Google.Associate Cloud Engineer.vJun-2024.by.Reen162q

Number: Associate Cloud Enginee
Passing Score: 800
Time Limit: 120
File Version: 5.0

Certification: Associate Cloud Engineer
Certification Full Name: Associate Cloud Engineer

Exam A

QUESTION 1
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.

Correct Answer: C
Section:
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 2
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.

Correct Answer: B
Section:
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 3

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.

Correct Answer: B
Section:
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 4
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.

Correct Answer: B
Section:
Explanation:
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

QUESTION 5
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.

Correct Answer: A
Section:
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 6
You are configuring service accounts for an application that spans multiple projects. Virtual machines (VMs) running in the web-applications project need access to BigQuery datasets in crm-databases-proj. You want to follow Google-recommended practices to give access to the service account in the web-applications project. What should you do?

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

Correct Answer: C
Section:
Explanation:
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 7
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 the user.

Correct Answer: C
Section:
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 8
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.

Correct Answer: A
Section:
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 9
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.

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/kubernetes-engine/docs/concepts/deployment

QUESTION 10
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.

Correct Answer: D
Section:
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 11
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.

Correct Answer: A
Section:
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 12
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.

Correct Answer: D
Section:
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 13
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.

Correct Answer: D
Section:
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 14
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.

Correct Answer: A
Section:
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 15
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.

Correct Answer: A
Section:
Explanation:
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 16
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

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/deployment-manager/docs/configuration/create-basic-configuration

QUESTION 17
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.

Correct Answer: C
Section:

QUESTION 18

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.

Correct Answer: C
Section:

QUESTION 19
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>

Correct Answer: B
Section:

QUESTION 20
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.

Correct Answer: B
Section:
Explanation:
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 21
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.

Correct Answer: B
Section:
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 22
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.

Correct Answer: D
Section:
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 23
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.

Correct Answer: A
Section:
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 24
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.

Correct Answer: A
Section:
Explanation:
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 25
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.

Correct Answer: D
Section:
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 26
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.

Correct Answer: A
Section:
Explanation:
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 27
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.

Correct Answer: D
Section:

QUESTION 28
You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?

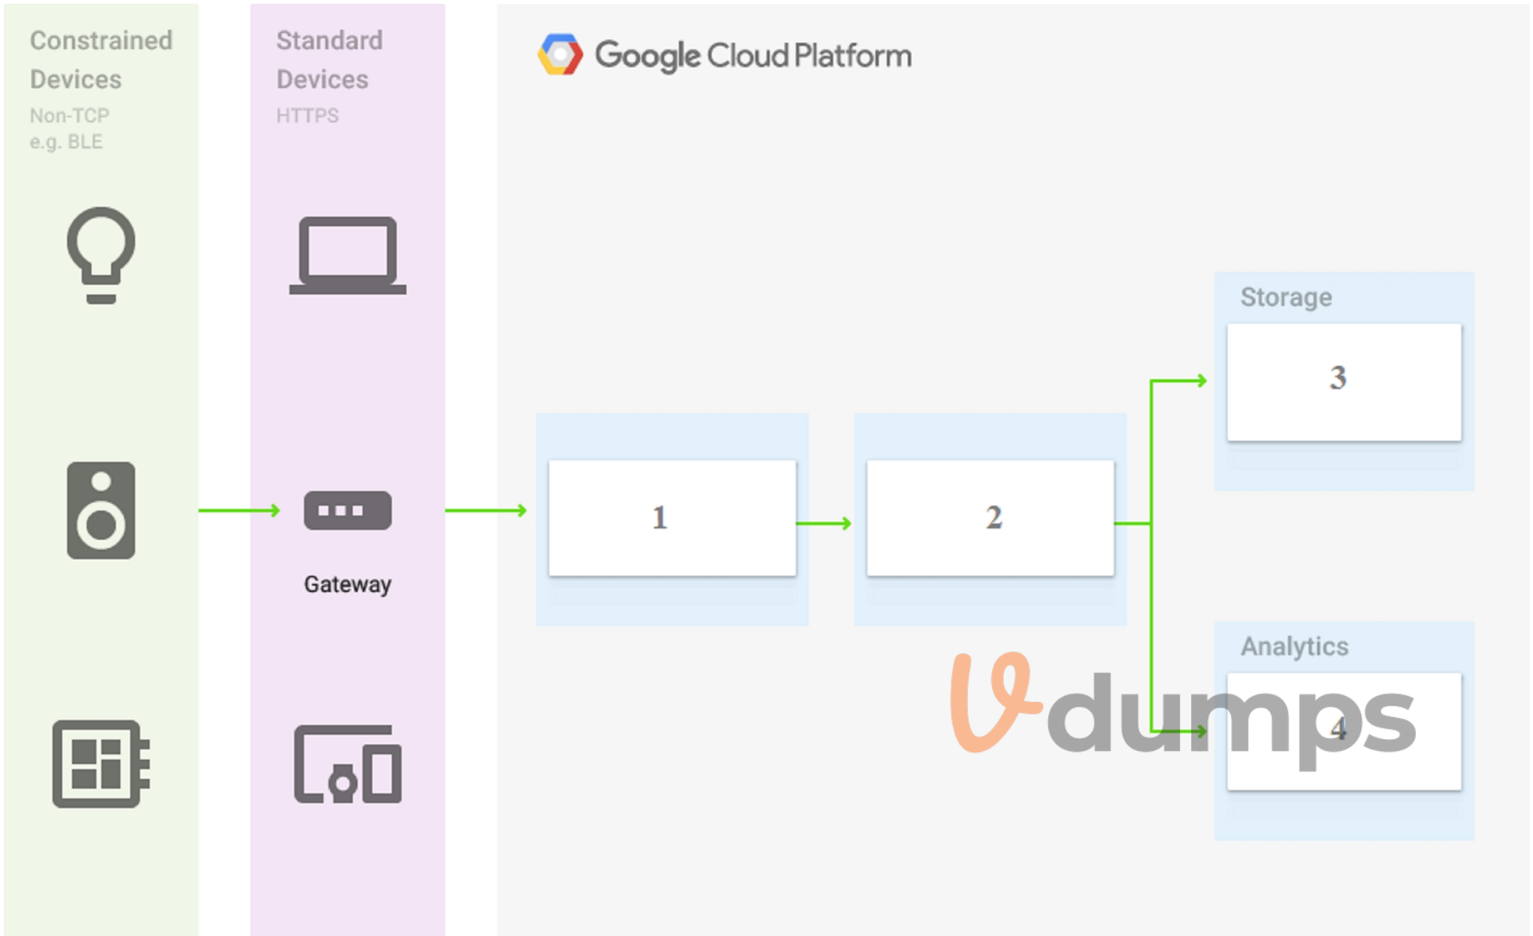

A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery

B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery

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

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

Correct Answer: D
Section:
Explanation:
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 29
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.

Correct Answer: A
Section:
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 30
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.

Correct Answer: A
Section:
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 31
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.

Correct Answer: C
Section:
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 32
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.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/storage/docs/audit-logs
https://cloud.google.com/compute/docs/logging/audit-logging#audited_operations

QUESTION 33
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

Correct Answer: B
Section:
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 34
You have an object in a Cloud Storage bucket that you want to share with an external company. The object contains sensitive data. You want access to the content to be removed after four hours. The external company does not have a Google account to which you can grant specific user-based access privileges. You want to use the most secure method that requires the fewest steps. What should you do?

A. Create a signed URL with a four-hour expiration and share the URL with the company.

B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.

C. Configure the storage bucket as a static website and furnish the object's URL to the 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.

Correct Answer: A
Section:
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 35
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.

Correct Answer: B
Section:
Explanation:
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 36
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.

Correct Answer: A
Section:
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 37
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.

Correct Answer: C
Section:
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 38
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.

Correct Answer: A
Section:
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 39
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.

Correct Answer: A
Section:
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 40
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.

Correct Answer: B
Section:
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 41
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.

Correct Answer: C
Section:
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 42
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.

Correct Answer: D
Section:
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 43
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.

Correct Answer: B
Section:
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 44
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.

Correct Answer: B
Section:
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 45
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.

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

QUESTION 46
You built an application on Google Cloud Platform that uses Cloud Spanner. Your support team needs to monitor the environment but should not have access to table data. You need a streamlined solution to grant the correct permissions to your support team, and you want to follow Google-recommended practices. What should you do?

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

Correct Answer: A
Section:
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 47
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.

Correct Answer: C
Section:
Explanation:
1. In Stackdriver Logging, create a filter to view only Compute Engine logs. 2. Click Create Export. 3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

QUESTION 48
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.

Correct Answer: A
Section:
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 49
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.

Correct Answer: B
Section:
Explanation:
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 50
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.

Correct Answer: A
Section:
Explanation:
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 51
You deployed a new application inside your Google Kubernetes Engine cluster using the YAML file specified below.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp-deployment
spec:
  selector:
    matchLabels:
      app: myapp
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp
    spec:
      containers:
      - name: myapp
        image: myapp:1.1
        ports:
        - containerPort: 80
```

```
apiVersion: v1
kind: Service
metadata:
  name: myapp-service
spec:
  ports:
  - port: 8000
    targetPort: 80
    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 -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
```

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.

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

QUESTION 52
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.

Correct Answer: B
Section:
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 53
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.

Correct Answer: A
Section:
Explanation:
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 54
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 .
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 to verify the project value, and then run gcloud services list --available.

Correct Answer: A
Section:
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 55
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.

Correct Answer: D
Section:
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 56
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.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/compute/docs/disks/local-ssd

QUESTION 57
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.

Correct Answer: A
Section:
Explanation:
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 58
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.

Correct Answer: B
Section:
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 59
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.

Correct Answer: C
Section:
Explanation:
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 60
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.

Correct Answer: D
Section:
Explanation:

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 61
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.

Correct Answer: A
Section:
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 62
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.

Correct Answer: B
Section:
Explanation:
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 63
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.

Correct Answer: D
Section:
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 on-premises 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 64
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.

Correct Answer: D
Section:
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 cost-effective, as it offers the

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

QUESTION 65
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.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/bigquery/docs/quickstarts/quickstart-web-ui?authuser=4

QUESTION 66
You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectl to examine the status of your Pod and observe that one of them is still in Pending status:

```
$ kubectl get pods -l app=myapp
NAME                                READY     STATUS    RESTART   AGE
myapp-deployment-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.

Correct Answer: B
Section:
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 67
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.

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/monitoring/audit-logging

QUESTION 68
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.

Correct Answer: B, E
Section:
Explanation:
https://cloud.google.com/bigquery/docs/custom-quotas https://cloud.google.com/bigquery/pricing#flat_rate_pricing

QUESTION 69
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.

Correct Answer: C
Section:
Explanation:
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. Multi-tenant 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 70
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 INT64 NOT NULL,     // sequential number based on number of registration
     account_creation_date DATE,   // system date
     birthdate DATE,               // customer birthdate
     firstname STRING (255),           // first name
     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 profile_picture field from the table.
B. Add a secondary index on the person_id 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

Correct Answer: C
Section:
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 71
Your finance team wants to view the billing report for your projects. You want to make sure that the finance team does not get additional permissions to the project. What should you do?

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

Correct Answer: C
Section:
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 72
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.

Correct Answer: D
Section:

QUESTION 73
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.

Correct Answer: D
Section:
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 74
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.

Correct Answer: D
Section:
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 75
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.

Correct Answer: D
Section:
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 76
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.

Correct Answer: C
Section:
Explanation:

QUESTION 77
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.

Correct Answer: B
Section:
Explanation:
https://support.google.com/cloudidentity/answer/6262987?hl=en&ref_topic=7558767

QUESTION 78
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.

Correct Answer: D
Section:
Explanation:
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 79
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?
B. Coldline Storage
C. Nearline Storage
D. Regional Storage
E. Multi-Regional Storage

Correct Answer: A
Section:
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

Google Cloud Storage Classes in the Organization

This slide represents the different types of storage classes such as multi-regional, regional, storage nearline, and storage cold line of the Google Cloud.

Storage Class	Characteristics	Use Cases	Price (Per Gb Per Month)*
Multi-Regional Storage	• 99.95% availability • Geo-redundant	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	$0.026 per GB/Month
Regional Storage	• 99.9% availability • Low cost per GB stored • Data storage in a small region	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	$0.02 per GB/Month
Storage Nearline	• 99.0% availability • Very low cost per GB • Data fetching costs • Higher per-task costs • 30-day minimum storage duration	Keeps data that is not accessed is often ideal for data backups	$0.01 per GB/Month
Storage Cold line	• 99.0% availability • Lowest cost per GB • Data fetching costs • Higher per-task costs • 90-day minimum storage duration	Keeps information that is infrequently ideal for disaster recovery or archived data	$0.007 per GB/Month

This slide is 100% editable. Adapt it to your needs and capture your audience's attention.

QUESTION 80
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.

Correct Answer: A
Section:
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 81
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.

Correct Answer: A
Section:
Explanation:
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 82
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.

Correct Answer: A
Section:
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 83
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.

Correct Answer: A
Section:
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 84
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.

Correct Answer: D
Section:
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 85
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.

Correct Answer: B
Section:
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 86
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.

Correct Answer: B
Section:
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 87
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?
B. Enable the Identity Aware Proxy API on the project.
C. Scan the bucker using the Data Loss Prevention API.
D. Allow only a single Service Account access to read the data.
E. Enable Data Access audit logs for the Cloud Storage API.

Correct Answer: D
Section:
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 88
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.

Correct Answer: C
Section:
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

QUESTION 89
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.

Correct Answer: D
Section:
Explanation:

QUESTION 90
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.

Correct Answer: C
Section:
Explanation:

QUESTION 91
You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:
* Restrict access so that suppliers can access only their own data.
* Give suppliers write access to data only for 30 minutes.
* Delete data that is over 45 days old.
You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

A. Build a lifecycle policy to delete Cloud Storage objects after 45 days.

B. Use signed URLs to allow suppliers limited time access to store their objects.

C. Set up an SFTP server for your application, and create a separate user for each supplier.

D. Build a Cloud function that triggers a timer of 45 days to delete objects that have expired.

E. Develop a script that loops through all Cloud Storage buckets and deletes any buckets that are older than 45 days.

Correct Answer: A, B
Section:
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 92
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

Correct Answer: C
Section:
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 93
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

Correct Answer: D
Section:

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

QUESTION 94
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

Correct Answer: C
Section:
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 95
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

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/kuberun/docs/architecture-overview#components_in_the_default_installation

QUESTION 96
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

Correct Answer: B
Section:
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 97
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 Cloud Debugger to find out why the data was not refreshed correctly

Correct Answer: C
Section:
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 98
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

Correct Answer: C
Section:
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 99
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.

Correct Answer: A
Section:
Explanation:
Keyword, Version, traffic splitting, App Engine supports traffic splitting for versions before releasing.

QUESTION 100
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

Correct Answer: A
Section:

QUESTION 101
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

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

QUESTION 102
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

Correct Answer: B
Section:
Explanation:
Keywords: - continually -> Standard - mission-critical analytics -> dual-regional

QUESTION 103
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

Correct Answer: B
Section:

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

QUESTION 104
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.

Correct Answer: C
Section:

QUESTION 105
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.

Correct Answer: D
Section:
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 106
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

Correct Answer: B
Section:

QUESTION 107
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.

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/spanner/docs/cpu-utilization#recommended-max

QUESTION 108
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

Correct Answer: A
Section:

QUESTION 109
You need to manage a third-party application that will run on a Compute Engine instance. Other Compute Engine instances are already running with default configuration. Application installation files are hosted on Cloud Storage. You need to access these files from the new instance without allowing other virtual machines (VMs) to access these files. What should you do?

A. Create the instance with the default Compute Engine service account Grant the service account permissions on Cloud Storage.

B. Create the instance with the default Compute Engine service account Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

C. Create a new service account and assign this service account to the new instance Grant the service account permissions on Cloud Storage.

D. Create a new service account and assign this service account to the new instance Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

Correct Answer: B
Section:
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 110
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

Correct Answer: B
Section:

QUESTION 111
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.

Correct Answer: A
Section:

QUESTION 112
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

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/iap/docs/using-tcp-forwarding

QUESTION 113
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

Correct Answer: B
Section:

QUESTION 114
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

Correct Answer: B

Section:
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 115
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?
B. Create a Cloud Bigtable cluster and use the HBase API
C. Deploy MongoDB Alias from the Google Cloud Marketplace
D. Download a MongoDB installation package and run it on Compute Engine instances
E. Download a MongoDB installation package, and run it on a Managed Instance Group

Correct Answer: B
Section:
Explanation:
https://console.cloud.google.com/marketplace/details/gc-launcher-for-mongodb-atlas/mongodb-atlas

QUESTION 116
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

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/architecture/identity/migrating-consumer-accounts

QUESTION 117
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

Correct Answer: A
Section:
Explanation:

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

QUESTION 118
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 folder that contains the cluster name.
D. Create a file called defaults.json in the ~/.gcloud folder that contains the cluster name.

Correct Answer: A
Section:
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 119
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

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

QUESTION 120
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

Correct Answer: A
Section:

QUESTION 121
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.

Correct Answer: A
Section:

QUESTION 122
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.

Correct Answer: A
Section:
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 123
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.

Correct Answer: D
Section:

QUESTION 124
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.

Correct Answer: B
Section:
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 125
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.

Correct Answer: D
Section:

QUESTION 126
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.

Correct Answer: A
Section:
Explanation:
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 127
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.

Correct Answer: C
Section:
Explanation:
Logged onto console and followed the steps and was able to see all the assigned users and roles.

QUESTION 128
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.

Correct Answer: D
Section:

QUESTION 129
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.

Correct Answer: B
Section:

QUESTION 130
Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises 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

Correct Answer: B
Section:

QUESTION 131
You are building a data lake on Google Cloud for your Internet of Things (IoT) application. The IoT application has millions of sensors that are constantly streaming structured and unstructured data to your backend in the cloud. You want to build a highly available and resilient architecture based on Google-recommended practices. What should you do?

A. Stream data to Pub/Sub, and use Dataflow to send data to Cloud Storage
B. Stream data to Pub/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.

Correct Answer: B
Section:

QUESTION 132

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.

Correct Answer: D
Section:

QUESTION 133
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.

Correct Answer: C
Section:

QUESTION 134
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.

Correct Answer: A
Section:

QUESTION 135
Your company's security vulnc'ability managen->enl policy wonts 3 member of the security tearr ;o have visibility into vulnerabi ities and other OS metadata for a specific Compute Engine instance This Compute Engine instance hosts a critical application in your Gongle Cloud p'oject. You need to irrplement your company's security vulnerability management policy. What should you dc?

A. * Ensuie that the Ops Agent Is Installed on the Compute Engine instance. * Create a custom metric in the Ctoud Monitoring dashboard. * Provide the security team menber with access to this dashboard.
B. * Ensure :hat 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.

Correct Answer: C
Section:

QUESTION 136
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.

Correct Answer: C
Section:

QUESTION 137
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.

Correct Answer: A
Section:

QUESTION 138
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.

Correct Answer: B
Section:
Explanation:
Autopilot is more reliable and stable release gives more time to fix issues in new version of GKE

QUESTION 139
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.

Correct Answer: D
Section:
Explanation:
https://www.cloudskillsboost.google/focuses/1035?parent=catalog#:~:text=custom%20role%20at%20the%20organization%20level

QUESTION 140
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.
```
gcloud container clusters get-credentials dev-cluster
kubectl get nodes
```

B.
```
gcloud container clusters update -generate-password dev-cluster
kubectl get nodes
```

C.
```
kubectl config set-context dev-cluster
kubectl cluster-info
```

D.
```
kubectl config set-credentials dev-cluster
kubectl cluster-info
```

Correct Answer: C
Section:

QUESTION 141
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.

Correct Answer: B
Section:

QUESTION 142
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.

Correct Answer: B
Section:

QUESTION 143
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.

Correct Answer: D
Section:
Explanation:

QUESTION 144
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.

Correct Answer: D
Section:
Explanation:

QUESTION 145
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.

Correct Answer: D
Section:
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 146
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.

Correct Answer: B
Section:
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.
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 147
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.

Correct Answer: C
Section:
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 totrueorfalse, or use a condition to apply the constraint to specific service accounts. By setting this constraint tofalsefor 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 as86400for one day. By setting this constraint to86400for 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.
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 148
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.

Correct Answer: C
Section:
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.
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 149
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.

Correct Answer: A, E
Section:
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.
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 150
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.

Correct Answer: D
Section:
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 151
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.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/monitoring/alerts#alerting-example

QUESTION 152
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.

Correct Answer: D
Section:
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 153
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.

Correct Answer: B
Section:
Explanation:
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 154
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.

Correct Answer: A

Section:
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 155
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.

Correct Answer: C
Section:
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 156
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.

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

QUESTION 157
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.

Correct Answer: C
Section:
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 158
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.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/iam/docs/understanding-roles#billing-roles

QUESTION 159
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.

Correct Answer: A
Section:
Explanation:
Ref:https://cloud.google.com/compute/docs/machine-types#n1_machine_type

QUESTION 160
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.

Correct Answer: A
Section:
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.
Managed instance groups documentation
Choosing a compute option for Google Cloud

QUESTION 161
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 1AM & admin page for all resources, and validate the results.
D. Use the gcloud projects get-iam-policy command to view the current role assignments.

Correct Answer: D
Section:
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 162
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.

Correct Answer: D
Section:
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.
HTTP(S) Load Balancing documentation
Setting up DNS records for HTTP(S) load balancing
Choosing a load balancer