Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1
Your managed instance group raised an alert stating that new instance creation has failed to create new instances. You need to maintain the number of running instances specified by the template to be able to process expected application traffic. What should you do?

A. Create an instance template that contains valid syntax which will be used by the instance grou
B. Delete any persistent disks with the same name as instance names.
C. Create an instance template that contains valid syntax that will be used by the instance grou
D. Verify that the instance name and persistent disk name values are not the same in the template.
E. Verify that the instance template being used by the instance group contains valid synta
F. Delete any persistent disks with the same name as instance name
G. Set the disks.autoDelete property to true in the instance template.
H. Delete the current instance template and replace it with a new instance templat
I. Verify that the instance name and persistent disk name values are not the same in the templat
J. Set the disks.autoDelete property to true in the instance template.

Answer: A

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

NEW QUESTION 2
You have created a new project in Google Cloud through the gcloud command line interface (CLI) and linked a billing account. You need to create a new Compute Engine instance using the CLI. You need to perform the prerequisite steps. What should you do?

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

Answer: C

NEW QUESTION 3
You have experimented with Google Cloud using your own credit card and expensed the costs to your company. Your company wants to streamline the billing process and charge the costs of your projects to their monthly invoice. What should you do?

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

Answer: D

NEW QUESTION 4
You are working for a hospital that stores Its medical images in an on-premises data room. The hospital wants to use Cloud Storage for archival storage of these images. The hospital wants an automated process to upload any new medical images to Cloud Storage. You need to design and implement a solution. What should you do?

A. Deploy a Dataflow job from the batch template "Datastore lo Cloud Storage" Schedule the batch job on the desired interval
B. In the Cloud Console, go to Cloud Storage Upload the relevant images to the appropriate bucket
C. Create a script that uses the gsutil command line interface to synchronize the on-premises storage with Cloud Storage Schedule the script as a cron job
D. Create a Pub/Sub topic, and enable a Cloud Storage trigger for the Pub/Sub topi
E. Create an application that sends all medical images to the Pub/Sub lope

Answer: C

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

NEW QUESTION 5
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 compan
D. Delete the object from the storage bucket after four hours.
E. Create a new Cloud Storage bucket specifically for the external company to acces
F. Copy the object to that bucke
G. Delete the bucket after four hours have passed.

Answer: A

Explanation:

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

NEW QUESTION 6

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

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

Answer: D

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

NEW QUESTION 7

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

A. Create a Compute Engine snapshot of your base V
B. Create your images from that snapshot.
C. Create a Compute Engine snapshot of your base V
D. Create your instances from that snapshot.
E. Create a custom Compute Engine image from a snapsho
F. Create your images from that image.
G. Create a custom Compute Engine image from a snapsho
H. Create your instances from that image.

Answer: D

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

NEW QUESTION 8

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

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

Answer: B

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

NEW QUESTION 9

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

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

Answer: D

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

NEW QUESTION 10

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

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

Answer: A

Explanation:

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

NEW QUESTION 10
You need to create a custom VPC with a single subnet. The subnet's range must be as large as possible. Which range should you use?

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

Answer: B

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

NEW QUESTION 14
For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

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

Answer: C

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

NEW QUESTION 18
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

Answer: C

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

NEW QUESTION 20
Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1–standard–2 nodes. You need to deploy additional pods requiring n2–highmem–16 nodes without any downtime. What should you do?

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

Answer: B

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

NEW QUESTION 24
You have sensitive data stored in three Cloud Storage buckets and have enabled data access logging. You want to verify activities for a particular user for these buckets, using the fewest possible steps. You need to verify
the addition of metadata labels and which files have been viewed from those buckets. What should you do?

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

Answer: A

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

NEW QUESTION 29
You need to set up permissions for a set of Compute Engine instances to enable them to write data into a particular Cloud Storage bucket. You want to follow Google-recommended practices. What should you do?

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

Answer: C

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

NEW QUESTION 31
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 Engin
B. Use a single-node cluster with a small instance type.
C. Select Google Kubernetes Engin
D. Use a three-node cluster with micro instance types.
E. Select Compute Engin
F. Use preemptible VM instances of the appropriate standard machine type.
G. Select Compute Engin
H. Use VM instance types that support micro bursting.

Answer: C

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

NEW QUESTION 35
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.

Answer: C

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.

NEW QUESTION 37
You have an application that receives SSL-encrypted TCP traffic on port 443. Clients for this application are located all over the world. You want to minimize latency for the clients. Which load balancing option should you use?

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

Answer: C

NEW QUESTION 40
You created a Kubernetes deployment by running kubectl run nginx image=nginx replicas=1. After a few days, you decided you no longer want this deployment. You identified the pod and deleted it by running kubectl delete pod. You noticed the pod got recreated.

≫ $ kubectl get pods

≫ NAME READY STATUS RESTARTS AGE

≫ nginx-84748895c4-nqqmt 1/1 Running 0 9m41s

≫ $ kubectl delete pod nginx-84748895c4-nqqmt

≫ pod nginx-84748895c4-nqqmt deleted

≫ $ kubectl get pods

≫ NAME READY STATUS RESTARTS AGE

≫ nginx-84748895c4-k6bzl 1/1 Running 0 25s
What should you do to delete the deployment and avoid pod getting recreated?

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

Answer: A

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

≫ $ kubectl delete deployment nginx

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

NEW QUESTION 44
You are the project owner of a GCP project and want to delegate control to colleagues to manage buckets and files in Cloud Storage. You want to follow Google-recommended practices. Which IAM roles should you grant your colleagues?

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

Answer: B

Explanation:
Storage Admin (roles/storage.admin) Grants full control of buckets and objects.
When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.
firebase.projects.get resourcemanager.projects.get resourcemanager.projects.list storage.buckets.* storage.objects.*

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

NEW QUESTION 46
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 BigQuer
C. Run a SQL query on this table and drop this table after you complete your request.
D. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.
E. Create a Hadoop cluster and copy the AVRO file to NDFS by compressing i
F. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Answer: C

Explanation:
https://cloud.google.com/bigquery/external-data-sources
An external data source is a data source that you can query directly from BigQuery, even though the data is not stored in BigQuery storage.
BigQuery supports the following external data sources: Amazon S3
Azure Storage Cloud Bigtable Cloud Spanner Cloud SQL Cloud Storage Drive

NEW QUESTION 51
Your company set up a complex organizational structure on Google Could Platform. The structure includes hundreds of folders and projects. Only a few team members should be able to view the hierarchical structure. You need to assign minimum permissions to these team members and you want to follow Google-recommended practices. What should you do?

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

Answer: C

Explanation:
We need to apply the GCP Best practices. roles/browser Browser Read access to browse the hierarchy for a project, including the folder, organization, and IAM policy. This role doesn't include permission to view resources in the project. https://cloud.google.com/iam/docs/understanding-roles

NEW QUESTION 55
You have been asked to set up the billing configuration for a new Google Cloud customer. Your customer wants to group resources that share common IAM policies. What should you do?

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

Answer: B

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

NEW QUESTION 60
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 Registr
D. Create a Deployment YAML file to point to that imag
E. Use kubectl to create the deployment with that file.
F. Create a docker image from the Dockerfile and upload it to Cloud Storag
G. Create a Deployment YAML file to point to that imag
H. Use kubectl to create the deployment with that file.

Answer: C

NEW QUESTION 63
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 projec

D. Use Shared VPC to share the current network with the new project.
E. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Answer: A

Explanation:
https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range
gcloud compute networks subnets expand-ip-range - expand the IP range of a Compute Engine subnetwork gcloud compute networks subnets expand-ip-range
NAME --prefix-length=PREFIX_LENGTH
[--region=REGION] [GCLOUD_WIDE_FLAG …]

NEW QUESTION 64
You need to immediately change the storage class of an existing Google Cloud bucket. You need to reduce service cost for infrequently accessed files stored in that bucket and for all files that will be added to that bucket in the future. What should you do?

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

Answer: B

NEW QUESTION 66
Your organization is a financial company that needs to store audit log files for 3 years. Your organization has hundreds of Google Cloud projects. You need to implement a cost-effective approach for log file retention. What should you do?

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

Answer: B

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

NEW QUESTION 71
You are designing an application that uses WebSockets and HTTP sessions that are not distributed across the web servers. You want to ensure the application runs properly on Google Cloud Platform. What should you do?

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

Answer: A

Explanation:
➢ Google HTTP(S) Load Balancing has native support for the WebSocket protocol when you use HTTP or HTTPS, not HTTP/2, as the protocol to the backend.
Ref: https://cloud.google.com/load-balancing/docs/https#websocket_proxy_support
➢ We dont need to convert WebSocket code to use HTTP streaming or Redesign the application, as
WebSocket support is offered by Google HTTP(S) Load Balancing. Reviewing the encryption requirements is a good idea but it has nothing to do with WebSockets.

NEW QUESTION 74
You created a Kubernetes deployment by running kubectl run nginx image=nginx labels=app=prod. Your Kubernetes cluster is also used by a number of other deployments. How can you find the identifier of the pods for this nginx deployment?

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

Answer: C

Explanation:
This command correctly lists pods that have the label app=prod. When creating the deployment, we used the label app=prod so listing pods that have this label retrieve the pods belonging to nginx deployments. You can list pods by using Kubernetes CLI kubectl get pods.
Ref: https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/
Ref: https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/#list-containe

NEW QUESTION 79
An employee was terminated, but their access to Google Cloud Platform (GCP) was not removed until 2 weeks later. You need to find out this employee accessed any sensitive customer information after their termination. What should you do?

A. View System Event Logs in Stackdrive

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

Answer: C

Explanation:
https://cloud.google.com/logging/docs/audit
Data Access audit logs Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data.
https://cloud.google.com/logging/docs/audit#data-access

NEW QUESTION 81
You are monitoring an application and receive user feedback that a specific error is spiking. You notice that the error is caused by a Service Account having insufficient permissions. You are able to solve the problem but want to be notified if the problem recurs. What should you do?

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

Answer: C

NEW QUESTION 85
Your company has an internal application for managing transactional orders. The application is used exclusively by employees in a single physical location. The application requires strong consistency, fast queries, and ACID guarantees for multi-table transactional updates. The first version of the application is implemented inPostgreSQL, and you want to deploy it to the cloud with minimal code changes. Which database is most appropriate for this application?

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

Answer: B

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

NEW QUESTION 90
You are developing a new application and are looking for a Jenkins installation to build and deploy your source code. You want to automate the installation as quickly and easily as possible. What should you do?

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

Answer: A

Explanation:
 Installing Jenkins
In this section, you use Cloud Marketplace to provision a Jenkins instance. You customize this instance to use the agent image you created in the previous section.
Go to the Cloud Marketplace solution for Jenkins. Click Launch on Compute Engine.
Change the Machine Type field to 4 vCPUs 15 GB Memory, n1-standard-4. Machine type selection for Jenkins deployment.
Click Deploy and wait for your Jenkins instance to finish being provisioned. When it is finished, you will see: Jenkins has been deployed.
https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine#installing_jenkins

NEW QUESTION 94
You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost. What should you do?

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

Answer: B

Explanation:

The key to understand the requirement is : "The objects are written once and accessed frequently for 30 days" Standard Storage
Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.
Archive Storage
Archive Storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days. Archive Storage is the best choice for data that you plan to access less than once a year.
https://cloud.google.com/storage/docs/storage-classes#standard

NEW QUESTION 98
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 Clou
D. Configure both as proxy servers using their public IP addresses.
E. Set up Cloud VPN between the infrastructure on-premises and Google Cloud.

Answer: D

Explanation:
"Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization."
https://cloud.google.com/vpc/docs/vpc-peering while
"Cloud Interconnect provides low latency, high availability connections that enable you to reliably transfer data between your on-premises and Google Cloud Virtual Private Cloud (VPC) networks."
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview
and "HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your 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

NEW QUESTION 100
You need to add a group of new users to Cloud Identity. Some of the users already have existing Google accounts. You want to follow one of Google's recommended practices and avoid conflicting accounts. What should you do?

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

Answer: A

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

NEW QUESTION 103
You manage an App Engine Service that aggregates and visualizes data from BigQuery. The application is deployed with the default App Engine Service account. The data that needs to be visualized resides in a different project managed by another team. You do not have access to this project, but you want your application to be able to read data from the BigQuery dataset. What should you do?

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

Answer: B

Explanation:
The resource that you need to get access is in the other project. roles/bigquery.dataViewer BigQuery Data Viewer
When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.
This role cannot be applied to individual models or routines.
When applied to a dataset, this role provides permissions to: Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables.
When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

NEW QUESTION 104
You have deployed multiple Linux instances on Compute Engine. You plan on adding more instances in the coming weeks. You want to be able to access all of these instances through your SSH client over me Internet without having to configure specific access on the existing and new instances. You do not want the Compute Engine instances to have a public IP. What should you do?

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

Answer: B

Explanation:

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

NEW QUESTION 106
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.databaseUse
B. Add the users to the role.
C. Run gcloud iam roles describe roles/spanner.databaseUse
D. Add the users to a new grou
E. Add the group to the role.
F. Run gcloud iam roles describe roles/spanner.viewer --project my-projec
G. Add the users to the role.
H. Run gcloud iam roles describe roles/spanner.viewer --project my-projec
I. Add the users to a new group.Add the group to the role.

Answer: B

Explanation:
https://cloud.google.com/spanner/docs/iam#spanner.databaseUser
Using the gcloud tool, execute the gcloud iam roles describe roles/spanner.databaseUser command on Cloud Shell. Attach the users to a newly created Google group and add the group to the role.

NEW QUESTION 107
You built an application on your development laptop that uses Google Cloud services. Your application uses Application Default Credentials for authentication and works fine on your development laptop. You want to migrate this application to a Compute Engine virtual machine (VM) and set up authentication using Google-recommended practices and minimal changes. What should you do?

A. Assign appropriate access for Google services to the service account used by the Compute Engine VM.
B. Create a service account with appropriate access for Google services, and configure the application to use this account.
C. Store credentials for service accounts with appropriate access for Google services in a config file, and deploy this config file with your application.
D. Store credentials for your user account with appropriate access for Google services in a config file, and deploy this config file with your application.

Answer: B

Explanation:
In general, Google recommends that each instance that needs to call a Google API should run as a service account with the minimum permissions necessary for that instance to do its job. In practice, this means you should configure service accounts for your instances with the following process: Create a new service account rather than using the Compute Engine default service account. Grant IAM roles to that service account for only the resources that it needs. Configure the instance to run as that service account. Grant the instance the https://www.googleapis.com/auth/cloud-platform scope to allow full access to all Google Cloud APIs, so that the IAM permissions of the instance are completely determined by the IAM roles of the service account. Avoid granting more access than necessary and regularly check your service account permissions to make sure they are up-to-date.
https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances#best_practices

NEW QUESTION 110
You have been asked to migrate a docker application from datacenter to cloud. Your solution architect has suggested uploading docker images to GCR in one project and running an application in a GKE cluster in a separate project. You want to store images in the project img-278322 and run the application in the project prod-278986. You want to tag the image as acme_track_n_trace:v1. You want to follow Google-recommended practices. What should you do?

A. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace
B. Run gcloud builds submit --tag gcr.io/img-278322/acme_track_n_trace:v1
C. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace
D. Run gcloud builds submit --tag gcr.io/prod-278986/acme_track_n_trace:v1

Answer: B

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

NEW QUESTION 113
You are deploying an application to App Engine. You want the number of instances to scale based on request rate. You need at least 3 unoccupied instances at all times. Which scaling type should you use?

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

Answer: D

NEW QUESTION 115
You need to create a Compute Engine instance in a new project that doesn't exist yet. What should you do?

A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.
B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the —project flag to specify a new project.
C. Using the Cloud SDK, create the new instance, and use the —project flag to specify the new project.Answer yes when prompted by Cloud SDK to enable the

Compute Engine API.
D. Enable the Compute Engine API in the Cloud Consol
E. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Answer: A

Explanation:
https://cloud.google.com/sdk/gcloud/reference/projects/create Quickstart: Creating a New Instance Using the Command Line Before you begin
* 1. In the Cloud Console, on the project selector page, select or create a Cloud project.
* 2. Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project.
To use the gcloud command-line tool for this quickstart, you must first install and initialize the Cloud SDK:
* 1. Download and install the Cloud SDK using the instructions given on Installing Google Cloud SDK.
* 2. Initialize the SDK using the instructions given on Initializing Cloud SDK.
To use gcloud in Cloud Shell for this quickstart, first activate Cloud Shell using the instructions given on Starting Cloud Shell.
https://cloud.google.com/ai-platform/deep-learning-vm/docs/quickstart-cli#before-you-begin

NEW QUESTION 119
You have an application that uses Cloud Spanner as a backend database. The application has a very predictable traffic pattern. You want to automatically scale up or down the number of Spanner nodes depending on traffic. What should you do?

A. Create a cron job that runs on a scheduled basis to review stackdriver monitoring metrics, and then resize the Spanner instance accordingly.
B. Create a Stackdriver alerting policy to send an alert to oncall SRE emails when Cloud Spanner CPU exceeds the threshol
C. SREs would scale resources up or down accordingly.
D. Create a Stackdriver alerting policy to send an alert to Google Cloud Support email when Cloud Spanner CPU exceeds your threshol
E. Google support would scale resources up or down accordingly.
F. Create a Stackdriver alerting policy to send an alert to webhook when Cloud Spanner CPU is over or under your threshol
G. Create a Cloud Function that listens to HTTP and resizes Spanner resources accordingly.

Answer: D

Explanation:
As to mexblood1's point, CPU utilization is a recommended proxy for traffic when it comes to Cloud Spanner. See: Alerts for high CPU utilization The following table specifies our recommendations for maximum CPU usage for both single-region and multi-region instances. These numbers are to ensure that your instance has enough compute capacity to continue to serve your traffic in the event of the loss of an entire zone (for single-region instances) or an entire region (for multi-region instances). - https://cloud.google.com/spanner/docs/cpu-utilization

NEW QUESTION 123
You created a cluster.YAML file containing

> resources:
> name: cluster
> type: container.v1.cluster
> properties:
> zone: europe-west1-b
> cluster:
> description: My GCP ACE cluster
> initialNodeCount: 2

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

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

Answer: D

Explanation:
gcloud deployment-manager deployments create creates deployments based on the configuration file. (Infrastructure as code). All the configuration related to the artifacts is in the configuration file. This command correctly creates a cluster based on the provided cluster.yaml configuration file.
Ref: https://cloud.google.com/sdk/gcloud/reference/deployment-manager/deployments/create

NEW QUESTION 126
You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

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

Answer: A

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

NEW QUESTION 128
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 Organizatio
B. Create a new VPC and add all instances.
C. Verify that both projects are in a GCP Organizatio
D. Share the VPC from one project and request that the Compute Engine instances in the other project use this shared VPC.
E. Verify that you are the Project Administrator of both project
F. Create two new VPCs and add all instances.
G. Verify that you are the Project Administrator of both project
H. Create a new VPC and add all instances.

Answer: B

Explanation:
Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network. When you use Shared VPC, you designate a project as a host project and attach one or more other service projects to it. The VPC networks in the host project are called Shared VPC networks. Eligible resources from service projects can use subnets in the Shared VPC network
https://cloud.google.com/vpc/docs/shared-vpc
"For example, an existing instance in a service project cannot be reconfigured to use a Shared VPC network, but a new instance can be created to use available subnets in a Shared VPC network."

NEW QUESTION 131
You are given a project with a single virtual private cloud (VPC) and a single subnetwork in the us-central1 region. There is a Compute Engine instance hosting an application in this subnetwork. You need to deploy a new instance in the same project in the europe-west1 region. This new instance needs access to the application. You want to follow Google-recommended practices. What should you do?

A. 1. Create a subnetwork in the same VPC, in europe-west1.2. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
B. 1. Create a VPC and a subnetwork in europe-west1.2. Expose the application with an internal load balancer.3. Create the new instance in the new subnetwork and use the load balancer's address as the endpoint.
C. 1. Create a subnetwork in the same VPC, in europe-west1.2. Use Cloud VPN to connect the two subnetworks.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.
D. 1. Create a VPC and a subnetwork in europe-west1.2. Peer the 2 VPCs.3. Create the new instance in the new subnetwork and use the first instance's private address as the endpoint.

Answer: C

Explanation:
➤ Given that the new instance wants to access the application on the existing compute engine instance, these applications seem to be related so they should be within the same VPC. It is possible to have them in different VPCs and peer the VPCs but this is a lot of additional work and we can simplify this by choosing the option below (which is the answer)
* 1. Create a subnet in the same VPC, in europe-west1.
* 2. Create the new instance in the new subnet and use the first instance subnets private address as the endpoint. is the right answer.

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

NEW QUESTION 132
Your company wants to standardize the creation and management of multiple Google Cloud resources using Infrastructure as Code. You want to minimize the amount of repetitive code needed to manage the environment What should you do?

A. Create a bash script that contains all requirement steps as gcloud commands
B. Develop templates for the environment using Cloud Deployment Manager
C. Use curl in a terminal to send a REST request to the relevant Google API for each individual resource.
D. Use the Cloud Console interface to provision and manage all related resources

Answer: B

Explanation:
You can use Google Cloud Deployment Manager to create a set of Google Cloud resources and manage them as a unit, called a deployment. For example, if your team's development environment needs two virtual machines (VMs) and a BigQuery database, you can define these resources in a configuration file, and use Deployment Manager to create, change, or delete these resources. You can make the configuration file part of your team's code repository, so that anyone can create the same environment with consistent results. https://cloud.google.com/deployment-manager/docs/quickstart

NEW QUESTION 135
Your company runs one batch process in an on-premises server that takes around 30 hours to complete. The task runs monthly, can be performed offline, and must be restarted if interrupted. You want to migrate this workload to the cloud while minimizing cost. What should you do?

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

Answer: D

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

NEW QUESTION 140
You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command should you use?

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

Answer: B

NEW QUESTION 143
The sales team has a project named Sales Data Digest that has the ID acme-data-digest You need to set up similar Google Cloud resources for the marketing team but their resources must be organized independently of the sales team. What should you do?

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

Answer: C

NEW QUESTION 145
Your company has workloads running on Compute Engine and on-premises. The Google Cloud Virtual Private Cloud (VPC) is connected to your WAN over a Virtual Private Network (VPN). You need to deploy a new Compute Engine instance and ensure that no public Internet traffic can be routed to it. What should you do?

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

Answer: A

Explanation:
VMs cannot communicate over the internet without a public IP address. Private Google Access permits access to Google APIs and services in Google's production infrastructure.
https://cloud.google.com/vpc/docs/private-google-access

NEW QUESTION 149
You need to produce a list of the enabled Google Cloud Platform APIs for a GCP project using the gcloud command line in the Cloud Shell. The project name is my-project. What should you do?

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

Answer: A

Explanation:
`gcloud services list --available` returns not only the enabled services in the project but also services that CAN be enabled.
https://cloud.google.com/sdk/gcloud/reference/services/list#--available
Run the following command to list the enabled APIs and services in your current project: gcloud services list
whereas, Run the following command to list the APIs and services available to you in your current project: gcloud services list –available
https://cloud.google.com/sdk/gcloud/reference/services/list#--available
--available
Return the services available to the project to enable. This list will include any services that the project has already enabled.
To list the services the current project has enabled for consumption, run: gcloud services list --enabled
To list the services the current project can enable for consumption, run: gcloud services list –available

NEW QUESTION 151
Your web application has been running successfully on Cloud Run for Anthos. You want to evaluate an updated version of the application with a specific percentage of your production users (canary deployment). What should you do?

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

H. Add an HTTP Load Balancer in front of both revisions.

Answer: B

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

NEW QUESTION 155
You have created an application that is packaged into a Docker image. You want to deploy the Docker image as a workload on Google Kubernetes Engine. What should you do?

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

Answer: D

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

NEW QUESTION 160
You are building an application that will run in your data center. The application will use Google Cloud Platform (GCP) services like AutoML. You created a service account that has appropriate access to AutoML. You need to enable authentication to the APIs from your on-premises environment. What should you do?

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

Answer: B

NEW QUESTION 162
You have one GCP account running in your default region and zone and another account running in a
non-default region and zone. You want to start a new Compute Engine instance in these two Google Cloud Platform accounts using the command line interface.
What should you do?

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

Answer: A

Explanation:
"Run gcloud configurations list to start the Compute Engine instances". How the heck are you expecting to "start" GCE instances doing "configuration list".
Each gcloud configuration has a 1 to 1 relationship with the region (if a region is defined). Since we have two different regions, we would need to create two separate configurations using gcloud config configurations createRef: https://cloud.google.com/sdk/gcloud/reference/config/configurations/create
Secondly, you can activate each configuration independently by running gcloud config configurations activate [NAME]Ref:
https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate
Finally, while each configuration is active, you can run the gcloud compute instances start [NAME] command to start the instance in the configurations
region.https://cloud.google.com/sdk/gcloud/reference/compute/instances/start

NEW QUESTION 165
You have deployed an application on a single Compute Engine instance. The application writes logs to disk. Users start reporting errors with the application. You want to diagnose the problem. What should you do?

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

Answer: D

NEW QUESTION 167
Your company runs its Linux workloads on Compute Engine instances. Your company will be working with a new operations partner that does not use Google Accounts. You need to grant access to the instances to your operations partner so they can maintain the installed tooling. What should you do?

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

Answer: A

Explanation:
IAP controls access to your App Engine apps and Compute Engine VMs running on Google Cloud. It leverages user identity and the context of a request to determine if a user should be allowed access. IAP is a building block toward BeyondCorp, an enterprise security model that enables employees to work from untrusted networks without using a VPN.
By default, IAP uses Google identities and IAM. By leveraging Identity Platform instead, you can authenticate users with a wide range of external identity providers, such as:
Email/password
OAuth (Google, Facebook, Twitter, GitHub, Microsoft, etc.) SAML
OIDC
Phone number Custom Anonymous
This is useful if your application is already using an external authentication system, and migrating your users to Google accounts is impractical.
https://cloud.google.com/iap/docs/using-tcp-forwarding#grant-permission

NEW QUESTION 169
You are building a new version of an application hosted in an App Engine environment. You want to test the new version with 1% of users before you completely switch your application over to the new version. What should you do?

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

Answer: D

Explanation:
GCP App Engine natively offers traffic splitting functionality between versions. You can use traffic splitting to specify a percentage distribution of traffic across two or more of the versions within a service. Splitting traffic allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.
Ref: https://cloud.google.com/appengine/docs/standard/python/splitting-traffic

NEW QUESTION 171
You want to deploy an application on Cloud Run that processes messages from a Cloud Pub/Sub topic. You want to follow Google-recommended practices. What should you do?

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

Answer: C

Explanation:
https://cloud.google.com/run/docs/tutorials/pubsub#integrating-pubsub
* 1. Create a service account. 2. Give the Cloud Run Invoker role to that service account for your Cloud Run application. 3. Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint.

NEW QUESTION 174
You need to host an application on a Compute Engine instance in a project shared with other teams. You want to prevent the other teams from accidentally causing downtime on that application. Which feature should you use?

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

Answer: D

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

NEW QUESTION 177
Your organization has strict requirements to control access to Google Cloud projects. You need to enable your Site Reliability Engineers (SREs) to approve requests from the Google Cloud support team when an SRE opens a support case. You want to follow Google-recommended practices. What should you do?

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

Answer: D

NEW QUESTION 181
You are the organization and billing administrator for your company. The engineering team has the Project Creator role on the organization. You do not want the engineering team to be able to link projects to the billing account. Only the finance team should be able to link a project to a billing account, but they should not be able to make any other changes to projects. What should you do?

A. Assign the finance team only the Billing Account User role on the billing account.
B. Assign the engineering team only the Billing Account User role on the billing account.
C. Assign the finance team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.
D. Assign the engineering team the Billing Account User role on the billing account and the Project Billing Manager role on the organization.

Answer: C

Explanation:
From this source:
https://cloud.google.com/billing/docs/how-to/custom-roles#permission_association_and_inheritance
"For example, associating a project with a billing account requires the billing.resourceAssociations.create permission on the billing account and also the resourcemanager.projects.createBillingAssignment permission on the project. This is because project permissions are required for actions where project owners control access, while billing account permissions are required for actions where billing account administrators control access. When both should be involved, both permissions are necessary."

NEW QUESTION 185
You are about to deploy a new Enterprise Resource Planning (ERP) system on Google Cloud. The application holds the full database in-memory for fast data access, and you need to configure the most appropriate resources on Google Cloud for this application. What should you do?

A. Provision preemptible Compute Engine instances.
B. Provision Compute Engine instances with GPUs attached.
C. Provision Compute Engine instances with local SSDs attached.
D. Provision Compute Engine instances with M1 machine type.

Answer: D

Explanation:
M1 machine series Medium in-memory databases such as SAP HANA Tasks that require intensive use of memory with higher memory-to-vCPU ratios than the general-purpose high-memory machine types.
In-memory databases and in-memory analytics, business warehousing (BW) workloads, genomics analysis, SQL analysis services. Microsoft SQL Server and similar databases.
https://cloud.google.com/compute/docs/machine-types
https://cloud.google.com/compute/docs/machine-types#:~:text=databases%20such%20as-,SAP%20HANA,-In%
https://www.sap.com/india/products/hana.html#:~:text=is%20SAP%20HANA-,in%2Dmemory,-database%3F

NEW QUESTION 188
You have a Compute Engine instance hosting a production application. You want to receive an email if the instance consumes more than 90% of its CPU resources for more than 15 minutes. You want to use Google services. What should you do?

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

Answer: B

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

NEW QUESTION 193
You have one project called proj-sa where you manage all your service accounts. You want to be able to use a service account from this project to take snapshots of VMs running in another project called proj-vm. What should you do?

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

Answer: C

Explanation:

https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0
You create the service account in proj-sa and take note of the service account email, then you go to proj-vm in IAM > ADD and add the service account's email as new member and give it the Compute Storage Admin role.
https://cloud.google.com/compute/docs/access/iam#compute.storageAdmin

NEW QUESTION 196
You have created a code snippet that should be triggered whenever a new file is uploaded to a Cloud Storage bucket. You want to deploy this code snippet. What should you do?

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

Answer: B

Explanation:
Google Cloud Storage Triggers
Cloud Functions can respond to change notifications emerging from Google Cloud Storage. These notifications can be configured to trigger in response to various events inside a bucket—object creation, deletion, archiving and metadata updates.
Note: Cloud Functions can only be triggered by Cloud Storage buckets in the same Google Cloud Platform project.
Event types
Cloud Storage events used by Cloud Functions are based on Cloud Pub/Sub Notifications for Google Cloud Storage and can be configured in a similar way.
Supported trigger type values are: google.storage.object.finalize google.storage.object.delete google.storage.object.archive google.storage.object.metadataUpdate
Object Finalize
Trigger type value: google.storage.object.finalize
This event is sent when a new object is created (or an existing object is overwritten, and a new generation of that object is created) in the bucket.
https://cloud.google.com/functions/docs/calling/storage#event_types

NEW QUESTION 200
Your auditor wants to view your organization's use of data in Google Cloud. The auditor is most interested in auditing who accessed data in Cloud Storage buckets. You need to help the auditor access the data they need. What should you do?

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

Answer: C

Explanation:
Types of audit logs Cloud Audit Logs provides the following audit logs for each Cloud project, folder, and organization: Admin Activity audit logs Data Access audit logs System Event audit logs Policy Denied audit logs ***Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data. https://cloud.google.com/logging/docs/audit#types
https://cloud.google.com/logging/docs/audit#data-access Cloud Storage: When Cloud Storage usage logs are enabled, Cloud Storage writes usage data to the Cloud Storage bucket, which generates Data Access audit logs for the bucket. The generated Data Access audit log has its caller identity redacted.

NEW QUESTION 201
Your company developed a mobile game that is deployed on Google Cloud. Gamers are connecting to the game with their personal phones over the Internet. The game sends UDP packets to update the servers about the gamers' actions while they are playing in multiplayer mode. Your game backend can scale over multiple virtual machines (VMs), and you want to expose the VMs over a single IP address. What should you do?

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

Answer: D

Explanation:
cell phones are sending UDP packets and the only that can receive that type of traffic is a External Network TCP/UDP https://cloud.google.com/load-balancing/docs/network
https://cloud.google.com/load-balancing/docs/choosing-load-balancer#lb-decision-tree

NEW QUESTION 203
You have a website hosted on App Engine standard environment. You want 1% of your users to see a new test version of the website. You want to minimize complexity. What should you do?

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

Answer: B

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

NEW QUESTION 206
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 objec
C. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
D. Store the database password inside a ConfigMap objec
E. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
F. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

Answer: B

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

NEW QUESTION 211
Your company publishes large files on an Apache web server that runs on a Compute Engine instance. The Apache web server is not the only application running in the project. You want to receive an email when the egress network costs for the server exceed 100 dollars for the current month as measured by Google Cloud Platform (GCP). What should you do?

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

Answer: C

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

NEW QUESTION 212
The core business of your company is to rent out construction equipment at a large scale. All the equipment that is being rented out has been equipped with multiple sensors that send event information every few seconds. These signals can vary from engine status, distance traveled, fuel level, and more. Customers are billed based on the consumption monitored by these sensors. You expect high throughput – up to thousands of events per hour per device – and need to retrieve consistent data based on the time of the event. Storing and retrieving individual signals should be atomic. What should you do?

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

Answer: D

Explanation:
Keyword need to look for

- "High Throughput",
- "Consistent",
- "Property based data insert/fetch like ngine status, distance traveled, fuel level, and more." which can be designed in column,
- "Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds" This will go for pera bytes situation,
- Export data based on the time of the event.
- Atomic
o BigTable will fit all requirement. o DataStore is not fully Atomic
o CloudStorage is not a option where we can export data based on time of event. We need another solution to do that
o FireStore can be used with MobileSDK.

NEW QUESTION 213
You need to assign a Cloud Identity and Access Management (Cloud IAM) role to an external auditor. The auditor needs to have permissions to review your Google Cloud Platform (GCP) Audit Logs and also to review your Data Access logs. What should you do?

A. Assign the auditor the IAM role roles/logging.privateLogViewe
B. Perform the export of logs to Cloud Storage.
C. Assign the auditor the IAM role roles/logging.privateLogViewe
D. Direct the auditor to also review the logs for changes to Cloud IAM policy.
E. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permissio
F. Perform the export of logs to Cloud Storage.
G. Assign the auditor's IAM user to a custom role that has logging.privateLogEntries.list permissio
H. Direct the auditor to also review the logs for changes to Cloud IAM policy.

Answer: B

Explanation:
Google Cloud provides Cloud Audit Logs, which is an integral part of Cloud Logging. It consists of two log streams for each project: Admin Activity and Data Access, which are generated by Google Cloud services to help you answer the question of who did what, where, and when? within your Google Cloud projects.
Ref: https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors

NEW QUESTION 215
You have a Compute Engine instance hosting an application used between 9 AM and 6 PM on weekdays. You want to back up this instance daily for disaster recovery purposes. You want to keep the backups for 30 days. You want the Google-recommended solution with the least management overhead and the least number of
services. What should you do?

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

Answer: B

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

NEW QUESTION 216
You want to select and configure a cost-effective solution for relational data on Google Cloud Platform. You are working with a small set of operational data in one geographic location. You need to support point-in-time recovery. What should you do?

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

Answer: A

NEW QUESTION 217
You need to manage multiple Google Cloud Platform (GCP) projects in the fewest steps possible. You want to configure the Google Cloud SDK command line interface (CLI) so that you can easily manage multiple GCP projects. What should you?

A. * 1. Create a configuration for each project you need to manage.* 2. Activate the appropriate configuration when you work with each of your assigned GCP projects.
B. * 1. Create a configuration for each project you need to manage.* 2. Use gcloud init to update the configuration values when you need to work with a non-default project
C. * 1. Use the default configuration for one project you need to manage.* 2. Activate the appropriate configuration when you work with each of your assigned GCP projects.
D. * 1. Use the default configuration for one project you need to manage.* 2. Use gcloud init to update the configuration values when you need to work with a non-default project.

Answer: A

Explanation:
https://cloud.google.com/sdk/gcloud https://cloud.google.com/sdk/docs/configurations#multiple_configurations

NEW QUESTION 222
You need to create an autoscaling managed instance group for an HTTPS web application. You want to make sure that unhealthy VMs are recreated. What should you do?

A. Create a health check on port 443 and use that when creating the Managed Instance Group.
B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.
C. In the Instance Template, add the label 'health-check'.
D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Answer: A

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

NEW QUESTION 227
Your development team needs a new Jenkins server for their project. You need to deploy the server using the fewest steps possible. What should you do?

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

Answer: D

NEW QUESTION 231
You have a workload running on Compute Engine that is critical to your business. You want to ensure that the data on the boot disk of this workload is backed up regularly. You need to be able to restore a backup as quickly as possible in case of disaster. You also want older backups to be cleaned automatically to save on cost. You want to follow Google-recommended practices. What should you do?

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

Answer: B

Explanation:
Best practices for persistent disk snapshots
You can create persistent disk snapshots at any time, but you can create snapshots more quickly and with greater reliability if you use the following best practices.
Creating frequent snapshots efficiently
Use snapshots to manage your data efficiently.
Create a snapshot of your data on a regular schedule to minimize data loss due to unexpected failure. Improve performance by eliminating excessive snapshot downloads and by creating an image and reusing it. Set your snapshot schedule to off-peak hours to reduce snapshot time.
Snapshot frequency limits
Creating snapshots from persistent disks
You can snapshot your disks at most once every 10 minutes. If you want to issue a burst of requests to snapshot your disks, you can issue at most 6 requests in 60 minutes.
If the limit is exceeded, the operation fails and returns the following error: https://cloud.google.com/compute/docs/disks/snapshot-best-practices

NEW QUESTION 233
You are setting up a Windows VM on Compute Engine and want to make sure you can log in to the VM via RDP. What should you do?

A. After the VM has been created, use your Google Account credentials to log in into the VM.
B. After the VM has been created, use gcloud compute reset-windows-password to retrieve the login credentials for the VM.
C. When creating the VM, add metadata to the instance using 'windows-password' as the key and a password as the value.
D. After the VM has been created, download the JSON private key for the default Compute Engine service accoun
E. Use the credentials in the JSON file to log in to the VM.

Answer: B

Explanation:
You can generate Windows passwords using either the Google Cloud Console or the gcloud command-line tool. This option uses the right syntax to reset the windows password.
gcloud compute reset-windows-password windows-instance
Ref: https://cloud.google.com/compute/docs/instances/windows/creating-passwords-for-windows-instances#gc

NEW QUESTION 236
You need to configure IAM access audit logging in BigQuery for external auditors. You want to follow Google-recommended practices. What should you do?

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

Answer:

A

Explanation:
https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors
Because if you directly add users to the IAM roles, then if any users left the organization then you have to remove the users from multiple places and need to revoke his/her access from multiple places. But, if you put a user into a group then its very easy to manage these type of situations. Now, if any user left then you just need to remove the user from the group and all the access got revoked
The organization creates a Google group for these external auditors and adds the current auditor to the group. This group is monitored and is typically granted access to the dashboard application. During normal access, the auditors' Google group is only granted access to view the historic logs stored in BigQuery. If any anomalies are discovered, the group is granted permission to view the actual Cloud Logging Admin Activity logs via the dashboard's elevated access mode. At the end of each audit period, the group's access is then revoked. Data is redacted using Cloud DLP before being made accessible for viewing via the dashboard application. The table below explains IAM logging roles that an Organization Administrator can grant to the service account used by the dashboard, as well as the resource level at which the role is granted.

NEW QUESTION 238
You need to deploy an application, which is packaged in a container image, in a new project. The application exposes an HTTP endpoint and receives very few requests per day. You want to minimize costs. What should you do?

A. Deploy the container on Cloud Run.
B. Deploy the container on Cloud Run on GKE.
C. Deploy the container on App Engine Flexible.
D. Deploy the container on Google Kubernetes Engine, with cluster autoscaling and horizontal pod autoscaling enabled.

Answer: A

Explanation:
Cloud Run takes any container images and pairs great with the container ecosystem: Cloud Build, Artifact Registry, Docker. ... No infrastructure to manage: once deployed, Cloud Run manages your services so you can sleep well. Fast autoscaling. Cloud Run automatically scales up or down from zero to N depending on traffic.
https://cloud.google.com/run

NEW QUESTION 241
Your company has developed a new application that consists of multiple microservices. You want to deploy the application to Google Kubernetes Engine (GKE), and you want to ensure that the cluster can scale as more applications are deployed in the future. You want to avoid manual intervention when each new application is deployed. What should you do?

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

Answer: C

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

NEW QUESTION 246
Your company uses Cloud Storage to store application backup files for disaster recovery purposes. You want to follow Google's recommended practices. Which storage option should you use?

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

Answer: D

NEW QUESTION 251
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-purposemachine type nodes for the other microservices
C. Use the node pool with general-purpose machine type nodes for lite mage rendering microservice Create a nodepool with compute-optimized machine type nodes for the other microservices
D. Configure the required amount of CPU and memory in the resource requests specification of the image rendering microservice deployment Keep the resource requests for the other microservices at the default

Answer: B

NEW QUESTION 255
You have a batch workload that runs every night and uses a large number of virtual machines (VMs). It is fault- tolerant and can tolerate some of the VMs being terminated. The current cost of VMs is too high. What should you do?

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

Answer: A

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

NEW QUESTION 260
Your finance team wants to view the billing report for your projects. You want to make sure that the finance team does not get additional permissions to the project. What should you do?

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

Answer: C

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

NEW QUESTION 264
You are working with a Cloud SQL MySQL database at your company. You need to retain a month-end copy of the database for three years for audit purposes. What should you do?

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

Answer: C

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

NEW QUESTION 266
Your organization needs to grant users access to query datasets in BigQuery but prevent them from accidentally deleting the datasets. You want a solution that follows Google-recommended practices. What should you do?

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

Answer: D

Explanation:
https://cloud.google.com/bigquery/docs/access-control#custom_roles
Custom roles enable you to enforce the principle of least privilege, ensuring that the user and service accounts in your organization have only the permissions essential to performing their intended functions.

NEW QUESTION 269
You are using Data Studio to visualize a table from your data warehouse that is built on top of BigQuery. Data is appended to the data warehouse during the day. At night, the daily summary is recalculated by overwriting the table. You just noticed that the charts in Data Studio are broken, and you want to analyze the problem. What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 274
You have designed a solution on Google Cloud Platform (GCP) that uses multiple GCP products. Your company has asked you to estimate the costs of the solution. You need to provide estimates for the monthly total cost. What should you do?

A. For each GCP product in the solution, review the pricing details on the products pricing pag
B. Use the pricing calculator to total the monthly costs for each GCP product.
C. For each GCP product in the solution, review the pricing details on the products pricing pag
D. Create a Google Sheet that summarizes the expected monthly costs for each product.
E. Provision the solution on GC
F. Leave the solution provisioned for 1 wee
G. Navigate to the Billing Report page in the Google Cloud Platform Consol
H. Multiply the 1 week cost to determine the monthly costs.
I. Provision the solution on GC
J. Leave the solution provisioned for 1 wee
K. Use Stackdriver to determine the provisioned and used resource amount
L. Multiply the 1 week cost to determine the monthly costs.

Answer: A

Explanation:
You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You dont incur any charges for doing so.
Ref: https://cloud.google.com/products/calculator

NEW QUESTION 276
You are managing a Data Warehouse on BigQuery. An external auditor will review your company's processes, and multiple external consultants will need view access to the data. You need to provide them with view access while following Google-recommended practices. What should you do?

A. Grant each individual external consultant the role of BigQuery Editor
B. Grant each individual external consultant the role of BigQuery Viewer
C. Create a Google Group that contains the consultants and grant the group the role of BigQuery Editor
D. Create a Google Group that contains the consultants, and grant the group the role of BigQuery Viewer

Answer: D

NEW QUESTION 280
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 categor
B. Filter the Resource type to Service Account.
C. Filter the Activity log to view the Configuration categor
D. Filter the Resource type to Google Project.
E. Filter the Activity log to view the Data Access categor
F. Filter the Resource type to Service Account.
G. Filter the Activity log to view the Data Access categor
H. Filter the Resource type to Google Project.

Answer: A

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

NEW QUESTION 281
You want to configure a solution for archiving data in a Cloud Storage bucket. The solution must be
cost-effective. Data with multiple versions should be archived after 30 days. Previous versions are accessed once a month for reporting. This archive data is also occasionally updated at month-end. What should you do?

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

Answer: B

NEW QUESTION 283
You need to reduce GCP service costs for a division of your company using the fewest possible steps. You need to turn off all configured services in an existing GCP project. What should you do?

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

Answer: A

Explanation:
https://cloud.google.com/run/docs/tutorials/gcloud https://cloud.google.com/resource-manager/docs/creating-managing-projects
https://cloud.google.com/iam/docs/understanding-roles#primitive_roles
You can shut down projects using the Cloud Console. When you shut down a project, this immediately happens: All billing and traffic serving stops, You lose access to the project, The owners of the project will be notified and can stop the deletion within 30 days, The project will be scheduled to be deleted after 30 days. However, some resources may be deleted much earlier.

NEW QUESTION 288
Your company is moving from an on-premises environment to Google Cloud Platform (GCP). You have multiple development teams that use Cassandra environments as backend databases. They all need a development environment that is isolated from other Cassandra instances. You want to move to GCP quickly and with minimal support effort. What should you do?

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

Answer: B

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

NEW QUESTION 291
Your projects incurred more costs than you expected last month. Your research reveals that a development GKE container emitted a huge number of logs, which resulted in higher costs. You want to disable the logs quickly using the minimum number of steps. What should you do?

A. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE container resource.
B. 1. Go to the Logs ingestion window in Stackdriver Logging, and disable the log source for the GKE Cluster Operations resource.
C. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Logging.
D. 1. Go to the GKE console, and delete existing clusters.2. Recreate a new cluster.3. Clear the option to enable legacy Stackdriver Monitoring.

Answer: A

Explanation:
https://cloud.google.com/logging/docs/api/v2/resource-list GKE Containers have more log than GKE Cluster Operations:
-GKE Containe:
cluster_name: An immutable name for the cluster the container is running in. namespace_id: Immutable ID of the cluster namespace the container is running in. instance_id: Immutable ID of the GCE instance the container is running in. pod_id: Immutable ID of the pod the container is running in. container_name: Immutable name of the container. zone: The GCE zone in which the instance is running. VS
-GKE Cluster Operations
project_id: The identifier of the GCP project associated with this resource, such as "my-project". cluster_name: The name of the GKE Cluster.
location: The location in which the GKE Cluster is running.

NEW QUESTION 295
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

* Associate-Cloud-Engineer Questions and Answers Updated Frequently

* Associate-Cloud-Engineer Practice Questions Verified by Expert Senior Certified Staff

* Associate-Cloud-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Associate-Cloud-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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