Number: Associate Cloud Engineer
Passing Score: 800
Time Limit: 120
File Version: 21.0

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

Exam A

QUESTION 1
Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

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

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/blog/products/gcp/introducing-coldline-and-a-unified-platform-for-data-storage
Cloud Storage Coldline: a low-latency storage class for long-term archiving Coldline is a new Cloud Storage class designed for long-term archival and disaster recovery. Coldline is perfect for the archival needs of big data or multimedia content, allowing businesses to archive years of data. Coldline provides fast and instant (millisecond) access to data and changes the way that companies think about storing and accessing their cold data.

QUESTION 2
Several employees at your company have been creating projects with Cloud Platform and paying for it with their personal credit cards, which the company reimburses. The company wants to centralize all these projects under a single, new billing account. What should you do?

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

Correct Answer: D
Section:
Explanation:
(https://cloud.google.com/resource-manager/docs/project-migration#change_billing_account)
https://cloud.google.com/billing/docs/concepts
https://cloud.google.com/resource-manager/docs/project-migration

QUESTION 3
You have an application that looks for its licensing server on the IP 10.0.3.21. You need to deploy the licensing server on Compute Engine. You do not want to change the configuration of the application and want the application to be able to reach the licensing server. What should you do?

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

Correct Answer: A
Section:
Explanation:

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

QUESTION 4
You are deploying an application to App Engine. You want the number of instances to scale based on request rate. You need at least 3 unoccupied instances at all times. Which scaling type should you use?

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

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/appengine/docs/standard/go/config/appref
'App Engine calculates the number of instances necessary to serve your current application traffic based on scaling settings such as target_cpu_utilization and target_throughput_utilization. Setting min_idle_instances specifies the number of instances to run in addition to this calculated number. For example, if App Engine calculates that 5 instances are necessary to serve traffic, and min_idle_instances is set to 2, App Engine will run 7 instances (5, calculated based on traffic, plus 2 additional per min_idle_instances).'
Automatic scaling creates dynamic instances based on request rate, response latencies, and other application metrics. However, if you specify the number of minimum idle instances, that specified number of instances run as resident instances while any additional instances are dynamic.
Ref:https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed

QUESTION 5
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 6
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 7
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 8
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 9
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 10
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 11
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 12
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 13
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 14
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 15
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 16
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 17
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 18
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?

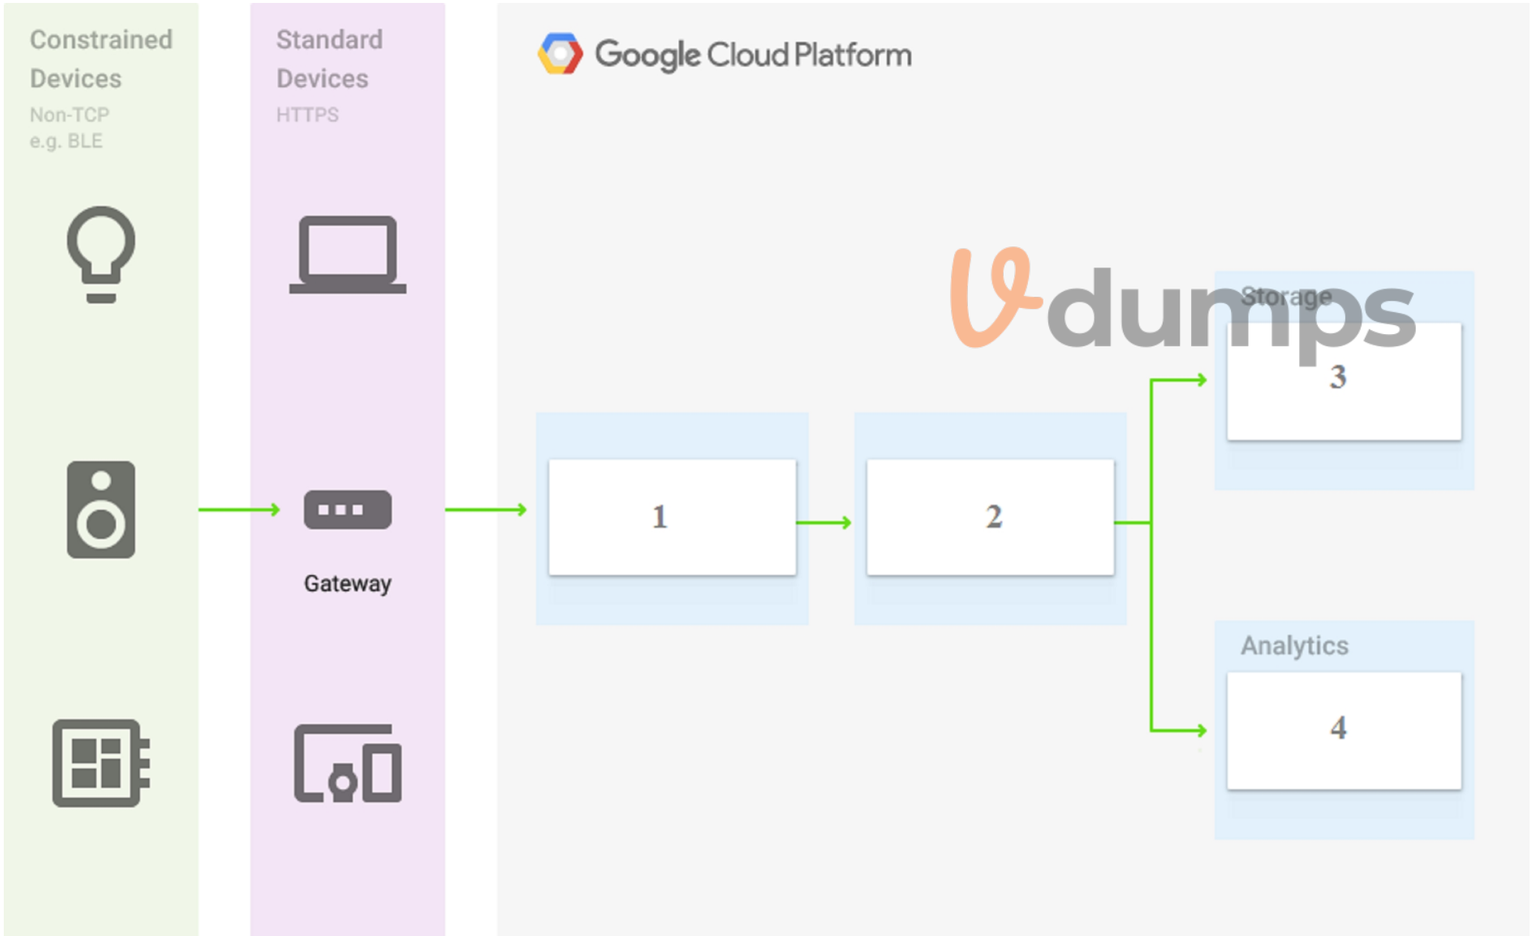

A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery

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 19
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 20
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 21
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 22
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 23
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 24
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 25
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
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 26
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 27
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 28
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 29
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 30
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 31

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 32
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 33
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 34
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 35
You have an instance group that you want to load balance. You want the load balancer to terminate the client SSL session. The instance group is used to serve a public web application over HTTPS. You want to follow Google-recommended practices. What should you do?

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

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

QUESTION 36
You have 32 GB of data in a single file that you need to upload to a Nearline Storage bucket. The WAN connection you are using is rated at 1 Gbps, and you are the only one on the connection. You want to use as much of the rated 1 Gbps as possible to transfer the file rapidly. How should you upload the file?

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

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/storage/docs/parallel-composite-uploads
https://cloud.google.com/storage/docs/uploads-downloads#parallel-composite-uploads

QUESTION 37
You've deployed a microservice called myapp1 to a Google Kubernetes Engine cluster using the YAML file specified below:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp1-deployment
spec:
  selector:
    matchLabels:
      app: myapp1
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp1
    spec:
      containers:
      - name: main-container
        image: gcr.io/my-company-repo/myapp1:1.4
        env:
        - name: DB_PASSWORD
          value: "t0ugh2guess!"
        ports:
        - containerPort: 8080
```

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

A. Store the database password inside the Docker image of the container, not in the YAML file.
B. Store the database password inside a Secret object. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
C. Store the database password inside a ConfigMap object. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
D. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

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

QUESTION 38
You are running an application on multiple virtual machines within a managed instance group and have autoscaling enabled. The autoscaling policy is configured so that additional instances are added to the group if the CPU utilization of instances goes above 80%. VMs are added until the instance group reaches its maximum limit of five VMs or until CPU utilization of instances lowers to 80%. The initial delay for HTTP health checks against the instances is set to 30 seconds. The virtual machine instances take around three minutes to become available for users. You observe that when the instance group autoscales, it adds more instances then necessary to support the levels of end-user traffic. You want to properly maintain instance group sizes when autoscaling. What should you do?

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

Correct Answer: D
Section:
Explanation:

So setting this to 200 ensures that it waits until the instance is up (around 180-second mark) and then starts forwarding traffic to this instance. Even after a cool out period, if the CPU utilization is still high, the autoscaler can again scale up but this scale-up is genuine and is based on the actual load.
Initial Delay Seconds This setting delays autohealing from potentially prematurely recreating the instance if the instance is in the process of starting up. The initial delay timer starts when the currentAction of the instance is VERIFYING. Ref:https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs

QUESTION 39
You need to select and configure compute resources for a set of batch processing jobs. These jobs take around 2 hours to complete and are run nightly. You want to minimize service costs. What should you do?

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

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

QUESTION 40
You recently deployed a new version of an application to App Engine and then discovered a bug in the release. You need to immediately revert to the prior version of the application. What should you do?

A. Run gcloud app restore.
B. On the App Engine page of the GCP Console, select the application that needs to be reverted and click Revert.
C. On the App Engine Versions page of the GCP Console, route 100% of the traffic to the previous version.
D. Deploy the original version as a separate application. Then go to App Engine settings and split traffic between applications so that the original version serves 100% of the requests.

Correct Answer: C
Section:

QUESTION 41
You deployed an App Engine application using gcloud app deploy, but it did not deploy to the intended project. You want to find out why this happened and where the application deployed. What should you do?

A. Check the app.yaml file for your application and check project settings.
B. Check the web-application.xml file for your application and check project settings.
C. Go to Deployment Manager and review settings for deployment of applications.
D. Go to Cloud Shell and run gcloud config list to review the Google Cloud configuration used for deployment.

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

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

QUESTION 42
You want to configure 10 Compute Engine instances for availability when maintenance occurs. Your requirements state that these instances should attempt to automatically restart if they crash. Also, the instances should be highly available including during system maintenance. What should you do?

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

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

QUESTION 43
You host a static website on Cloud Storage. Recently, you began to include links to PDF files on this site. Currently, when users click on the links to these PDF files, their browsers prompt them to save the file onto their local system. Instead, you want the clicked PDF files to be displayed within the browser window directly, without prompting the user to save the file locally. What should you do?

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

Correct Answer: C
Section:
Explanation:
https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_the_correct_mime_type

QUESTION 44
You have a virtual machine that is currently configured with 2 vCPUs and 4 GB of memory. It is running out of memory. You want to upgrade the virtual machine to have 8 GB of memory. What should you do?

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

Correct Answer: D
Section:
Explanation:
In Google compute engine, if predefined machine types don't meet your needs, you can create an instance with custom virtualized hardware settings. Specifically, you can create an instance with a custom number of vCPUs and custom memory, effectively using a custom machine type. Custom machine types are ideal for the following scenarios: 1. Workloads that aren't a good fit for the predefined machine types that are available to you. 2. Workloads that require more processing power or more memory but don't need all of the upgrades that are provided by the next machine type level. In our scenario, we only need a memory upgrade. Moving to a bigger instance would also bump up the CPU which we don't need so we have to use a custom machine type. It is not possible to change memory while the instance is running so you need to first stop the instance, change the

memory and then start it again. See below a screenshot that shows how CPU/Memory can be customized for an instance that has been stopped. Ref:https://cloud.google.com/compute/docs/instances/creating-instance-with-custom-machine-type

QUESTION 45
You have production and test workloads that you want to deploy on Compute Engine. Production VMs need to be in a different subnet than the test VMs. All the VMs must be able to reach each other over internal IP without creating additional routes. You need to set up VPC and the 2 subnets. Which configuration meets these requirements?

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

Correct Answer: A
Section:
Explanation:
When we create subnets in the same VPC with different CIDR ranges, they can communicate automatically within VPC. Resources within a VPC network can communicate with one another by using internal (private) IPv4 addresses, subject to applicable network firewall rules
Ref:https://cloud.google.com/vpc/docs/vpc

QUESTION 46
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 47
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 48
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 49
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 50
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 51
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 52
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 53

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 54
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 55
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 56
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 57
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 58
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 59
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 60
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 61
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 62
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 63
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 64
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 65
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 66
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 67
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 68
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 69
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 70
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 71
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 72
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 73
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 74
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 75
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 76
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 77
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 78
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 79
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 80
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 81
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 82
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 83
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 84
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 85
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 86
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 87
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 88
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 89
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 90
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: A
Section:
Explanation:
This answer is the simplest and most effective way to monitor unexpected firewall changes and instance creation in Google Cloud. Cloud Logging filters allow you to specify the criteria for the log entries that you want to view or export. You can use the Logging query language to write filters based on the LogEntry fields, such as resource.type, severity, or protoPayload.methodName. For example, you can filter for firewall-related events by using the following query:
resource.type=''gce_subnetwork'' logName=''projects/PROJECT_ID/logs/compute.googleapis.com%2Ffirewall''
You can filter for instance-related events by using the following query:
resource.type=''gce_instance'' logName=''projects/PROJECT_ID/logs/compute.googleapis.com%2Factivity_log''
You can create log-based metrics from these filters to measure the rate or count of log entries that match the filter. Log-based metrics can be used to create charts and dashboards in Cloud Monitoring, or to set up alerts based on the metric values. For example, you can create an alert policy that triggers when the log-based metric for firewall changes exceeds a certain threshold in a given time interval. This way, you can get notified of any unexpected or malicious changes to your firewall rules.
Option B is incorrect because it is unnecessarily complex and costly. Installing Kibana on a compute instance requires additional configuration and maintenance. Creating a log sink to forward Cloud Audit Logs to Pub/Sub also incurs additional charges for the Pub/Sub service. Analyzing the logs on Kibana in real time may not be feasible or efficient, as it requires constant monitoring and manual intervention.
Option C is incorrect because Google Cloud firewall rules logging is a different feature from Cloud Audit Logs. Firewall rules logging allows you to audit, verify, and analyze the effects of your firewall rules by creating connection records for each rule that applies to traffic. However, firewall rules logging does not log the insert, update, or delete events for the firewall rules themselves. Those events are logged by Cloud Audit Logs, which record the administrative activities in your Google Cloud project.
Option D is incorrect because it is not a real-time solution. Creating a log sink to forward Cloud Audit Logs to Cloud Storage requires additional storage space and charges. Using BigQuery to periodically analyze log events in the storage bucket also incurs additional costs for the BigQuery service. Moreover, this option does not provide any alerting mechanism to notify you of any unexpected or malicious changes to your firewall rules or instances.

QUESTION 91
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 92
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 93
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 94
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 95
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 96
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 97
Your application stores files on Cloud Storage by using the Standard Storage class. The application only requires access to files created in the last 30 days. You want to automatically save costs on files that are no longer accessed by the application. What should you do?

A. Enable object versioning on the storage bucket and add lifecycle rules to expire non-current versions after 30 days.
B. Create an object lifecycle on the storage bucket to change the storage class to Archive Storage for objects with an age over 30 days.
C. Create a retention policy on the storage bucket of 30 days, and lock the bucket by using a retention policy lock.
D. Create a cron job in Cloud Scheduler to call a Cloud Functions instance every day to delete files older than 30 days.

Correct Answer: B
Section:
Explanation:
Object lifecycle management is a feature of Cloud Storage that allows you to automatically manage the storage class and retention of your objects. You can use object lifecycle management to reduce the cost of storing your data by transitioning objects to lower-cost storage classes, such as Archive Storage, which is designed for long-term storage of data that is rarely accessed. You can also use object lifecycle management to delete objects that are no longer needed after a certain period of time. To use object lifecycle management, you need to create a lifecycle configuration, which is a set of rules that specify the actions to take on objects that match certain conditions. You can apply a lifecycle configuration to a storage bucket, and it will affect all objects in that bucket.
In this scenario, you want to automatically save costs on files that are no longer accessed by the application after 30 days. Therefore, you should create an object lifecycle on the storage bucket to change the storage class to Archive Storage for objects with an age over 30 days. This way, you can keep the files in the bucket, but pay less for storing them. Archive Storage has the lowest cost per GB among all storage classes, but also has the highest retrieval and early deletion fees. Therefore, it is suitable for data that is rarely accessed and has a long retention period.
The other options are not correct because they either do not reduce the cost of storage, or they delete the files that you may still need. Option A is not correct because enabling object versioning on the storage bucket and adding lifecycle rules to expire non-current versions after 30 days will only affect the old versions of the files, not the current ones. Therefore, you will still pay the same amount for storing the current files in the Standard

Storage class. Option C is not correct because creating a retention policy on the storage bucket of 30 days, and locking the bucket by using a retention policy lock will prevent you from deleting or modifying any object in the bucket for 30 days. This will not reduce the cost of storage, and it will also limit your flexibility to manage your data. Option D is not correct because creating a cron job in Cloud Scheduler to call a Cloud Functions instance every day to delete files older than 30 days will permanently remove the files from the bucket. This may not be what you want, as you may still need to access the files for backup, audit, or compliance purposes.
Object Lifecycle Management
Storage Classes
Archive Storage