloud Function triggered by Cloud Storage bucket events to submit the file URI to the Google Speech-to-Text API.

Answer: B

Explanation:

Question: 269

You need to migrate invoice documents stored on-premises to Cloud Storage. The documents have the following storage requirements:

- Documents must be kept for five years.
- Up to five revisions of the same invoice document must be stored, to allow for corrections.
- Documents older than 365 days should be moved to lower cost storage tiers.

You want to follow Google-recommended practices to minimize your operational and development costs. What should you do?

A. Enable retention policies on the bucket, and use Cloud Scheduler to invoke a Cloud Function to move or delete your documents based on their metadata.

B. Enable retention policies on the bucket, use lifecycle rules to change the storage classes of the objects, set the number of versions, and delete old files.

C. Enable object versioning on the bucket, and use Cloud Scheduler to invoke a Cloud Functions instance to move or delete your documents based on their metadata.

D. Enable object versioning on the bucket, use lifecycle conditions to change the storage class of the objects, set the number of versions, and delete old files.

Answer: D

Explanation:

https://cloud.google.com/storage/docs/object-versioning

Question: 270

You have a VM instance running in a VPC with single-stack subnets. You need to ensure that the VM instance has a fixed IP address so that other services hosted in the same VPC can communicate with the VM. You want to follow Google-recommended practices while minimizing cost. What should you do?

A. Reserve a new static external IP address and assign the new IP address to the VM.
B. Promote the existing IP address of the VM to become a static external IP address.
C. Reserve a new static external IPv6 address and assign the new IP address to the VM.
D. Promote the existing IP address of the VM to become a static internal IP address.

Answer: B

Explanation:

Question: 271

Your company uses BigQuery to store and analyze dat

a. Upon submitting your query in BigQuery, the query fails with a quotaExceeded error. You need to diagnose the issue causing the error. What should you do?

Choose 2 answers

A. Search errors in Cloud Audit Logs to analyze the issue.
B. Configure Cloud Trace to analyze the issue.
C. View errors in Cloud Monitoring to analyze the issue.
D. Use the information schema views to analyze the underlying issue.
E. Use BigQuery Bl Engine to analyze the issue.

Answer: AC

Explanation:

When encountering a quotaExceeded error in BigQuery, you should follow these steps to diagnose and mitigate the issue:
Understand the Error:

The error message indicates that a quota was exceeded (either a short-term rate limit or a longer- term limit).
The response payload contains information about which quota was reached.

Quotas can fall into two categories:

rateLimitExceeded: Short-term limits. Retry the operation after a few seconds using exponential backoff.
quotaExceeded: Longer-term limits. Wait 10 minutes or longer before retrying the operation.

Search Errors in Cloud Audit Logs (Option A):

Cloud Audit Logs provide detailed information about API requests and responses.
By searching the logs, you can identify the specific API call that triggered the quotaExceeded error.
This helps you understand which resource or operation exceeded the quota.
View Errors in Cloud Monitoring (Option C):
Cloud Monitoring (formerly known as Stackdriver) provides insights into your Google Cloud resources.
Check the monitoring dashboard for any alerts related to BigQuery quotas.
You can set up custom monitoring rules to track specific quotas and receive notifications.
Other Options:
B . Configure Cloud Trace: Cloud Trace is used for performance analysis and latency tracking. It's not directly related to quota issues.
D . Use Information Schema Views: Information schema views provide metadata about your datasets and tables but won't help diagnose quota errors.
E . Use BigQuery Bl Engine: There is no such tool called "BigQuery Bl Engine." This option is invalid. Remember that some quotas replenish incrementally over a 24-hour period, so you don't always need to wait a full 24 hours after reaching the limit. If you consistently hit longer-term quotas, consider workload optimization or requesting a quota increase

Question: 272

Your web application is hosted on Cloud Run and needs to query a Cloud SOL database. Every morning during a traffic spike, you notice API quota errors in Cloud SOL logs. The project has already reached the maximum API quot
a. You want to make a configuration change to mitigate the issue. What should you do?
A. Modify the minimum number of Cloud Run instances.
B. Set a minimum concurrent requests environment variable for the application.
C. Modify the maximum number of Cloud Run instances.
D. Use traffic splitting.

Answer: C

Explanation:

Question: 273

You are planning to migrate the following on-premises data management solutions to Google Cloud:

- One MySQL cluster for your main database
- Apache Kafka for your event streaming platform
- One Cloud SOL for PostgreSOL database for your analytical and reporting needs

You want to implement Google-recommended solutions for the migration. You need to ensure that the new solutions provide global scalability and require minimal operational and infrastructure management. What should you do?

A. Migrate from MySQL to Cloud SQL, from Kafka to Memorystore, and from Cloud SQL for PostgreSQL to Cloud SQL
B. Migrate from MySQL to Cloud Spanner, from Kafka to Memorystore, and from Cloud SOL for PostgreSQL to Cloud SQL
C. Migrate from MySQL to Cloud SOL, from Kafka to Pub/Sub, and from Cloud SOL for PostgreSQL to BigQuery.
D. Migrate from MySQL to Cloud Spanner, from Kafka to Pub/Sub. and from Cloud SQL for PostgreSQL to BigQuery

Answer: D

Explanation:

Question: 274

You need to deploy a single stateless web application with a web interface and multiple endpoints. For security reasons, the web application must be reachable from an internal IP address from your company's private VPC and on-premises network. You also need to update the web application multiple times per day with minimal effort and want to manage a minimal amount of cloud infrastructure. What should you do?

A. Deploy the web application on Google Kubernetes Engine standard edition with an internal ingress.
B. Deploy the web application on Cloud Run with Private Google Access configured
C. Deploy the web application to GKE Autopilot with Private Google Access configured
D. Deploy the web application on Cloud Run with Private Service Connect configured.

Answer: A

Explanation:

Question: 275

Your application stores files on Cloud Storage by using the Standard Storage class. The application only requires access to files created in the last 30 days. You want to automatically save costs on files that are no longer accessed by the application. What should you do?

A. Create a retention policy on the storage bucket of 30 days, and lock the bucket by using a retention policy lock.
B. Enable object versioning on the storage bucket and add lifecycle rules to expire non-current versions after 30 days
C. Create an object lifecycle on the storage bucket to change the storage class to Archive Storage for objects with an age over 30 days.
D. Create a cron job in Cloud Scheduler to call a Cloud Functions instance every day to delete files older than 30 days.

Answer: C

Explanation:

Question: 276

You are a Google Cloud organization administrator. You need to configure organization policies and log sinks on Google Cloud projects that cannot be removed by project users to comply with your company's security policies. The security policies are different for each company department Each company department has a user with the Project Owner role assigned to their projects. What should you do?

A. Organize projects under folders for each department. Configure both organization policies and log sinks on the folders
B. Organize projects under folders for each department. Configure organization policies on the organization and log sinks on the folders.
C. Use a standard naming convention for projects that includes the department name. Configure organization policies on the organization and log sinks on the projects.
D. Use a standard naming convention for projects that includes the department name. Configure both organization policies and log sinks on the projects.

Answer: A

Explanation:

Question: 277

You are running out of primary internal IP addresses in a subnet for a custom mode VPC. The subnet has the IP range 10.0.0.0/20. and the IP addresses are primarily used by virtual machines in the project. You need to provide more IP addresses for the virtual machines. What should you do?

A. Change the subnet IP range from 10.0.0.0/20 to 10.0.0.0/22.
B. Change the subnet IP range from 10.0 0.0/20 to 10.0.0.0718.
C. Add a secondary IP range 10.1.0.0/20 to the subnet.
D. Convert the subnet IP range from IPv4 to IPv6

Answer: B

Explanation:

Question: 278

Your company is running a three-tier web application on virtual machines that use a MySQL database. You need to create an estimated total cost of cloud infrastructure to run this application on Google Cloud instances and Cloud SQL. What should you do?

A. Use the Google Cloud Pricing Calculator to determine the cost of every Google Cloud resource you expect to use. Use similar size instances for the web server, and use your current on-premises machines as a comparison for Cloud SQL.
B. Implement a similar architecture on Google Cloud, and run a reasonable load test on a smaller scale. Check the billing information, and calculate the estimated costs based on the real load your system usually handles.
C. Use the Google Cloud Pricing Calculator and select the Cloud Operations template to define your web application with as much detail as possible.
D. Create a Google spreadsheet with multiple Google Cloud resource combinations. On a separate sheet, import the current Google Cloud prices and use these prices for the calculations within formulas.

Answer: C

Explanation:

Question: 279

You want to enable your development team to deploy new features to an existing Cloud Run service in production. To minimize the risk associated with a new revision, you want to reduce the number of customers who might be affected by an outage without introducing any development or operational costs to your customers. You want to follow Google-recommended practices for managing revisions to a service. What should you do9

A. Deploy your application to a second Cloud Run service, and ask your customers to use the second Cloud Run service.
B. Ask your customers to retry access to your service with exponential backoff to mitigate any potential problems after the

new revision is deployed.

C. Gradually roll out the new revision and split customer traffic between the revisions to allow rollback in case a problem occurs.

D. Send all customer traffic to the new revision, and roll back to a previous revision if you witness any problems in production.

Answer: C

Explanation:

Question: 280

You have two Google Cloud projects: project-a with VPC vpc-a (10.0.0.0/16) and project-b with VPC vpc-b (10.8.0.0/16). Your frontend application resides in vpc-a and the backend API services ate deployed in vpc-b. You need to efficiently and cost-effectively enable communication between these Google Cloud projects. You also want to follow Google-recommended practices. What should you do?

A. Configure a Cloud Router in vpc-a and another Cloud Router in vpc-b.

B. Configure a Cloud Interconnect connection between vpc-a and vpc-b.

C. Create VPC Network Peering between vpc-a and vpc-b.

D. Create an OpenVPN connection between vpc-a and vpc-b.

Answer: C

Explanation:

Question: 281

You are building a backend service for an ecommerce platform that will persist transaction data from mobile and web clients. After the platform is launched, you expect a large volume of global transactions. Your business team wants to run SQL queries to analyze the dat

a. You need to build a highly available and scalable data store for the platform. What should you do?

A. Create a multi-region Cloud Spanner instance with an optimized schema.

B. Create a multi-region Firestore database with aggregation query enabled.

C. Create a multi-region Cloud SQL for PostgreSQL database with optimized indexes.

D. Create a multi-region BigQuery dataset with optimized tables.

Answer: A

Explanation:

Question: 282

Your company requires that Google Cloud products are created with a specific configuration to comply with your company's security policies You need to implement a mechanism that will allow software engineers at your company to deploy and update Google Cloud products in a preconfigured

and approved manner. What should you do?

A. Create Java packages that utilize the Google Cloud Client Libraries for Java to configure Google Cloud products. Store and share the packages in a source code repository.
B. Create bash scripts that utilize the Google Cloud CLI to configure Google Cloud products. Store and share the bash scripts in a source code repository.
C. Create Terraform modules that utilize the Google Cloud Terraform Provider to configure Google Cloud products. Store and share the modules in a source code repository.
D. Use the Google Cloud APIs by using curl to configure Google Cloud products. Store and share the curl commands in a source code repository.

Answer: C

Explanation:

Question: 283

You use Cloud Logging lo capture application logs. You now need to use SOL to analyze the application logs in Cloud Logging, and you want to follow Google-recommended practices. What should you do?

A. Develop SQL queries by using Gemini for Google Cloud.
B. Enable Log Analytics for the log bucket and create a linked dataset in BigQuery.
C. Create a schema for the storage bucket and run SQL queries for the data in the bucket.
D. Export logs to a storage bucket and create an external view in BigQuery.

Answer: B

Explanation:

Question: 284

Your preview application, deployed on a single-zone Google Kubernetes Engine (GKE) cluster in us- centrall, has gained popularity. You are now ready to make the application generally available. You need to deploy the application to production while ensuring high availability and resilience. You also want to follow Google-recommended practices. What should you do?

A. Use the gcloud container clusters create command with the options--enable-multi-networking and--enable- autoscaling to create an autoscaling zonal cluster and deploy the application to it.
B. Use the gcloud container clusters create-auto command to create an autopilot cluster and deploy the application to it.
C. Use the gcloud container clusters update command with the option—region us-centrall to update the cluster and deploy the application to it.
D. Use the gcloud container clusters update command with the option—node-locations us-centrall- a,us-centrall-b to update the cluster and deploy the application to the nodes.

Answer: B

Explanation:

Question: 285

You want to deploy a new containerized application into Google Cloud by using a Kubernetes manifest. You want to have full control over the Kubernetes deployment, and at the same time, you want to minimize configuring infrastructure. What should you do?

A. Deploy the application on GKE Autopilot.
B. Deploy the application on GKE Standard.
C. Deploy the application on Cloud Functions.
D. Deploy the application on Cloud Run.

Answer: A

Explanation:

Question: 286

You need to deploy a third-party software application onto a single Compute Engine VM instance. The application requires the highest speed read and write disk access for the internal database. You need to ensure the instance will recover on failure. What should you do?

A. Create an instance template. Set the disk type to be an SSD Persistent Disk. Launch the instance template as part of a stateful managed instance group.
B. Create an instance template. Set the disk type to be an SSD Persistent Disk. Launch the instance template as part of a stateless managed instance group.
C. Create an instance template. Set the disk type to be Hyperdisk Extreme. Launch the instance template as part of a stateful managed instance group.
D. Create an instance template. Set the disk type to be Hyperdisk Extreme. Launch the instance template as part of a stateless managed instance group.

Answer: A

Explanation:

Question: 287

Your company is running a critical workload on a single Compute Engine VM instance. Your company's disaster recovery policies require you to backup the entire instance's disk data every day. The backups must be retained for 7 days. You must configure a backup solution that complies with your company's security policies and requires minimal setup and configuration. What should you do?

A. Configure the instance to use persistent disk asynchronous replication.
B. Configure daily scheduled persistent disk snapshots with a retention period of 7 days.
C. Configure Cloud Scheduler to trigger a Cloud Function each day that creates a new machine image

and deletes machine images that are older than 7 days.

D. Configure a bash script using gsutil to run daily through a cron job. Copy the disk's files to a Cloud Storage bucket with archive storage class and an object lifecycle rule to delete the objects after 7 days.

Answer: B

Explanation:

Question: 288

You have several hundred microservice applications running in a Google Kubernetes Engine (GKE) cluster. Each microservice is a deployment with resource limits configured for each container in the deployment. You've observed that the resource limits for memory and CPU are not appropriately set for many of the microservices. You want to ensure that each microservice has right sized limits for memory and CPU. What should you do?

A. Modify the cluster's node pool machine type and choose a machine type with more memory and CPU.
B. Configure a Horizontal Pod Autoscaler for each microservice.
C. Configure GKE cluster autoscaling.
D. Configure a Vertical Pod Autoscaler for each microservice.

Answer: D

Explanation:

Question: 289

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 the crm- databases project. You want to follow Google-recommended practices to grant access to the service account in the web-applications project. What should you do?

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

Answer: C

Explanation:

Question: 290

You are deploying an application on Google Cloud that requires a relational database for storage. To satisfy your company's security policies, your application must connect to your database through an encrypted and authenticated connection that requires minimal management and integrates with Identity and Access Management (IAM). What should you do?

A. Deploy a Cloud SQL database with the SSL mode set to encrypted only, configure SSL/TLS client certificates, and configure a database user and password.
B. Deploy a Cloud SOL database and configure IAM database authentication. Access the database through the Cloud SQL Auth

Proxy.

C. Deploy a Cloud SQL database with the SSL mode set to encrypted only, configure SSL/TLS client certificates, and configure IAM database authentication.

D. Deploy a Cloud SQL database and configure a database user and password. Access the database through the Cloud SQL Auth Proxy.

Answer: B

Explanation:
Cloud SQL Auth Proxy: This proxy ensures secure connections to your Cloud SQL database by automatically handling encryption (SSL/TLS) and IAM-based authentication. It simplifies the management of secure connections without needing to manage SSL/TLS certificates manually. IAM Database Authentication: This allows you to use IAM credentials to authenticate to the database, providing a unified and secure authentication mechanism that integrates seamlessly with Google Cloud IAM.

Question: 291

(You are managing the security configuration of your company's Google Cloud organization. The Operations team needs specific permissions on both a Google Kubernetes Engine (GKE) cluster and a Cloud SQL instance. Two predefined Identity and Access Management (IAM) roles exist that contain a subset of the permissions needed by the team. You need to configure the necessary IAM permissions for this team while following Google-recommended practices. What should you do?)

A. Grant the team the two predefined IAM roles.

B. Create a custom IAM role that combines the permissions from the two relevant predefined roles. C. Create a custom IAM role that includes only the required permissions from the predefined roles. D. Grant the team the IAM roles of Kubernetes Engine Admin and Cloud SQL Admin.

Answer: C

Explanation:
Granting more permissions than necessary violates the principle of least privilege, a fundamental security best practice. While option A grants the necessary permissions (as subsets exist in two predefined roles), it might also grant more permissions than the Operations team strictly requires for their tasks on GKE and Cloud SQL. Option D is too broad; 'Admin' roles grant extensive permissions that likely exceed the specific needs.
Google Cloud's best practices strongly recommend adhering to the principle of least privilege.

Creating a custom role allows you to precisely define the set of permissions the Operations team needs for their specific tasks on the GKE cluster and the Cloud SQL instance, without granting any unnecessary permissions. This minimizes the potential blast radius in case of accidental or malicious actions.
Google Cloud Documentation Reference:

IAM best practices: https://cloud.google.com/iam/docs/best-practices - This document explicitly recommends granting the minimum necessary permissions.
Creating and managing custom roles: https://cloud.google.com/iam/docs/creating-managing- custom-roles - This explains how to create roles tailored to specific job functions.
Understanding roles: https://cloud.google.com/iam/docs/understanding-roles - This outlines the concepts of predefined and custom roles and their use cases.

Question: 292

(Your digital media company stores a large number of video files on-premises. Each video file ranges from 100 MB to 100 GB. You are currently storing 150 TB of video data in your on-premises network, with no room for expansion. You need to migrate all infrequently accessed video files older than one year to Cloud Storage to ensure that on-premises storage remains available for new files. You must also minimize costs and control bandwidth usage. What should you do?)

A. Create a Cloud Storage bucket. Establish an Identity and Access Management (IAM) role with write permissions to the bucket. Use the gsutil tool to directly copy files over the network to Cloud Storage.

B. Set up a Cloud Interconnect connection between the on-premises network and Google Cloud. Establish a private endpoint for Filestore access. Transfer the data from the existing Network File System (NFS) to Filestore.

C. Use Transfer Appliance to request an appliance. Load the data locally, and ship the appliance back to Google for ingestion into Cloud Storage.

D. Use Storage Transfer Service to move the data from the selected on-premises file storage systems to a Cloud Storage bucket.

Answer: D

Explanation:
Let's analyze each option:

A. Using gsutil: While gsutil can transfer data to Cloud Storage, for 150 TB of infrequently accessed data, direct transfer over the network might be slow and consume significant bandwidth, potentially impacting other network operations. It also lacks built-in mechanisms for filtering files based on age. B. Using Cloud Interconnect and Filestore: Cloud Interconnect provides a dedicated connection, but Filestore is a fully managed NFS service primarily designed for high-performance file sharing for applications running in Google Cloud. Migrating 150 TB of infrequently accessed data to Filestore would be cost-inefficient compared to Cloud Storage and doesn't directly address the requirement of moving older than one year files.

C. Using Transfer Appliance: Transfer Appliance is suitable for very large datasets (petabytes) or when network connectivity is poor or unreliable. While it addresses bandwidth concerns, it involves a physical appliance and might be an overkill for 150 TB of data, especially if network connectivity is reasonable.

D. Using Storage Transfer Service: Storage Transfer Service is specifically designed for moving large amounts of data between online storage systems, including on-premises file systems and Cloud Storage. It offers features like filtering by file age, scheduling transfers, and bandwidth control, directly addressing all the requirements of the question: migrating infrequently accessed files older than one year to Cloud Storage, minimizing costs (by using appropriate Cloud Storage classes for infrequent access), and controlling bandwidth usage.

Google Cloud Documentation Reference:

Storage Transfer Service Overview: https://cloud.google.com/storage-transfer- service/docs/overview - This page details the capabilities and use cases of Storage Transfer Service, including transferring from on-premises.

Storage Transfer Service for on-premises data: https://cloud.google.com/storage-transfer- service/docs/on-prem-overview - This specifically covers transferring data from on-premises file systems.

Cloud Storage Classes: https://cloud.google.com/storage/docs/storage-classes - Understanding the different storage classes (Standard, Nearline, Coldline, Archive) is crucial for cost optimization of infrequently accessed dat

a. Storage Transfer Service can be configured to move data to a cost-effective class like Nearline or Coldline.

Question: 293

(You are developing an internet of things (IoT) application that captures sensor data from multiple devices that have already been set up. You need to identify the global data storage product your company should use to store this data. You must ensure that the storage solution you choose meets your requirements of sub-millisecond latency. What should you do?)

A. Store the IoT data in Spanner. Use caches to speed up the process and avoid latencies.
B. Store the IoT data in Bigtable.
C. Capture IoT data in BigQuery datasets.
D. Store the IoT data in Cloud Storage. Implement caching by using Cloud CDN.

Answer: B

Explanation:

Let's evaluate each option based on the requirement of sub-millisecond latency for globally stored IoT data:

A. Spanner with Caching: While Spanner offers strong consistency and global scalability, the base latency might not consistently be sub-millisecond for all read/write operations globally. Introducing caching adds complexity and doesn't guarantee sub-millisecond latency for all initial reads or cache misses.

B. Bigtable: Bigtable is a highly scalable NoSQL database service designed for low-latency, high- throughput workloads. It excels at storing and retrieving large volumes of time-series data, which is
typical for IoT sensor data. Its architecture is optimized for single-key lookups and scans, providing consistent sub-millisecond latency, making it a strong candidate for this use case.

C. BigQuery: BigQuery is a fully managed, serverless data warehouse designed for analytical queries on large datasets. While it's excellent for analyzing IoT data in batch, it's not optimized for the low- latency, high-throughput ingestion and retrieval required for real-time IoT applications with submillisecond latency needs.

D. Cloud Storage with Cloud CDN: Cloud Storage is object storage and is not designed for low-latency transactional workloads. Cloud CDN is a content delivery network that caches content closer to users for faster delivery, but it's not suitable for the primary storage of rapidly incoming IoT sensor data requiring sub-millisecond write latency.

Google Cloud Documentation Reference:

Cloud Bigtable Overview: https://cloud.google.com/bigtable/docs/overview - This document highlights Bigtable's suitability for low-latency and high-throughput applications, including IoT. It mentions its ability to handle massive amounts of data with consistent performance.

Spanner Overview: https://cloud.google.com/spanner/docs/overview - While Spanner offers low latency, Bigtable is generally preferred for extremely high-throughput, low-latency use cases like raw sensor data ingestion due to its optimized architecture for such workloads.

BigQuery Overview: https://cloud.google.com/bigquery/docs/introduction - This emphasizes BigQuery's analytical capabilities rather than low-latency operational workloads.

Cloud Storage Overview: https://cloud.google.com/storage/docs/overview - This describes Cloud Storage as object storage, not ideal for sub-millisecond latency reads and writes required for realtime IoT data.

Question: 294

(You are managing an application deployed on Cloud Run. The development team has released a new version of the application. You want to deploy and redirect traffic to this new version of the application. To ensure traffic to the new version of the application is served with no startup time, you want to ensure that there are two idle instances available for incoming traffic before adjusting the traffic flow. You also want to minimize administrative overhead. What should you do?)

A. Ensure the checkbox "Serve this revision immediately" is unchecked when deploying the new revision. Before changing the traffic rules, use a traffic simulation tool to send load to the new revision.

B. Configure service autoscaling and set the minimum number of instances to 2.

C. Configure revision autoscaling for the new revision and set the minimum number of instances to 2. D. Configure revision autoscaling for the existing revision and set the minimum number of instances to 2.

Answer: C

Explanation:

Let's analyze each option to find the one that meets the requirements of no startup time for new traffic, two idle instances, and minimal administrative overhead:

A. Unchecking "Serve this revision immediately" and using a traffic simulation tool: Unchecking

"Serve this revision immediately" does prevent the new revision from receiving traffic immediately. However, manually using a traffic simulation tool adds administrative overhead. It also doesn't guarantee that two idle instances will be ready before traffic is shifted; you would need to monitor and adjust traffic manually based on the simulation.

B. Configuring service autoscaling and setting the minimum number of instances to 2: Service-level autoscaling applies to all revisions of the service. Setting the minimum instances at the service level would ensure at least two instances are running across all active revisions, not specifically for the new revision before traffic shift.

C. Configuring revision autoscaling for the new revision and setting the minimum number of instances to 2: This is the correct approach. By configuring revision autoscaling specifically for the new revision and setting the minimum number of instances to 2, Cloud Run will ensure that at least two instances of the new version are running and ready to serve traffic before you redirect any traffic to it. This eliminates startup latency when you do shift traffic. It also minimizes administrative overhead as Cloud Run manages the instance scaling based on this configuration.

D. Configuring revision autoscaling for the existing revision and setting the minimum number of instances to 2: This would ensure the existing version has at least two idle instances, which doesn't directly address the requirement of having idle instances ready for the new version before traffic redirection.

Google Cloud Documentation Reference:

Cloud Run Autoscaling: https://cloud.google.com/run/docs/configuring/min-instances - This document explains how to configure minimum and maximum instances for Cloud Run services and revisions. It clarifies that you can set minimum instances at the revision level to ensure instances are always ready.

Cloud Run Traffic Management: https://cloud.google.com/run/docs/managing/traffic - This describes how to deploy new revisions and gradually shift traffic between them. Combining minimum instances on the new revision with traffic splitting allows for zero-downtime deployments with prewarmed instances.

Question: 295

(You need to migrate multiple PostgreSQL databases from your on-premises data center to Google Cloud. You want to significantly improve the performance of your databases while minimizing changes to your data schema and application code. You expect to exceed 150 TB of data per geographical region. You want to follow Google-recommended practices and minimize your operational costs. What should you do?)

A. Migrate your data to AlloyDB.

B. Migrate your data to Spanner.

C. Migrate your data to Firebase.

D. Migrate your data to Bigtable.

Answer: A

Explanation:
Let's analyze each option based on the requirements: PostgreSQL compatibility, significant performance improvement, minimal schema/code changes, handling large data volumes, Google- recommended practices, and cost minimization:

A. Migrate your data to AlloyDB: AlloyDB for PostgreSQL is a fully managed, PostgreSQL-compatible database service that offers significant performance improvements over standard PostgreSQL due to its architectural optimizations. It is designed to handle large data volumes and minimizes the need for schema and application code changes as it's wire-compatible with PostgreSQL. This aligns well with the requirements for performance improvement, minimal changes, large data, and being a Google-recommended option for PostgreSQL workloads.
B. Migrate your data to Spanner: Spanner is a globally distributed, horizontally scalable database with strong consistency. While it offers excellent scalability and performance, it's not directly PostgreSQL-compatible. Migrating to Spanner would likely require significant schema and application code changes due to differences in data modeling and SQL dialect.
C. Migrate your data to Firebase: Firebase is a suite of mobile and web development tools, with its primary database offering being Firestore (a NoSQL document database) and Realtime Database. These are not PostgreSQL-compatible and would require substantial changes to the data model and application code.
D. Migrate your data to Bigtable: Bigtable is a highly scalable NoSQL wide-column store. It's not compatible with PostgreSQL and requires a completely different data model and application logic. Therefore, AlloyDB is the most suitable option as it provides PostgreSQL compatibility for minimal migration effort, significant performance improvements, scalability for large data volumes, and is a recommended Google Cloud database service for PostgreSQL workloads.
Google Cloud Documentation Reference:

AlloyDB for PostgreSQL Overview: https://cloud.google.com/alloydb/docs/overview - This document highlights AlloyDB's PostgreSQL compatibility, performance benefits, scalability, and suitability for migrating existing PostgreSQL workloads.
Spanner Overview: https://cloud.google.com/spanner/docs/overview - This emphasizes Spanner's unique features and differences from traditional relational databases like PostgreSQL.
Firebase Documentation: https://firebase.google.com/docs - This outlines the features of Firebase, including Firestore and Realtime Database, highlighting their NoSQL nature and incompatibility with PostgreSQL.
Cloud Bigtable Overview: https://cloud.google.com/bigtable/docs/overview - This describes Bigtable as a NoSQL database, emphasizing its differences from relational databases like PostgreSQL.

Question: 296

(You are deploying an application to Google Kubernetes Engine (GKE). The application needs to make API calls to a private Cloud Storage bucket. You need to configure your application Pods to authenticate to the Cloud Storage API, but your organization policy prevents the usage of service account keys. You want to follow Google-recommended practices. What should you do?)

A. Create the GKE cluster and deploy the application. Request a security exception to create a Google service account key. Set the constraints/iam.serviceAccountKeyExpiryHours organization policy to 8 hours.
B. Create the GKE cluster and deploy the application. Request a security exception to create a Google service account key. Set the constraints/iam.serviceAccountKeyExpiryHours organization policy to 24 hours.
C. Create the GKE cluster with Workload Identity Federation. Configure the default node service account to access the bucket. Deploy the application into the cluster so the application can use the
node service account permissions. Use Identity and Access Management (IAM) to grant the service account access to the bucket.
D. Create the GKE cluster with Workload Identity Federation. Create a Google service account and a Kubernetes ServiceAccount, and configure both service accounts to use Workload Identity Federation. Attach the Kubernetes ServiceAccount to the application

Pods and configure the Google service account to access the bucket with Identity and Access Management (IAM).

Answer: D

Explanation:
The organization policy explicitly prevents the use of service account keys, so options A and B, which involve requesting exceptions to create them, are not in line with the policy and Google's recommended practices for secure authentication.

Question: 297

(Your company is migrating its workloads to Google Cloud due to an expiring data center contract. The on-premises environment and Google Cloud are not connected. You have decided to follow a lift- and-shift approach, and you plan to modernize the workloads in a future project. Several old applications connect to each other through hard-coded internal IP addresses. You want to migrate these workloads quickly without modifying the application code. You also want to maintain all functionality. What should you do?)

A. Create a VPC with non-overlapping CIDR ranges compared to your on-premises network. When migrating individual workloads, assign each workload a new static internal IP address.
B. Migrate your DNS server first. Configure Cloud DNS with a forwarding zone to your migrated DNS server. Then migrate all other workloads with ephemeral internal IP addresses.
C. Migrate all workloads to a single VPC subnet. Configure Cloud NAT for the subnet and manually assign a static IP address to the Cloud NAT gateway.
D. Create a VPC with the same CIDR ranges as your on-premises network. When migrating individual workloads, assign each workload the same static internal IP address.

Answer: D

Explanation:
Comprehensive and Detailed In Depth Explanation:
The key requirement is to migrate applications that rely on hard-coded internal IP addresses without modifying the application code. To achieve this, the migrated VMs in Google Cloud need to retain their original internal IP addresses.

A. Non-overlapping CIDR ranges and new static IPs: This option requires changing the IP addresses of the migrated workloads, which would necessitate modifying the application code to reflect these new addresses. This violates a core requirement.
B. Migrating DNS and using ephemeral IPs: While migrating DNS can be beneficial in the long run, using ephemeral internal IP addresses for the migrated workloads means their IPs could change

upon restart, breaking the hard-coded IP address dependencies.
C. Single subnet with Cloud NAT and static NAT IP: Cloud NAT allows instances without external IP addresses to access the internet, but it doesn't help in preserving the internal IP addresses that the applications use to communicate with each other. The internal IP addresses of the VMs would still be within the VPC subnet range and might conflict if they are the same as the on-premises IPs.
D. Same CIDR ranges and same static IPs: Creating a VPC with the same CIDR ranges as the on-premises network and assigning the same static internal IP addresses to the migrated workloads is the only way to ensure that the applications can continue to communicate using their hard-coded IP addresses without any code changes. This approach effectively extends the on-premises network's IP address space into Google Cloud (though without direct connectivity initially, as stated in the problem). Once the workloads are migrated, future steps can involve establishing connectivity (e.g., using VPN or Interconnect) if needed for hybrid scenarios.

Google Cloud Documentation Reference:

VPC Network Overview: https://cloud.google.com/vpc/docs/vpc - This document explains the fundamentals of VPC networks and their IP addressing. While it doesn't explicitly detail lift-and-shift scenarios with identical IP ranges without connectivity, it lays the groundwork for understanding VPC configuration.
Considerations for planning IP address ranges: https://cloud.google.com/vpc/docs/subnets#ip- ranges - This section discusses IP address planning, and while overlapping ranges are generally discouraged for connected networks, for isolated migration scenarios as described, it's a necessary step to avoid application changes. The problem statement explicitly says the environments are not connected during the initial migration.

Question: 298

(Your company has a rapidly growing social media platform and a user base primarily located in North Americ
a. Due to increasing demand, your current on-premises PostgreSQL database, hosted in your United States headquarters data center, no longer meets your needs. You need to identify a cloud-based database solution that offers automatic scaling, multi-region support for future expansion, and maintains low latency.)

A. Use Bigtable.
B. Use BigQuery.
C. Use Spanner.
D. Use Cloud SQL for PostgreSQL.

Answer: C

Explanation:
Comprehensive and Detailed In Depth Explanation:
Let's evaluate each database option against the requirements: automatic scaling, multi-region support, and low latency for a growing social media platform:

A. Bigtable: Bigtable is a highly scalable NoSQL database designed for large analytical and operational

workloads with low latency. It offers excellent horizontal scalability and can be deployed across multiple regions for high availability and lower latency for a global user base. However, it's a NoSQL database and might require significant changes to your existing PostgreSQL data model and application code.
B. BigQuery: BigQuery is a fully managed, serverless data warehouse optimized for analytical queries on large datasets. It's not designed for low-latency transactional workloads that a social media platform would require for real-time user interactions. While it's globally available, its primary use case is not operational database needs.
C. Spanner: Spanner is a globally distributed, horizontally scalable relational database service with strong consistency. It offers automatic scaling, built-in multi-region and multi-continental configurations for high availability and low latency across a global user base, and supports standard SQL (with some extensions). This makes it a strong candidate for a rapidly growing platform needing scalability, global presence, and low latency. While it's not directly PostgreSQL, it offers a relational model and tools to aid migration.
D. Cloud SQL for PostgreSQL: Cloud SQL offers managed PostgreSQL instances with automatic scaling capabilities. It supports high availability within a region and cross-region read replicas for disaster recovery and read scaling. However, its multi-region capabilities for write operations and automatic scaling across regions are more limited compared to Spanner. For a rapidly growing platform with a primarily North American user base but future global expansion in mind and a need for low latency, Spanner's architecture is better

suited for true multi-region write capabilities and consistent low latency globally.

Considering the requirements for automatic scaling, multi-region support for both reads and writes with low latency for a growing user base, Spanner is the most appropriate choice.

Google Cloud Documentation Reference:

Cloud Spanner Overview: https://cloud.google.com/spanner/docs/overview - This document highlights Spanner's global scalability, strong consistency, and multi-region capabilities.

Cloud Bigtable Overview: https://cloud.google.com/bigtable/docs/overview - While scalable and low-latency, it's a NoSQL database, which might require significant application changes.

BigQuery Overview: https://cloud.google.com/bigquery/docs/introduction - Focuses on analytics, not low-latency transactional workloads.

Cloud SQL for PostgreSQL Overview: https://cloud.google.com/sql/docs/postgres/overview - While it offers scaling and regional HA, its multi-region write capabilities are not as robust as Spanner's.

Question: 299

(Your company uses a multi-cloud strategy that includes Google Cloud. You want to centralize application logs in a third-party software-as-a-service (SaaS) tool from all environments. You need to integrate logs originating from Cloud Logging, and you want to ensure the export occurs with the least amount of delay possible. What should you do?)

A. Use a Cloud Scheduler cron job to trigger a Cloud Function that queries Cloud Logging and sends the logs to the SaaS tool.

B. Create a Cloud Logging sink and configure Pub/Sub as the destination. Configure the SaaS tool to subscribe to the Pub/Sub topic to retrieve the logs.

C. Create a Cloud Logging sink and configure Cloud Storage as the destination. Configure the SaaS tool to read the Cloud Storage bucket to retrieve the logs.

D. Create a Cloud Logging sink and configure BigQuery as the destination. Configure the SaaS tool to query BigQuery to retrieve the logs.

Answer: B

Explanation:

Comprehensive and Detailed In Depth Explanation:

The requirement is to export logs from Cloud Logging to a third-party SaaS tool with the least amount of delay possible. Let's analyze each option:

A. Cloud Scheduler, Cloud Function, and querying Cloud Logging: This approach introduces a delay based on the Cloud Scheduler's cron job frequency. The Cloud Function would periodically query Cloud Logging, which might not capture the logs in real-time. This does not meet the "least amount of delay possible" requirement.

B. Cloud Logging sink to Pub/Sub, SaaS tool subscribing to Pub/Sub: Cloud Logging sinks can be configured to export logs in near real-time as they are ingested into Cloud Logging. Pub/Sub is a messaging service designed for asynchronous and near real-time message delivery. By configuring the sink to send logs to a Pub/Sub topic, and having the SaaS tool subscribe to this topic, logs can be delivered to the SaaS tool with minimal delay. This aligns with the requirement for immediate export.

C. Cloud Logging sink to Cloud Storage, SaaS tool reading Cloud Storage: Exporting logs to Cloud Storage involves a batch-oriented approach. Logs are typically written to files periodically. The SaaS tool would then need to poll or be configured to read these files, introducing a significant delay compared to a streaming approach.

D. Cloud Logging sink to BigQuery, SaaS tool querying BigQuery: Similar to Cloud Storage, exporting to BigQuery is more suitable for analytical purposes. The SaaS tool would need to periodically query BigQuery, which introduces latency and is not the most efficient way to achieve near real-time log delivery.

Therefore, configuring a Cloud Logging sink to Pub/Sub and having the SaaS tool subscribe to the Pub/Sub topic provides the lowest latency for exporting logs.

Google Cloud Documentation Reference:

Cloud Logging Sinks Overview: https://cloud.google.com/logging/docs/export/configure_export_v2 - This document explains how to create and manage Cloud Logging sinks, including the available destinations.

Pub/Sub Overview: https://cloud.google.com/pubsub/docs/overview - This highlights Pub/Sub's capabilities for real-time message delivery and its use cases in streaming data.

Exporting Logs with Cloud Logging: https://cloud.google.com/logging/docs/export - This provides a comprehensive guide to exporting logs from Cloud Logging to various destinations, emphasizing Pub/Sub for streaming.

Question: 300

(You manage a VPC network in Google Cloud with a subnet that is rapidly approaching its private IP address capacity. You expect the number of Compute Engine VM instances in the same region to double within a week. You need to implement a Google-recommended solution that minimizes operational costs and does not require downtime. What should you do?)

A. Create a second VPC with the same subnet IP range, and connect this VPC to the existing VPC by using VPC Network Peering.

B. Delete the existing subnet, and create a new subnet with double the IP range available.

C. Use the Google Cloud CLI tool to expand the primary IP range of your subnet.

D. Permit additional traffic from the expected range of private IP addresses to reach your VMs by configuring firewall rules.

Answer: C

Explanation:

Comprehensive and Detailed In Depth Explanation:

The problem states that a subnet is nearing its IP address capacity, and the requirement is to expand it without downtime and with minimal operational cost, following Google-recommended practices.

A. Creating a second VPC with the same subnet IP range and peering: While VPC Network Peering allows communication between VPCs, having overlapping IP ranges is generally not recommended and can lead to routing complexities if not managed carefully. It also adds operational overhead of managing two VPCs. This is not the most straightforward or cost-effective solution for simply expanding IP capacity within the same logical network.

B. Deleting and recreating the subnet: Deleting a subnet that contains active VM instances will cause downtime for those instances, violating a key requirement.

C. Using the Google Cloud CLI tool to expand the primary IP range of your subnet: Google Cloud allows you to expand the primary IP range of an existing subnet after it's created, as long as there are no conflicting subnets in the VPC. This operation does not require

deleting the subnet or restarting the existing VMs within it, thus avoiding downtime. It's a direct and cost-effective way to increase the available IP address space within the existing subnet. This is a Google-recommended practice for handling subnet capacity issues.

D. Permitting additional traffic with firewall rules: Firewall rules control network traffic based on IP ranges, protocols, and ports. They do not increase the number of available private IP addresses within the subnet. This option does not address the core issue of IP address exhaustion.

Therefore, expanding the primary IP range of the existing subnet using the Google Cloud CLI is the recommended solution that meets all the requirements: addressing IP capacity, minimizing operational costs, and avoiding downtime.

Google Cloud Documentation Reference:

Expanding Subnet IP Ranges: https://cloud.google.com/vpc/docs/expand-subnet - This documentation explicitly describes how to expand the IP range of an existing subnet without downtime. It outlines the prerequisites and steps involved using the gcloud CLI or the Google Cloud Console.

VPC Network Overview: https://cloud.google.com/vpc/docs/vpc - Provides context on subnet IP

ranges and their management.

Question: 301

(Your company was recently impacted by a service disruption that caused multiple Dataflow jobs to get stuck, resulting in significant downtime in downstream applications and revenue loss. You were able to resolve the issue by identifying and fixing an error you found in the code. You need to design a solution with minimal management effort to identify when jobs are stuck in the future to ensure that this issue does not occur again. What should you do?)

A. Set up Error Reporting to identify stack traces that indicate slowdowns in Dataflow jobs. Set up alerts based on these log entries.

B. Use the Personalized Service Health dashboard to identify issues with Dataflow jobs across regions.

C. Update the Dataflow job configurations to send messages to a Pub/Sub topic when there are delays. Configure a backup Dataflow job to process jobs that are delayed. Use Cloud Tasks to trigger an alert when messages are pushed to the Pub/Sub topic.

D. Set up Cloud Monitoring alerts on the data freshness metric for the Dataflow jobs to receive a notification when a certain threshold is reached.

Answer: D

Explanation:

Comprehensive and Detailed In Depth Explanation:

The goal is to proactively identify stuck Dataflow jobs with minimal management effort. Let's analyze each option:

A. Error Reporting for slowdowns: Error Reporting primarily focuses on capturing and aggregating exceptions and errors (stack traces). While a stuck job might eventually throw an error, it might also just become unresponsive without generating explicit errors. Relying solely on Error Reporting might not provide timely detection of stuck jobs. Identifying stack traces that indicate slowdowns can also be complex and require significant manual configuration and analysis.

B. Personalized Service Health dashboard: The Personalized Service Health dashboard provides information about Google Cloud service incidents that might be affecting your resources. While it can alert you to broader Dataflow service outages, it won't

specifically identify individual stuck jobs due to application-level errors or logic within your Dataflow pipeline.

C. Pub/Sub messages for delays, backup job, and Cloud Tasks alerts: This approach involves significant custom implementation and management. You would need to instrument your Dataflow jobs to detect delays, send messages to Pub/Sub, manage a backup job, and configure Cloud Tasks for alerting. This adds considerable operational overhead and complexity.

D. Cloud Monitoring alerts on data freshness metric: Dataflow provides built-in metrics, including "data freshness" (or similar metrics like "system lag" or "processing time"), which indicate how far behind the pipeline is in processing data. If a job gets stuck, the data freshness will deteriorate beyond an acceptable threshold. Cloud Monitoring allows you to easily set up alerts based on these built-in metrics. This requires minimal custom coding and leverages the platform's existing

monitoring capabilities, aligning with the "minimal management effort" requirement.

Therefore, setting up Cloud Monitoring alerts on relevant Dataflow metrics like data freshness is the most efficient and recommended way to detect stuck Dataflow jobs with minimal management overhead.

Google Cloud Documentation Reference:

Monitoring Dataflow Pipelines: https://cloud.google.com/dataflow/docs/guides/monitoring-your- pipeline - This document details the various metrics available for monitoring Dataflow jobs in Cloud Monitoring, including metrics related to processing time, system lag, and data freshness.

Creating Alerts in Cloud Monitoring: https://cloud.google.com/monitoring/alerts/create - Explains how to set up alerts based on metrics collected by Cloud Monitoring.

Dataflow Metrics: https://cloud.google.com/dataflow/docs/reference/monitoring-metrics - Provides a comprehensive list of Dataflow metrics that can be used for monitoring and alerting.

Question: 302

(You have an application running inside a Compute Engine instance. You want to provide the application with secure access to a BigQuery dataset. You must ensure that credentials are only valid for a short period of time, and your application will only have access to the intended BigQuery dataset. You want to follow Google-recommended practices and minimize your operational costs. What should you do?)

A. Attach a custom service account to the instance, and grant the service account the BigQuery Data Viewer IAM role on the project.

B. Attach a new service account to the instance every hour, and grant the service account the BigQuery Data Viewer IAM role on the dataset.

C. Attach a custom service account to the instance, and grant the service account the BigQuery Data Viewer IAM role on the dataset.

D. Attach a new service account to the instance every hour, and grant the service account the BigQuery Data Viewer IAM role on the project.

Answer: C

Explanation:

Comprehensive and Detailed In Depth Explanation:

The core requirements are secure access to a specific BigQuery dataset from a Compute Engine instance, using short-lived credentials, adhering to Google's best practices, and minimizing operational overhead.

A. Project-level IAM role: Granting the BigQuery Data Viewer role at the project level gives the service account broad access to all BigQuery datasets within that project. This violates the principle of least privilege, a fundamental security best practice, as the application should only have access to the designated dataset.

B. Hourly new service account with dataset-level role: While this aims to achieve short-lived

credentials, the operational burden of creating, attaching, and managing IAM policies for a new service account every hour is significant and not a Google-recommended practice for routine access. It introduces unnecessary complexity and potential for errors.

C. Custom service account with dataset-level IAM role: This is the recommended and most efficient approach. You create a dedicated Google Cloud service account specifically for this application. You then grant this service account the necessary IAM role (e.g., BigQuery Data Viewer, or a more specific custom role) directly on the target BigQuery dataset. When the Compute Engine instance runs as this service account, the Google Cloud client libraries automatically handle the acquisition and rotation of short-lived OAuth 2.0 access tokens from the instance's metadata server. This eliminates the need to manage long-lived credentials (like service account keys) and ensures the application only has access to the intended dataset. This adheres to the principle of least privilege and minimizes operational costs.

D. Hourly new service account with project-level role: This option combines the high operational overhead of frequently creating new service accounts with the security risk of granting overly permissive project-level access. It is not a recommended practice.

Therefore, the most secure, cost-effective, and operationally efficient solution is to create a custom service account, attach it to the Compute Engine instance, and grant it the appropriate BigQuery IAM role specifically on the target dataset. The platform handles the short-lived credentials automatically.

Google Cloud Documentation Reference:

Creating and enabling service accounts for instances: https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances - Explains how to create and associate service accounts with Compute Engine VMs.
Granting, changing, and revoking access to resources: https://cloud.google.com/iam/docs/granting- changing-revoking-access - Details how to manage IAM policies and grant roles to service accounts on specific resources (like BigQuery datasets).
BigQuery IAM roles: https://cloud.google.com/bigquery/docs/control-access - Provides information on the various IAM roles available for controlling access to BigQuery resources at different levels (project, dataset, table, etc.).
Best practices for using service accounts: https://cloud.google.com/iam/docs/best-practices-for- using-service-accounts - Emphasizes the importance of the principle of least privilege and avoiding the management of long-lived service account keys when possible (relying on the metadata server for short-lived tokens).

Question: 303

(You are migrating your company's on-premises compute resources to Google Cloud. You need to deploy batch processing jobs that run every night. The jobs require significant CPU and memory for several hours but can tolerate interruptions. You must ensure that the deployment is cost-effective. What should you do?)

A. Containerize the batch processing jobs and deploy them on Compute Engine.
B. Use custom machine types on Compute Engine.
C. Use the M1 machine series on Compute Engine.

D. Use Spot VMs on Compute Engine.

Answer: D

Explanation:

Spot VMs (formerly known as preemptible VMs) are Compute Engine virtual machine instances that are available at a much lower price than standard Compute Engine instances. However, Compute Engine might preempt (stop) these instances if it needs to reclaim those resources for other tasks. This makes Spot VMs ideal for batch processing jobs that are fault-tolerant and can handle interruptions, as they can be restarted when resources become available again. This directly addresses the requirement for a cost-effective solution for interruptible workloads.

Option A: While containerization offers portability and consistency, it doesn't inherently provide cost savings for compute resources. You would still need to choose a cost-effective underlying compute option.
Option B: Custom machine types allow you to tailor the CPU and memory configuration of your VMs, which can optimize costs to some extent by avoiding over-provisioning. However, they don't offer the significant cost reduction that Spot VMs provide.
Option C: The M1 machine series is a specific family of Compute Engine instances optimized for memory-intensive workloads. While potentially suitable for the job's requirements, it doesn't inherently address the cost-effectiveness requirement as directly as Spot VMs, which are priced lower regardless of the machine series.
Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
The concept and use cases for Spot VMs are explicitly covered in the Compute Engine section of the Google Cloud documentation, which is a key area for the Associate Cloud Engineer certification. The cost savings and suitability for fault-tolerant workloads are highlighted as primary benefits.

Question: 304

(Your company is modernizing its applications and refactoring them to containerized microservices. You need to deploy the infrastructure on Google Cloud so that teams can deploy their applications. The applications cannot be exposed publicly. You want to minimize management and operational overhead. What should you do?)

A. Provision a Standard zonal Google Kubernetes Engine (GKE) cluster.
B. Provision a fleet of Compute Engine instances and install Kubernetes.
C. Provision a Google Kubernetes Engine (GKE) Autopilot cluster.
D. Provision a Standard regional Google Kubernetes Engine (GKE) cluster.

Answer: C

Explanation:

GKE Autopilot is a mode of operation in GKE where Google manages the underlying infrastructure,

including nodes, node pools, and their upgrades. This significantly reduces the management and operational overhead for the user, allowing teams to focus solely on deploying and managing their containerized applications. Since the applications are not exposed publicly, the zonal or regional nature of the cluster primarily impacts availability within Google Cloud, and Autopilot is available for both. Autopilot minimizes the operational burden, which is a key requirement.

Option A: A Standard zonal GKE cluster requires you to manage the nodes yourself, including sizing, scaling, and upgrades, increasing operational overhead compared to Autopilot.

Option B: Manually installing and managing Kubernetes on a fleet of Compute Engine instances involves the highest level of management overhead, which contradicts the requirement to minimize it.
Option D: A Standard regional GKE cluster provides higher availability than a zonal cluster by replicating the control plane and nodes across multiple zones within a region. However, it still requires you to manage the underlying nodes, unlike Autopilot.
Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
The different modes of GKE operation, including Standard and Autopilot, and their respective management responsibilities and benefits, are clearly outlined in the Google Kubernetes Engine documentation, a core topic for the Associate Cloud Engineer certification. The emphasis on reduced operational overhead with Autopilot is a key differentiator.

Question: 305

(You are deploying a web application using Compute Engine. You created a managed instance group (MIG) to host the application. You want to follow Google-recommended practices to implement a secure and highly available solution. What should you do?)

A. Use a proxy Network Load Balancer for the MIG and an A record in your DNS private zone with the load balancer's IP address.
B. Use a proxy Network Load Balancer for the MIG and a CNAME record in your DNS public zone with the load balancer's IP address.
C. Use an Application Load Balancer for the MIG and a CNAME record in your DNS private zone with the load balancer's IP address.
D. Use an Application Load Balancer for the MIG and an A record in your DNS public zone with the load balancer's IP address.

Answer: D

Explanation:

For a web application (typically using HTTP/HTTPS), an Application Load Balancer is the recommended choice as it operates at Layer 7, providing features like content-based routing, SSL termination, and improved security. To expose the application publicly, you would need to use a public DNS zone. An A record in a public DNS zone maps a domain name to the public IP address of the Application Load Balancer. Using a CNAME record would also work but is generally recommended for aliasing one domain name to another, not directly to an IP address.

Option A & B: Network Load Balancers operate at Layer 4 (TCP/UDP) and lack the application-level features of an Application Load Balancer. Private DNS zones are for internal name resolution within

your VPC, not for public access.
Option C: While an Application Load Balancer is the correct type, using a private DNS zone wouldn't make the web application publicly accessible.
Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
The best practices for load balancing web applications on Google Cloud, including the use of Application Load Balancers for Layer 7 traffic and the configuration of public DNS records (A or CNAME) for public access, are detailed in the Google Cloud Load Balancing and Cloud DNS documentation, both important for the Associate Cloud Engineer certification.

Question: 306

(You are migrating your on-premises workload to Google Cloud. Your company is implementing its Cloud Billing configuration and

requires access to a granular breakdown of its Google Cloud costs. You need to ensure that the Cloud Billing datasets are available in BigQuery so you can conduct a detailed analysis of costs. What should you do?)

A. Enable the BigQuery API and ensure that the BigQuery User IAM role is selected. Change the BigQuery dataset to select a data location.

B. Create a Cloud Billing account. Enable the BigQuery Data Transfer Service API to export pricing data.

C. Enable Cloud Billing data export to BigQuery when you create a Cloud Billing account.

D. Enable Cloud Billing on the project and link a Cloud Billing account. Then view the billing data table in the BigQuery dataset.

Answer: C

Explanation:

The most direct and recommended way to get a granular breakdown of your Google Cloud costs in BigQuery is to enable Cloud Billing data export to BigQuery when you create or manage your Cloud Billing account. This automatically sets up a daily export of your billing data to a BigQuery dataset you specify.

Option A: Enabling the BigQuery API and managing IAM roles are necessary for interacting with BigQuery, but they don't automatically populate it with Cloud Billing data. Selecting a data location is also important for BigQuery datasets but is a separate step from enabling billing export.

Option B: The BigQuery Data Transfer Service is used for transferring data from various sources into BigQuery, but for Cloud Billing data, the direct export feature is the standard and simpler method. Option D: Enabling Cloud Billing and linking an account makes billing data available in the Cloud Billing console, but it doesn't automatically export it to BigQuery for detailed analysis. You need to explicitly configure the BigQuery export.

Reference to Google Cloud Certified - Associate Cloud Engineer Documents:

The process of setting up Cloud Billing export to BigQuery is clearly documented in the Google Cloud Billing documentation, which is a fundamental area for the Associate Cloud Engineer certification. Understanding how to access and analyze billing data is crucial for cost management.

Question: 307

(Your company's developers use an automation that you recently built to provision Linux VMs in Compute Engine within a Google Cloud project to perform various tasks. You need to manage the Linux account lifecycle and access for these users. You want to follow Google-recommended practices to simplify access management while minimizing operational costs. What should you do?)

A. Enable OS Login for all VMs. Use IAM roles to grant user permissions.

B. Enable OS Login for all VMs. Write custom startup scripts to update user permissions.

C. Require your developers to create public SSH keys. Make the owner of the public key the root user. D. Require your developers to create public SSH keys. Write custom startup scripts to update user permissions.

Answer: A

Explanation:

OS Login is a Google-recommended practice for managing access to Linux VMs in Compute Engine. It centralizes user account management by linking the Linux user accounts on the VMs to Google Cloud identities. You then use IAM roles to grant users the necessary permissions to access the VMs (e.g., roles/compute.osLogin or roles/compute.osAdminLogin). This simplifies management

as you control access through IAM policies rather than managing individual SSH keys on each VM, thus minimizing operational costs.

Option B: While enabling OS Login is a good first step, writing custom startup scripts to manage user permissions adds complexity and operational overhead, contradicting the goal of simplification and minimizing costs.
Option C: Requiring developers to manage their own SSH keys and making the owner root is a significant security risk and not a recommended practice. It also doesn't centralize management. Option D: This approach also involves managing individual SSH keys and custom scripts, which increases operational overhead and doesn't leverage the centralized management benefits of OS Login.
Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
OS Login and its benefits for simplified and secure Linux VM access management are detailed in the Compute Engine documentation, which is a key area for the Associate Cloud Engineer certification. The integration with IAM for permission control is a central aspect of this service.

Question: 308

(You are managing a stateful application deployed on Google Kubernetes Engine (GKE) that can only have one replic
a. You recently discovered that the application becomes unstable at peak times. You have identified that the application needs more CPU than what has been configured in the manifest at these peak times. You want Kubernetes to allocate the application sufficient CPU resources during these peak times, while ensuring cost efficiency during off-peak periods. What should you do?)

A. Enable cluster autoscaling on the GKE cluster.
B. Configure a Vertical Pod Autoscaler on the Deployment.
C. Configure a Horizontal Pod Autoscaler on the Deployment.
D. Enable node auto-provisioning on the GKE cluster.

Answer: B

Explanation:
The Vertical Pod Autoscaler (VPA) in Kubernetes automatically adjusts the CPU and memory requests and limits of the containers within a pod based on historical and real-time resource usage. In this scenario, where a single-replica stateful application needs more CPU during peak times, VPA can dynamically increase the CPU allocated to the pod when needed and potentially decrease it during off-peak periods to optimize resource utilization and cost efficiency.

Option A: Cluster autoscaling adds or removes nodes in your GKE cluster based on the resource requests of your pods. While it can help with overall cluster capacity, it oesn't directly address the need for more CPU for a specific pod.
Option C: Horizontal Pod Autoscaler (HPA) scales the number of pod replicas based on observed CPU utilization or other select metrics. Since the application can only have one replica, HPA is not suitable.
Option D: Node auto-provisioning is similar to cluster autoscaling, automatically creating and deleting node pools based on workload demands. It doesn't directly manage the resources of individual pods.

Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
The functionality and use cases of the Vertical Pod Autoscaler (VPA) are detailed in the Google Kubernetes Engine documentation, specifically within the resource management and autoscaling sections. Understanding how VPA can dynamically adjust pod resources is relevant to the Associate Cloud Engineer certification.

Question: 309

(You host your website on Compute Engine. The number of global users visiting your website is rapidly expanding. You need to minimize latency and support user growth in multiple geographical regions. You also want to follow Google-recommended practices and minimize operational costs. Which two actions should you take?

Choose 2 answers)

A. Deploy all of your VMs in a single Google Cloud region with the largest available CIDR range.
B. Deploy your VMs in multiple Google Cloud regions closest to your users' geographical locations.
C. Use an external Application Load Balancer in Regional mode.
D. Use an external Application Load Balancer in Global mode.
E. Use a Network Load Balancer.

Answer: BD

Explanation:
To minimize latency for a global user base, it's crucial to serve users from regions geographically close to them. Deploying VMs in multiple Google Cloud regions (Option B) achieves this by reducing

the network distance and thus the round-trip time for requests.

To support user growth and provide a single point of entry with global reach, a global external Application Load Balancer (Option D) is the recommended choice for web applications. It distributes traffic to backend instances across multiple regions based on user proximity, capacity, and health. Application Load Balancers also offer features like SSL termination, content-based routing, and security policies, which are important for modern web applications.

* Option A: Deploying in a single region, regardless of the CIDR range, will result in high latency for users far from that region.
* Option C: A regional external Application Load Balancer only distributes traffic within a single region, not across multiple global regions, thus not effectively minimizing latency for all global users.
* Option E: Network Load Balancers operate at Layer 4 and don't offer the application-level routing and features of an Application Load Balancer, which are generally preferred for web applications. While they can be global, Application Load Balancers are better suited for this scenario.

Reference to Google Cloud Certified - Associate Cloud Engineer Documents:
The concepts of multi-region deployments for low latency and the use of global load balancers (specifically Application Load Balancers for web traffic) for global reach and traffic management are core topics in the Compute Engine and Load Balancing sections of the Google Cloud documentation, which are essential for the Associate Cloud Engineer certification. The best practices for global application deployment are emphasized.

Question: 310

You have an application that is currently processing transactions by using a group of managed VM instances. You need to migrate the application so that it is serverless and scalable. You want to implement an asynchronous transaction processing system, while minimizing management overhead. What should you do?

A. Install Kafka on VM instances to acknowledge incoming transactions. Use Cloud Run to process transactions.
B. Install Kafka on VM Instances to acknowledge incoming transactions. Use VM Instances to process transactions.
C. Use Pub/Sub to acknowledge incoming transactions. Use VM instances to process transactions.
D. Use Pub/Sub to acknowledge incoming transactions. Use Cloud Run to process transactions.

Answer: D

Explanation:

The goal is to create a serverless, scalable, and asynchronous transaction processing system with minimal management overhead.

Serverless Requirement:Options involving installing Kafka on VMs (A, B) or using VM instances for processing (B, C) introduce management overhead associated with VMs (patching, scaling configuration, OS management) and Kafka cluster management, violating the serverless and minimal

management criteria.

Asynchronous Requirement:Both Kafka and Pub/Sub can handle asynchronous messaging. However, Pub/Sub is Google Cloud's fully managed, serverless messaging service, inherently minimizing management overhead compared to self-managed Kafka on VMs.

Scalability and Processing:Cloud Run is a fully managed, serverless platform that automatically scales based on traffic, suitable for processing transactions without managing underlying infrastructure. VM instances require manual scaling configuration or managed instance groups, adding overhead.

Combining Pub/Sub for asynchronous message ingestion (fully managed, serverless) and Cloud Run for processing (fully managed, serverless, scalable) directly meets all requirements: serverless, scalable, asynchronous, and minimal management overhead. Option D is the only one that uses fully serverless components for both ingestion and processing.

Reference:

Google Cloud Pub/Sub Overview: "Pub/Sub is an asynchronous and scalable messaging service..." - https://cloud.google.com/pubsub/docs/overview

Google Cloud Run Overview: "Cloud Run is a managed compute platform that lets you run containers directly on top of Google's scalable infrastructure." - https://cloud.google.com/run/docs/overview/what-is-cloud-run

Serverless Patterns (Pub/Sub + Cloud Run): This combination is a standard pattern for event-driven, serverless applications. - https://cloud.google.com/run/docs/triggering/pubsub-push

Question: 311

You are deploying an application to Google Kubernetes Engine (GKE) that needs to call an external third-party API. You need to provide the external API vendor with a list of IP addresses for their firewall to allow traffic from your application. You want to follow Google-recommended practices and avoid any risk of interrupting traffic to the API due to IP address changes. What should you do?

A. Configure your GKE cluster with one node, and set the node to have a static external IP address. Ensure that the GKE cluster autoscaler is off. Send the external IP address of the node to the vendor to be added to the allowlist.

B. Configure your GKE cluster with private nodes. Configure a Cloud NAT instance with static IP addresses. Provide these IP addresses to the vendor to be added to the allowlist.

C. Configure your GKE cluster with public nodes. Write a Cloud Function that pulls the public IP addresses of each node in the cluster. Trigger the function to run every day with Cloud Scheduler. Send the list to the vendor by email every day.

D. Configure your GKE cluster with private nodes. Configure a Cloud NAT instance with dynamic IP addresses. Provide these IP addresses to the vendor to be added to the allowlist.

Answer: B

Explanation:

The requirement is for a stable set of egress IP addresses from a GKE cluster for allowlisting by a third party, following best practices.

Option A is not recommended: Using a single node lacks scalability and high availability. Relying on a single node's static IP creates a single point of failure and doesn't align with GKE's design principles.
Disabling autoscaling hinders elasticity.
Option C is complex and unreliable: Public nodes typically have ephemeral external IPs (unless manually configured per node, which is difficult to manage with autoscaling). Dynamically tracking and emailing IPs daily is operationally burdensome and prone to race conditions where the allowlist might lag behind IP changes.
Option D uses Cloud NAT but with dynamic IPs. Dynamic IPs change over time, making them unsuitable for stable firewall allowlists.
Option B is the Google-recommended practice: Configuring the GKE cluster with private nodes enhances security as nodes don't have direct external IPs. Cloud NAT provides managed network address translation for these private nodes to access the internet. By configuring Cloud NAT with a static allocation of external IP addresses, all egress traffic from the private GKE nodes will appear to originate from this stable, predictable set of IPs. This set can be given to the vendor for allowlisting without worrying about node IP changes due to scaling or maintenance.
This approach decouples the application's egress IP from the individual nodes, providing stability and adhering to the principle of least privilege (private nodes).
Reference:

Cloud NAT Overview: "Cloud NAT lets certain resources without external IP addresses create outbound connections to the internet." - https://cloud.google.com/nat/docs/overview
Cloud NAT IP Addresses: "When you configure a NAT gateway... You can configure the NAT gateway to automatically allocate regional external IP addresses... Alternatively, you can manually assign a fixed number of static external IP addresses to the gateway." - https://cloud.google.com/nat/docs/overview#ip-addresses
GKE and Cloud NAT: "Configure Cloud NAT with GKE... Use Case: You want a GKE pod to deterministically egress traffic from a static set of IP addresses that you control." - https://cloud.google.com/nat/docs/gke-example
Private Clusters: "Private nodes do not have endpoint-accessible external IP addresses." - https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters

Question: 312

You are planning to migrate your on-premises VMs to Google Cloud. You need to set up a landing zone in Google Cloud before migrating the VMs. You must ensure that all VMs in your production environment can communicate with each other through private IP addresses. You need to allow all VMs in your Google Cloud organization to accept connections on specific TCP ports. You want to follow Google-recommended practices, and you need to minimize your operational costs. What should you do?

A. Create individual VPCs per Google Cloud project. Peer all the VPCs together. Apply organization policies on the organization level.
B. Create individual VPCs for each Google Cloud project. Peer all the VPCs together. Apply hierarchical firewall policies on the organization level.
C. Create a host VPC project with each production project as its service project. Apply organization policies on the organization level.
D. Create a host VPC project with each production project as its service project. Apply hierarchical firewall policies on the organization level.

Answer: D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:
The goal is to create a landing zone facilitating private IP communication across production projects and apply organization-wide firewall rules, following best practices and minimizing operational costs.

Network Structure:Individual VPCs with Peering (A, B): While VPC Peering allows private connectivity, managing a full mesh or complex peering topology across many projects becomes operationally complex and can hit peering limits. It's not the recommended pattern for centralized connectivity in a landing zone.
Shared VPC (C, D): This is the Google-recommended practice for scenarios where resources from multiple projects need to communicate privately within a common VPC network. A central host project owns the network, and service projects use it. This simplifies network administration and connectivity.
Firewall Rules:Organization Policies (A, C): These enforce organizational constraints (e.g., disable external IPs, restrict locations) but do not define specific network firewall rules (like allowing TCP ports).
Hierarchical Firewall Policies (B, D): These allow defining firewall rules at the Organization or Folder level, which are inherited by resources in descendant projects/folders. This is the mechanism to apply consistent firewall rules (like allowing specific TCP ports) across all VMs in the organization (or a specific folder) efficiently, without managing rules in each individual VPC or project.
Combining Shared VPC for the network structure (best practice for cross-project private communication and central management) with Hierarchical Firewall Policies (for applying organization-wide firewall rules) meets all requirements efficiently and follows Google recommendations.

Reference:

Shared VPC Overview: "Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network..." - https://cloud.google.com/vpc/docs/shared-vpc
Hierarchical firewall policies: "Hierarchical firewall policies let you create and enforce a consistent firewall policy across your organization... They can be configured to explicitly deny traffic, or allow traffic..." - https://cloud.google.com/firewall/docs/hierarchical-firewall-policies
Google Cloud security foundations guide: Often recommends Shared VPC and centralized firewall management (using Hierarchical Firewalls or traditional firewalls with tags in the host project) as part of a secure landing zone. - (Conceptual reference, specific document may vary)

Question: 313

Your organization has decided to deploy all its compute workloads to Kubernetes on Google Cloud and two other cloud providers. You want to build an infrastructure-as-code solution to automate the provisioning process for all cloud resources. What should you do?

A. Build the solution by using YAML manifests, and provision the resources.
B. Build the solution by using Terraform, and provision the resources.
C. Build the solution by using Python and the cloud SDKs from all providers to provision the resources.
D. Build the solution by using Config Connector, and provision the resources.

Answer: B

Explanation:

The requirement is for an infrastructure-as-code (IaC) solution that can manage Kubernetes resources and other cloud resources across multiple cloud providers (Google Cloud and two others).

Option A (YAML manifests): YAML manifests are primarily used for defining Kubernetes objects, not for provisioning general cloud resources (like VPCs, IAM policies, databases) across different cloud providers.
Option C (Python + SDKs): While possible, writing custom scripts using each provider's SDK requires significant development effort to handle state management, dependencies, and provider differences. It essentially reinvents much of what dedicated IaC tools provide and is not a standard IaC approach. Option D (Config Connector): Config Connector allows managing Google Cloud resources using Kubernetes-style manifests and APIs. It is specific to Google Cloud and cannot manage resources in other cloud providers.
Option B (Terraform): Terraform is an open-source IaC tool explicitly designed for building, changing, and versioning infrastructure safely and efficiently across multiple cloud providers and on-premises data centers. It uses providers for different platforms (GCP, AWS, Azure, Kubernetes, etc.), allowing a unified workflow to manage diverse resources across the required environments (Google Cloud, other clouds, Kubernetes).
Terraform is the standard tool for multi-cloud IaC automation as described in the scenario. Reference:

Terraform on Google Cloud: "Terraform is an open source infrastructure as code (IaC) tool...
Terraform lets you manage Google Cloud resources with declarative configuration files..." - https://cloud.google.com/docs/terraform
Terraform Providers (General): Terraform supports numerous providers for various cloud platforms and services. - https://registry.terraform.io/browse/providers
Config Connector Overview: "Config Connector is a Kubernetes addon that allows you to manage Google Cloud resources through Kubernetes." (Google Cloud specific) - https://cloud.google.com/config-connector/docs/overview

Question: 314

You assist different engineering teams in deploying their infrastructure on Google Cloud. Your company has defined certain practices required for all workloads. You need to provide the engineering teams with a solution that enables teams to deploy their infrastructure independently without having to know all implementation details of the company's required practices. What should you do?

A. Create a service account per team, and grant the service account the Project Editor role. Ask the teams to provision their infrastructure through the Google Cloud CLI (gcloud CLI), while

impersonating their dedicated service account.
B. Provide training for all engineering teams you work with to understand the company's required practices. Allow the engineering teams to provision the infrastructure to best meet their needs.
C. Configure organization policies to enforce your company's required practices. Ask the teams to provision their infrastructure by using the Google Cloud console.
D. Write Terraform modules for each component that are compliant with the company's required practices, and ask teams to implement their infrastructure through these modules.

Answer: D

Explanation:

The goal is to enable teams to deploy infrastructure independently while ensuring compliance with company practices, without requiring teams to understand the underlying details of those practices.

Option A provides deployment capability but doesn't enforce practices. The Editor role is overly broad, and using the gcloud CLI directly requires knowledge of how to configure resources compliantly.

Option B requires teams to learn all the practices, contradicting the requirement that they don't need to know the implementation details.

Option C (Organization Policies) is useful for setting constraints (e.g., disallowing public IPs, restricting regions), but it doesn't provide pre-configured, deployable components that embody best practices. Teams still need to figure out how to build compliant resources within the policy constraints.

Option D (Terraform Modules): This approach encapsulates the company's required practices within reusable infrastructure-as-code modules. Engineering teams can then use these modules as building blocks, providing only the necessary input parameters (like application name orsize). The module handles the compliant implementation details internally. This allows teams to deploy independently and ensures compliance without needing deep knowledge of every practice.

Using standardized, compliant modules is a common pattern for enabling self-service infrastructure deployment while maintaining standards and governance.

Reference:

Terraform Modules: "Modules are containers for multiple resources that are used together... Modules allow complex resources to be abstracted away behind a clean interface." - https://developer.hashicorp.com/terraform/language/modules

Google Cloud Architecture Framework - Security, privacy, and compliance: Recommends using IaC and pre-approved templates/modules to enforce security configurations. - https://cloud.google.com/architecture/framework/security-privacy-compliance/define-and-enforce- security-configurations

Organization Policy Service: "The Organization Policy Service gives you centralized and programmatic control over your organization's cloud resources... define constraints..." (Focuses on constraints, not providing deployable components). - https://cloud.google.com/resource- manager/docs/organization-policy/overview

Question: 315

You are deploying an application to Cloud Run. Your application requires the use of an API that runs on Google Kubernetes Engine (GKE). You need to ensure that your Cloud Run service can privately reach the API on GKE, and you want to follow Google-recommended practices. What should you do?

A. Deploy an ingress resource on the GKE cluster to expose the API to the internet. Use Cloud Armor to filter for IP addresses that can connect to the API. On the Cloud Run service, configure the application to fetch its public IP address and update the Cloud Armor policy on startup to allow this IP address to call the API on ports 80 and 443.

B. Create an egress firewall rule on the VPC to allow connections to 0.0.0.0/0 on ports 80 and 443.

C. Create an ingress firewall rule on the VPC to allow connections from 0.0.0.0/0 on ports 80 and 443.

D. Deploy an internal Application Load Balancer to expose the API on GKE to the VPC. Configure Cloud DNS with the IP address of the internal Application Load Balancer. Deploy a Serverless VPC Access connector to allow the Cloud Run service to call the API through the FQDN on Cloud DNS.

Answer: D

Explanation:

The requirement is for private communication between a Cloud Run service and a GKE API, following best practices.
Option A exposes the GKE API to the public internet, which violates the "privately reach" requirement. Relying on dynamic IP allowlisting with Cloud Armor is complex and less secure than private networking.
Options B and C configure overly permissive firewall rules (allowing all egress or ingress) and do not establish the necessary private network path between Cloud Run (which normally runs outside your VPC) and the GKE cluster within your VPC.
Option D describes the standard Google-recommended pattern for this scenario:
Internal Application Load Balancer (ILB): Expose the GKE service (API) using an ILB. This gives the service a private IP address accessible only within the VPC network (or connected networks).
Cloud DNS: Create a private DNS zone and record pointing a fully qualified domain name (FQDN) to the ILB's private IP address. This allows services to reach the API via a stable name instead of an IP. Serverless VPC Access Connector: This connector creates a bridge allowing serverless services like Cloud Run to send traffic into your VPC network.
Cloud Run Configuration: Configure the Cloud Run service to use the VPC Access connector. The application code can then call the GKE API using its private FQDN registered in Cloud DNS. This setup ensures traffic flows entirely over private networks (within the VPC via the ILB and through the VPC Access connector), meeting the private communication requirement securely and reliably.
Reference:
Serverless VPC Access: "Serverless VPC Access lets your serverless environment send requests to your VPC network..." - https://cloud.google.com/vpc/docs/serverless-vpc-access
Internal Application Load Balancer Overview: "Google Cloud internal Application Load Balancers are regional, proxy-based Layer 7 load balancers that enable you to run and scale your services behind an internal IP address..." - https://cloud.google.com/load-balancing/docs/internal
Connecting from Cloud Run to a VPC network: Documentation often outlines patterns using VPC Access Connectors and Internal Load Balancers or Private Service Connect. - https://cloud.google.com/run/docs/configuring/connecting-vpc
GKE Internal Load Balancing: How to expose GKE services internally. - https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing

Question: 316

Your company has many legacy third-party applications that rely on a shared NFS server for file sharing between these workloads. You want to modernize the NFS server by using a Google Cloud managed service. You need to select the solution that requires the least amount of change to the application. What should you do?

A. Configure Firestore. Configure all applications to use Firestore instead of the NFS server.
B. Deploy a Filestore instance. Replace all NFS mounts with a Filestore mount.
C. Create a Cloud Storage bucket. Configure all applications to use Cloud Storage client libraries instead of the NFS server.
D. Create a Compute Engine instance and configure an NFS server on the instance. Point all NFS mounts to the Compute Engine instance.

Answer: B

Explanation:

Question: 317

You are planning to migrate a database and a backend application to a Standard Google Kubernetes Engine (GKE) cluster. You need to prevent data loss and make sure there are enough nodes available for your backend application based on the demands of your workloads. You want to follow Google- recommended practices and minimize the amount of manual work required. What should you do?

A. Run your database as a StatefulSet. Configure cluster autoscaling to handle changes in the demands of your workloads.
B. Run your database as a single Pod. Run the resize command when you notice changes in the demands of your workloads.
C. Run your database as a Deployment. Configure cluster autoscaling to handle changes in the demands of your workloads.
D. Run your database as a DaemonSet. Run the resize command when you notice changes in the demands of your workloads.

Answer: A

Explanation:

Question: 318

You are planning to migrate your containerized workloads to Google Kubernetes Engine (GKE). You need to determine which GKE option to use. Your solution must have high availability, minimal

downtime, and the ability to promptly apply security updates to your nodes. You also want to pay only for the compute resources that your workloads use without managing nodes. You want to follow Google-recommended practices and minimize operational costs. What should you do?

A. Configure a Standard multi-zonal GKE cluster.
B. Configure an Autopilot GKE cluster.
C. Configure a Standard zonal GKE cluster.
D. Configure a Standard regional GKE cluster.

Answer: B

Explanation:

Question: 319

Your company stores data from multiple sources that have different data storage requirements. These data include:

1. Customer data that is structured and read with complex queries
2. Historical log data that is large in volume and accessed infrequently
3. Real-time sensor data with high-velocity writes, which needs to be available for analysis but can tolerate some data loss

You need to design the most cost-effective storage solution that fulfills all data storage requirements. What should you do?

A. Use Spanner for all data.
B. Use Cloud SQL for customer data, Cloud Storage (Coldline) for historical logs, and BigQuery for sensor data.
C. Use Cloud SQL for customer data, Cloud Storage (Archive) for historical logs, and Bigtable for sensor data.

D. Use Firestore for customer data, Cloud Storage (Nearline) for historical logs, and Bigtable for sensor data.

Answer: C

Explanation:

Question: 320

You are the Google Cloud systems administrator for your organization. User A reports that they received an error when attempting to access the Cloud SQL database in their Google Cloud project, while User B can access the database. You need to troubleshoot the issue for User A, while following Google-recommended practices.
What should you do first?

A. Confirm that network firewall rules are not blocking traffic for User A.

B. Review recent configuration changes that may have caused unintended modifications to permissions.
C. Verify that User A has the Identity and Access Management (IAM) Project Owner role assigned.
D. Review the error message that User A received.

Answer: D

Explanation:

Question: 321

Your company runs a variety of applications and workloads on Google Cloud and you are responsible for managing cloud costs. You need to identify a solution that enables you to perform detailed cost analysis You also must be able to visualize the cost data in multiple ways on the same dashboard What should you do?

A. Use the cost breakdown report with the available filters from Cloud Billing to visualize the data
B. Enable the Cloud Billing export to BigQuery. and use Looker Studio to visualize the data
C. Run Queries in Cloud Monitoring Create dashboards to visualize the billing metrics
D. Enable Cloud Monitoring metrics export to BigQuery and use Looker to visualize the data

Answer: B

Explanation:

Question: 322

You are planning to move your company's website and a specific asynchronous background job to Google Cloud Your website contains only static HTML content The background job is started through an HTTP endpoint and generates monthly invoices for your customers. Your website needs to be available in multiple geographic locations and requires autoscaling. You want to have no costs

when your workloads are not In use and follow recommended practices. What should you do?

A. Move your website to Google Kubemetes Engine (GKE). and move your background job to Cloud Functions
B. Move both your website and background job to Compute Engine
CMove both your website and background job to Cloud Run.
D. Move your website to Google Kubemetes Engine (GKE), and move your background job to Compute Engine

Answer: C

Explanation: