Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1
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 2
You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your organizational policies require you to use virtual machines directly. You need to ensure that the application scaling is operationally efficient and completed as quickly as possible. What should you do?

A. Create a Google Kubernetes Engine cluster, and use horizontal pod autoscaling to scale the application.
B. Create an instance template, and use the template in a managed instance group with autoscaling configured.
C. Create an instance template, and use the template in a managed instance group that scales up and down based on the time of day.
D. Use a set of third-party tools to build automation around scaling the application up and down, based on Stackdriver CPU usage monitoring.

Answer: B

Explanation:
Managed instance groups offer autoscaling capabilities that let you automatically add or delete instances from a managed instance group based on increases or decreases in load (CPU Utilization in this case). Autoscaling helps your apps gracefully handle increases in traffic and reduce costs when the need for resources is lower. You define the autoscaling policy and the autoscaler performs automatic scaling based on the measured load (CPU Utilization in this case). Autoscaling works by adding more instances to your instance group when there is more load (upscaling), and deleting instances when the need for instances is lowered (downscaling). Ref: https://cloud.google.com/compute/docs/autoscaler

NEW QUESTION 3
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 4
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 5
You are creating an application that will run on Google Kubernetes Engine. You have identified MongoDB as the most suitable database system for your application and want to deploy a managed MongoDB environment that provides a support SLA. What should you do?

A. Create a Cloud Bigtable cluster and use the HBase API
B. Deploy MongoDB Alias from the Google Cloud Marketplace

C. Download a MongoDB installation package and run it on Compute Engine instances
D. Download a MongoDB installation package, and run it on a Managed Instance Group

Answer: B

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

NEW QUESTION 6
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 7
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 8

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 9
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 10
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 10
You are analyzing Google Cloud Platform service costs from three separate projects. You want to use this information to create service cost estimates by service type, daily and monthly, for the next six months using standard query syntax. What should you do?

A. Export your bill to a Cloud Storage bucket, and then import into Cloud Bigtable for analysis.
B. Export your bill to a Cloud Storage bucket, and then import into Google Sheets for analysis.
C. Export your transactions to a local file, and perform analysis with a desktop tool.
D. Export your bill to a BigQuery dataset, and then write time window-based SQL queries for analysis.

Answer: D

Explanation:
"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "
https://cloud.google.com/billing/docs/how-to/export-data-bigquery
https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4

NEW QUESTION 13
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 15
You are asked to set up application performance monitoring on Google Cloud projects A, B, and C as a single pane of glass. You want to monitor CPU, memory, and disk. What should you do?

A. Enable API and then share charts from project A, B, and C.
B. Enable API and then give the metrics.reader role to projects A, B, and C.
C. Enable API and then use default dashboards to view all projects in sequence.
D. Enable API, create a workspace under project A, and then add project B and C.

Answer: D

Explanation:
 https://cloud.google.com/monitoring/settings/multiple-projects https://cloud.google.com/monitoring/workspaces

NEW QUESTION 18
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.
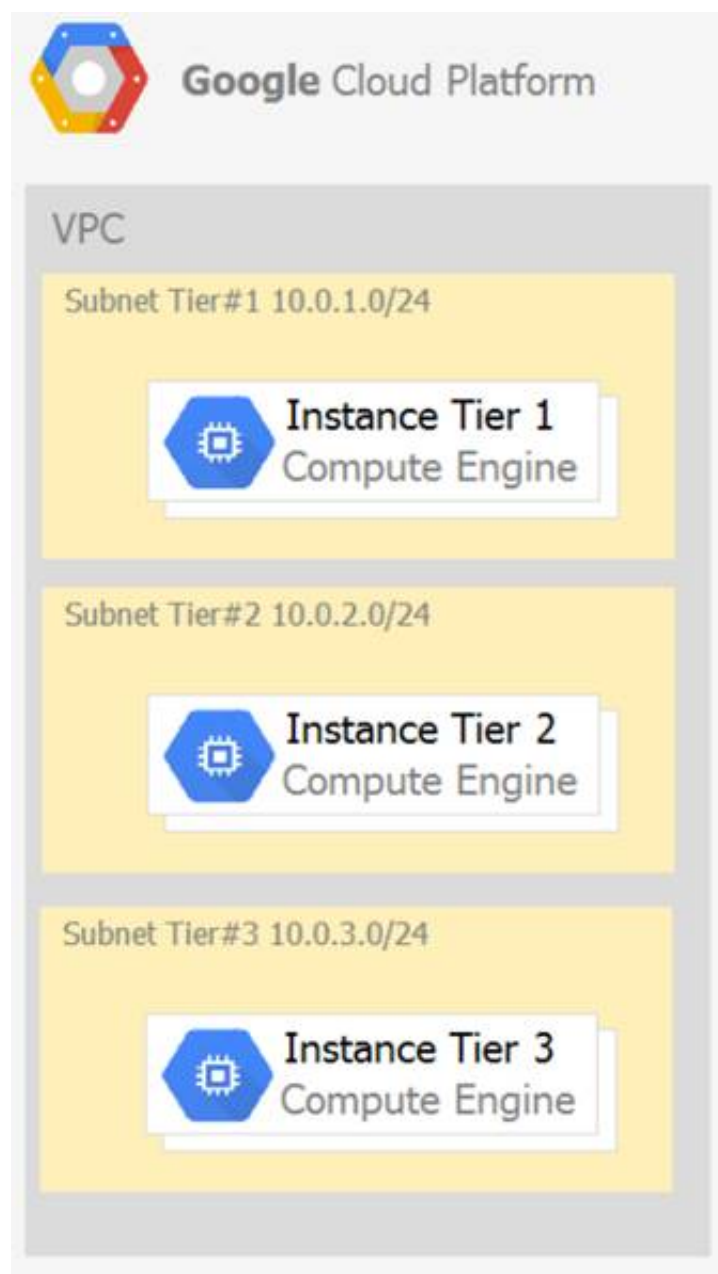D. Add the users to a group, and add this group to roles/iam.roleViewer role.

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 23
Your company has a 3-tier solution running on Compute Engine. The configuration of the current infrastructure is shown below.

Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:
• Instances in tier #1 must communicate with tier #2.
• Instances in tier #2 must communicate with tier #3. What should you do?

A. 1. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.2.0/24)• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances• Source filter: IP ranges (with the range set to 10.0.1.0/24)•Protocols: allow all

B. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow TCP:80802. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow TCP: 8080

C. 1. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #2 service account• Source filter: all instances with tier #1 service account• Protocols: allow all2. Create an ingress firewall rule with the following settings:• Targets: all instances with tier #3 service account• Source filter: all instances with tier #2 service account• Protocols: allow all

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

Answer: B

Explanation:
* 1. Create an ingress firewall rule with the following settings: "¢ Targets: all instances with tier #2 service account "¢ Source filter: all instances with tier #1 service account "¢ Protocols: allow TCP:8080 2. Create an ingress firewall rule with the following settings: "¢ Targets: all instances with tier #3 service account "¢ Source filter: all instances with tier #2 service account "¢ Protocols: allow TCP: 8080

NEW QUESTION 24
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 25
You have a number of applications that have bursty workloads and are heavily dependent on topics to decouple publishing systems from consuming systems. Your company would like to go serverless to enable developers to focus on writing code without worrying about infrastructure. Your solution architect has already identified Cloud Pub/Sub as a suitable alternative for decoupling systems. You have been asked to identify a suitable GCP Serverless service that is easy to use with Cloud Pub/Sub. You want the ability to scale down to zero when there is no traffic in order to minimize costs. You want to follow Google recommended practices. What should you suggest?

A. Cloud Run for Anthos
B. Cloud Run

C. App Engine Standard
D. Cloud Functions.

Answer: D

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

NEW QUESTION 29
Your company is moving its entire workload to Compute Engine. Some servers should be accessible through the Internet, and other servers should only be accessible over the internal network. All servers need to be able to talk to each other over specific ports and protocols. The current on-premises network relies on a demilitarized zone (DMZ) for the public servers and a Local Area Network (LAN) for the private servers. You need to design the networking infrastructure on Google Cloud to match these requirements. What should you do?

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

Answer: A

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

NEW QUESTION 33
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 37
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 38
You are storing sensitive information in a Cloud Storage bucket. For legal reasons, you need to be able to record all requests that read any of the stored data. You want to make sure you comply with these requirements. What should you do?

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

Answer: D

Explanation:
Logged information Within Cloud Audit Logs, there are two types of logs: Admin Activity logs: Entries for operations that modify the configuration or metadata of a project, bucket, or object. Data Access logs: Entries for operations that modify objects or read a project, bucket, or object. There are several sub-types of data access logs: ADMIN_READ: Entries for operations that read the configuration or metadata of a project, bucket, or object. DATA_READ: Entries for operations that read an object. DATA_WRITE: Entries for operations that create or modify an object. https://cloud.google.com/storage/docs/audit-logs#types

NEW QUESTION 39
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 41
You deployed an App Engine application using gcloud app deploy, but it did not deploy to the intended project. You want to find out why this happened and where the application deployed. What should you do?

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

Answer: D

Explanation:
 C:\GCP\appeng>gcloud config list [core]
account = xxx@gmail.com disable_usage_reporting = False project = my-first-demo-xxxx
https://cloud.google.com/endpoints/docs/openapi/troubleshoot-gce-deployment

NEW QUESTION 44
You have a web application deployed as a managed instance group. You have a new version of the application to gradually deploy. Your web application is currently receiving live web traffic. You want to ensure that the available capacity does not decrease during the deployment. What should you do?

A. Perform a rolling-action start-update with maxSurge set to 0 and maxUnavailable set to 1.
B. Perform a rolling-action start-update with maxSurge set to 1 and maxUnavailable set to 0.
C. Create a new managed instance group with an updated instance templat
D. Add the group to the backend service for the load balance
E. When all instances in the new managed instance group are healthy, delete the old managed instance group.
F. Create a new instance template with the new application versio
G. Update the existing managed instance group with the new instance templat
H. Delete the instances in the managed instance group to allow the managed instance group to recreate the instance using the new instance template.

Answer: B

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

NEW QUESTION 48
You are building an archival solution for your data warehouse and have selected Cloud Storage to archive your data. Your users need to be able to access this archived data once a quarter for some regulatory requirements. You want to select a cost-efficient option. Which storage option should you use?

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

Answer: A

Explanation:
Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Since we have a requirement to access data once a quarter and want to go with the most cost-efficient option, we should select Coldline Storage.
Ref: https://cloud.google.com/storage/docs/storage-classes#coldline

Google Cloud Storage Classes in the Organization

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

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

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

NEW QUESTION 53

You are building a product on top of Google Kubernetes Engine (GKE). You have a single GKE cluster. For each of your customers, a Pod is running in that cluster, and your customers can run arbitrary code inside their Pod. You want to maximize the isolation between your customers' Pods. What should you do?

A. Use Binary Authorization and whitelist only the container images used by your customers' Pods.
B. Use the Container Analysis API to detect vulnerabilities in the containers used by your customers' Pods.
C. Create a GKE node pool with a sandbox type configured to gviso
D. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.
E. Use the cos_containerd image for your GKE node
F. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Answer: C

NEW QUESTION 55

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 60

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 63
You need to provide a cost estimate for a Kubernetes cluster using the GCP pricing calculator for Kubernetes. Your workload requires high IOPs, and you will also be using disk snapshots. You start by entering the number of nodes, average hours, and average days. What should you do next?

A. Fill in local SS
B. Fill in persistent disk storage and snapshot storage.
C. Fill in local SS
D. Add estimated cost for cluster management.
E. Select Add GPU
F. Fill in persistent disk storage and snapshot storage.
G. Select Add GPU
H. Add estimated cost for cluster management.

Answer: A

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

NEW QUESTION 68
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 73
You have a Google Cloud Platform account with access to both production and development projects. You need to create an automated process to list all compute instances in development and production projects on a daily basis. What should you do?

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

Answer: A

Explanation:
You can create two configurations – one for the development project and another for the production project. And you do that by running "gcloud config configurations create" command.https://cloud.google.com/sdk/gcloud/reference/config/configurations/createIn your custom script, you can load these configurations one at a time and execute gcloud compute instances list to list Google Compute Engine instances in the project that is active in the gcloud configuration.Ref: https://cloud.google.com/sdk/gcloud/reference/compute/instances/listOnce you have this information, you can export it in a suitable format to a suitable target e.g. export as CSV or export to Cloud Storage/BigQuery/SQL, etc

NEW QUESTION 74
You created an instance of SQL Server 2017 on Compute Engine to test features in the new version. You want to connect to this instance using the fewest number of steps. What should you do?

A. Install a RDP client on your deskto
B. Verify that a firewall rule for port 3389 exists.
C. Install a RDP client in your deskto
D. Set a Windows username and password in the GCP Consol
E. Use the credentials to log in to the instance.
F. Set a Windows password in the GCP Consol
G. Verify that a firewall rule for port 22 exist
H. Click the RDP button in the GCP Console and supply the credentials to log in.
I. Set a Windows username and password in the GCP Consol
J. Verify that a firewall rule for port 3389 exist
K. Click the RDP button in the GCP Console, and supply the credentials to log in.

Answer: D

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

NEW QUESTION 75
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 77
Your coworker has helped you set up several configurations for gcloud. You've noticed that you're running commands against the wrong project. Being new to the company, you haven't yet memorized any of the projects. With the fewest steps possible, what's the fastest way to switch to the correct configuration?

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

Answer: C

Explanation:
as gcloud config configurations list can help check for the existing configurations and activate can help switch to the configuration.
gcloud config configurations list lists existing named configurations
gcloud config configurations activate activates an existing named configuration
Obtains access credentials for your user account via a web-based authorization flow. When this command completes successfully, it sets the active account in the current configuration to the account specified. If no configuration exists, it creates a configuration named default.

NEW QUESTION 78
Your organization has a dedicated person who creates and manages all service accounts for Google Cloud projects. You need to assign this person the minimum role for projects. What should you do?

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

Answer: D

NEW QUESTION 83
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 87
You want to send and consume Cloud Pub/Sub messages from your App Engine application. The Cloud Pub/Sub API is currently disabled. You will use a service account to authenticate your application to the API. You want to make sure your application can use Cloud Pub/Sub. What should you do?

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

Answer: A

Explanation:
Quickstart: using the Google Cloud Console
This page shows you how to perform basic tasks in Pub/Sub using the Google Cloud Console. Note: If you are new to Pub/Sub, we recommend that you start with

the interactive tutorial. Before you begin
Set up a Cloud Console project. Set up a project
Click to:
Create or select a project.
Enable the Pub/Sub API for that project.
You can view and manage these resources at any time in the Cloud Console. Install and initialize the Cloud SDK.
Note: You can run the gcloud tool in the Cloud Console without installing the Cloud SDK. To run the gcloud tool in the Cloud Console, use Cloud Shell .
https://cloud.google.com/pubsub/docs/quickstart-console

NEW QUESTION 90
You are managing several Google Cloud Platform (GCP) projects and need access to all logs for the past 60 days. You want to be able to explore and quickly analyze the log contents. You want to follow Google- recommended practices to obtain the combined logs for all projects. What should you do?

A. Navigate to Stackdriver Logging and select resource.labels.project_id="*"
B. Create a Stackdriver Logging Export with a Sink destination to a BigQuery datase
C. Configure the table expiration to 60 days.
D. Create a Stackdriver Logging Export with a Sink destination to Cloud Storag
E. Create a lifecycle rule to delete objects after 60 days.
F. Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuer
G. Configure the table expiration to 60 days.

Answer: B

Explanation:
≫ Navigate to Stackdriver Logging and select resource.labels.project_id=*. is not right.
Log entries are held in Stackdriver Logging for a limited time known as the retention period which is 30 days (default configuration). After that, the entries are deleted. To keep log entries longer, you need to export them outside of Stackdriver Logging by configuring log sinks.
Ref: https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging ≫ Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days. is not right.
While this works, it makes no sense to use Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery when Google provides a feature (export sinks) that does exactly the same thing and works out of the box.Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
≫ Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days. is not right.
You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and
Pub/Sub.Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organication, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud
organization.Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks
Either way, we now have the data in Cloud Storage, but querying logs information from Cloud Storage is harder than Querying information from BigQuery dataset. For this reason, we should prefer Big Query over Cloud Storage.
≫ Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days. is the right answer.
You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and
Pub/Sub.Ref: https://cloud.google.com/logging/docs/export/configure_export_v2
Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organication, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud
organization.Ref: https://cloud.google.com/logging/docs/export/aggregated_sinks
Either way, we now have the data in a BigQuery Dataset. Querying information from a Big Query dataset is easier and quicker than analyzing contents in Cloud Storage bucket. As our requirement is to Quickly analyze the log contents, we should prefer Big Query over Cloud Storage.
Also, You can control storage costs and optimize storage usage by setting the default table expiration for newly created tables in a dataset. If you set the property when the dataset is created, any table created in the dataset is deleted after the expiration period. If you set the property after the dataset is created, only new tables are deleted after the expiration period.For example, if you set the default table expiration to 7 days, older data is automatically deleted after 1 week.Ref: https://cloud.google.com/bigquery/docs/best-practices-storage

NEW QUESTION 92
You deployed a new application inside your Google Kubernetes Engine cluster using the YAML file specified below.

```
apiVersion: apps/v1              apiVersion: v1
kind: Deployment                 kind: Service
metadata:                        metadata:
  name: myapp-deployment           name: myapp-service
spec:                            spec:
  selector:                        ports:
    matchLabels:                   - port: 8000
      app: myapp                     targetPort: 80
  replicas: 2                        protocol: TCP
  template:                        selector:
    metadata:                        app: myapp
      labels:
        app: myapp
      spec:
        containers:
        - name: myapp
          image: myapp:1.1
          ports:
          - containerPort: 80
```

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

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

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

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

Answer: C

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

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

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

Answer: A

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

NEW QUESTION 98
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 99
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 102
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 103
A colleague handed over a Google Cloud Platform project for you to maintain. As part of a security checkup, you want to review who has been granted the Project Owner role. What should you do?

A. In the console, validate which SSH keys have been stored as project-wide keys.
B. Navigate to Identity-Aware Proxy and check the permissions for these resources.
C. Enable Audit Logs on the IAM & admin page for all resources, and validate the results.
D. Use the command gcloud projects get–iam–policy to view the current role assignments.

Answer: D

Explanation:
A simple approach would be to use the command flags available when listing all the IAM policy for a given project. For instance, the following command: `gcloud projects get-iam-policy $PROJECT_ID
--flatten="bindings[].members" --format="table(bindings.members)" --filter="bindings.role:roles/owner"` outputs all the users and service accounts associated with the role 'roles/owner' in the project in question. https://groups.google.com/g/google-cloud-dev/c/Z6sZs7TvygQ?pli=1

NEW QUESTION 104
Your company has an existing GCP organization with hundreds of projects and a billing account. Your
company recently acquired another company that also has hundreds of projects and its own billing account. You would like to consolidate all GCP costs of both GCP organizations onto a single invoice. You would like to consolidate all costs as of tomorrow. What should you do?

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

Answer: A

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

NEW QUESTION 106
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 110
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 111
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 115
You want to select and configure a solution for storing and archiving data on Google Cloud Platform. You need to support compliance objectives for data from one geographic location. This data is archived after 30 days and needs to be accessed annually. What should you do?

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

Answer: D

Explanation:
Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.
The Real description is about Coldline storage Class: Coldline Storage
Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is a better choice than Standard Storage or Nearline Storage in scenarios where slightly lower availability, a 90-day minimum storage duration, and higher costs for data access are acceptable trade-offs for lowered at-rest storage costs.
Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Note, however, that for data being kept entirely for backup or archiving purposes, Archive Storage is more cost-effective, as it offers the lowest storage costs.
https://cloud.google.com/storage/docs/storage-classes#coldline

NEW QUESTION 116
You have a project for your App Engine application that serves a development environment. The required testing has succeeded and you want to create a new project to serve as your production environment. What should you do?

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

Answer: A

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

NEW QUESTION 117
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 120
An application generates daily reports in a Compute Engine virtual machine (VM). The VM is in the project corp-iot-insights. Your team operates only in the project corp-aggregate-reports and needs a copy of the daily exports in the bucket corp-aggregate-reports-storage. You want to configure access so that the daily reports from the VM are available in the bucket corp-aggregate-reports-storage and use as few steps as possible while following Google-recommended practices. What should you do?

A. Move both projects under the same folder.
B. Grant the VM Service Account the role Storage Object Creator on corp-aggregate-reports-storage.
C. Create a Shared VPC network between both project
D. Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
E. Make corp-aggregate-reports-storage public and create a folder with a pseudo-randomized suffix name.Share the folder with the IoT team.

Answer: B

Explanation:
 Predefined roles
The following table describes Identity and Access Management (IAM) roles that are associated with Cloud Storage and lists the permissions that are contained in each role. Unless otherwise noted, these roles can be applied either to entire projects or specific buckets.
Storage Object Creator (roles/storage.objectCreator) Allows users to create objects. Does not give permission to view, delete, or overwrite objects.
https://cloud.google.com/storage/docs/access-control/iam-roles#standard-roles

NEW QUESTION 121
You have an application on a general-purpose Compute Engine instance that is experiencing excessive disk read throttling on its Zonal SSD Persistent Disk. The application primarily reads large files from disk. The disk size is currently 350 GB. You want to provide the maximum amount of throughput while minimizing costs. What should you do?

A. Increase the size of the disk to 1 TB.
B. Increase the allocated CPU to the instance.
C. Migrate to use a Local SSD on the instance.

D. Migrate to use a Regional SSD on the instance.

Answer: C

Explanation:
Standard persistent disks are efficient and economical for handling sequential read/write operations, but they aren't optimized to handle high rates of random input/output operations per second (IOPS). If your apps require high rates of random IOPS, use SSD persistent disks. SSD persistent disks are designed for single-digit millisecond latencies. Observed latency is application specific.

NEW QUESTION 126
You are using Deployment Manager to create a Google Kubernetes Engine cluster. Using the same Deployment Manager deployment, you also want to create a DaemonSet in the kube-system namespace of the cluster. You want a solution that uses the fewest possible services. What should you do?

A. Add the cluster's API as a new Type Provider in Deployment Manager, and use the new type to create the DaemonSet.
B. Use the Deployment Manager Runtime Configurator to create a new Config resource that contains the DaemonSet definition.
C. With Deployment Manager, create a Compute Engine instance with a startup script that uses kubectl to create the DaemonSet.
D. In the cluster's definition in Deployment Manager, add a metadata that has kube-system as key and the DaemonSet manifest as value.

Answer: A

Explanation:
Adding an API as a type provider
This page describes how to add an API to Google Cloud Deployment Manager as a type provider. To learn more about types and type providers, read the Types overview documentation.
A type provider exposes all of the resources of a third-party API to Deployment Manager as base types that you can use in your configurations. These types must be directly served by a RESTful API that supports Create, Read, Update, and Delete (CRUD).
If you want to use an API that is not automatically provided by Google with Deployment Manager, you must add the API as a type provider.
https://cloud.google.com/deployment-manager/docs/configuration/type-providers/creating-type-provider

NEW QUESTION 129
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 133
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 134
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 135
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 140
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 141
You have developed a containerized web application that will serve Internal colleagues during business hours. You want to ensure that no costs are incurred outside of the hours the application is used. You have just created a new Google Cloud project and want to deploy the application. What should you do?

A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.
C. Deploy the container on App Engine flexible environment with autoscalin
D. and set the value min_instances to zero in the app yaml
E. Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app yaml

Answer: B

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

NEW QUESTION 145
You need to set a budget alert for use of Compute Engineer services on one of the three Google Cloud Platform projects that you manage. All three projects are linked to a single billing account. What should you do?

A. Verify that you are the project billing administrato
B. Select the associated billing account and create a budget and alert for the appropriate project.
C. Verify that you are the project billing administrato
D. Select the associated billing account and create a budget and a custom alert.
E. Verify that you are the project administrato
F. Select the associated billing account and create a budget for the appropriate project.
G. Verify that you are project administrato
H. Select the associated billing account and create a budget and a custom alert.

Answer: A

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

NEW QUESTION 146
You want to add a new auditor to a Google Cloud Platform project. The auditor should be allowed to read, but not modify, all project items.
How should you configure the auditor's permissions?

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

Answer: C

NEW QUESTION 147
You are assigned to maintain a Google Kubernetes Engine (GKE) cluster named dev that was deployed on Google Cloud. You want to manage the GKE configuration using the command line interface (CLI). You have just downloaded and installed the Cloud SDK. You want to ensure that future CLI commands by default address this specific cluster. What should you do?

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

Answer: A

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

NEW QUESTION 149
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 152
You built an application on Google Cloud Platform that uses Cloud Spanner. Your support team needs to monitor the environment but should not have access to table data. You need a streamlined solution to grant the correct permissions to your support team, and you want to follow Google-recommended practices. What should you do?

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

Answer: A

Explanation:
roles/monitoring.viewer provides read-only access to get and list information about all monitoring data and configurations. This role provides monitoring access and fits our requirements. roles/monitoring.viewer. is the right answer.
Ref: https://cloud.google.com/iam/docs/understanding-roles#cloud-spanner-roles

NEW QUESTION 154
Every employee of your company has a Google account. Your operational team needs to manage a large number of instances on Compute Engine. Each member of this team needs only administrative access to the servers. Your security team wants to ensure that the deployment of credentials is operationally efficient and must be able to determine who accessed a given instance. What should you do?

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

Answer: C

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

NEW QUESTION 158
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 162
You are using multiple configurations for gcloud. You want to review the configured Kubernetes Engine cluster of an inactive configuration using the fewest possible steps. What should you do?

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

Answer: D

NEW QUESTION 164
You need to monitor resources that are distributed over different projects in Google Cloud Platform. You want to consolidate reporting under the same Stackdriver Monitoring dashboard. What should you do?

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

Answer: C

Explanation:
When you intially click on Monitoring(Stackdriver Monitoring) it creates a workspac(a stackdriver account) linked to the ACTIVE(CURRENT) Project from which it was clicked.
Now if you change the project and again click onto Monitoring it would create an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don't want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).
If you have accidently created two diff workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.
If we have only one workspace and two projects we can simply add other GCP Project under Monitoring > Settings > GCP Projects > Add GCP Projects.
https://cloud.google.com/monitoring/settings/multiple-projects
Nothing about groups https://cloud.google.com/monitoring/settings?hl=en

NEW QUESTION 169
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 174
You have a managed instance group comprised of preemptible VM's. All of the VM's keepdeleting and recreating themselves every minute. What is a possible cause of thisbehavior?

A. Your zonal capacity is limited, causing all preemptible VM's to be shutdown torecover capacit
B. Try deploying your group to another zone.
C. You have hit your instance quota for the region.
D. Your managed instance group's VM's are toggled to only last 1 minute inpreemptible settings.
E. Your managed instance group's health check is repeatedly failing, either to amisconfigured health check or misconfigured firewall rules not allowing the healthcheck to access the instance

Answer: D

Explanation:
as the instances (normal or preemptible) would be terminated and relaunched if the health check fails either due to application not configured properly or the instances firewall do not allow health check to happen.
GCP provides health check systems that connect to virtual machine (VM) instances on a configurable, periodic basis. Each connection attempt is called a probe. GCP records the success or failure of each probe.
Health checks and load balancers work together. Based on a configurable number of sequential successful or failed probes, GCP computes an overall health state for each VM in the load balancer. VMs that respond successfully for the configured number of times are considered healthy. VMs that fail to respond successfully for a separate number of times are unhealthy.
GCP uses the overall health state of each VM to determine its eligibility for receiving new requests. In addition to being able to configure probe frequency and health state thresholds, you can configure the criteria that define a successful probe.

NEW QUESTION 175
......

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