Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

About Exambible

Your Partner of IT Exam

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

    All examinations will be up to date.

* 24/7 Quality Support

    We will provide service round the clock.

* 100% Pass Rate

    Our guarantee that you will pass the exam.

* Unique Gurantee

    If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

You recently discovered that your developers are using many service account keys during their development process. While you work on a long term improvement, you need to quickly implement a process to enforce short-lived service account credentials in your company. You have the following requirements:
• All service accounts that require a key should be created in a centralized project called pj-sa.
• Service account keys should only be valid for one day.

You need a Google-recommended solution that minimizes cost. What should you do?

A. Implement a Cloud Run job to rotate all service account keys periodically in pj-s
B. Enforce an org policy to deny service account key creation with an exception to pj-sa.
C. Implement a Kubernetes Cronjob to rotate all service account keys periodicall
D. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.
E. Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hour
F. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.
G. Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hour
H. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

Answer: C

Explanation:
According to the Google Cloud documentation, you can use organization policy constraints to control the creation and expiration of service account keys. The constraints are:

⟩ constraints/iam.allowServiceAccountKeyCreation: This constraint allows you to specify which projects
or folders can create service account keys. You can set the value to true or false, or use a condition to apply the constraint to specific service accounts. By setting this constraint to false for the organization and adding an exception for the pj-sa project, you can prevent developers from creating service account keys in other projects.

⟩ constraints/iam.serviceAccountKeyMaxLifetime: This constraint allows you to specify the maximum lifetime of service account keys. You can set the value to a duration in seconds, such as 86400 for one day. By setting this constraint to 86400 for the organization, you can ensure that all service account ke expire after one day.

These constraints are recommended by Google Cloud as best practices to minimize the risk of service account key misuse or compromise. They also help you reduce the cost of managing service account keys, as you do not need to implement a custom solution to rotate or delete them.
References:

⟩ 1: Associate Cloud Engineer Certification Exam Guide | Learn - Google Cloud

⟩ 5: Create and delete service account keys - Google Cloud

⟩ Organization policy constraints for service accounts

NEW QUESTION 2

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 3

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

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

Answer: A

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

NEW QUESTION 4

Your learn wants to deploy a specific content management system (CMS) solution lo Google Cloud. You need a quick and easy way to deploy and install the solution. What should you do?

A. Search for the CMS solution in Google Cloud Marketplac
B. Use gcloud CLI to deploy the solution.
C. Search for the CMS solution in Google Cloud Marketplac
D. Deploy the solution directly from Cloud Marketplace.
E. Search for the CMS solution in Google Cloud Marketplac
F. Use Terraform and the Cloud Marketplace ID to deploy the solution with the appropriate parameters.
G. Use the installation guide of the CMS provide
H. Perform the installation through your configuration management system.

Answer: B

NEW QUESTION 5
You have an on-premises data analytics set of binaries that processes data files in memory for about 45 minutes every midnight. The sizes of those data files range from 1 gigabyte to 16 gigabytes. You want to migrate this application to Google Cloud with minimal effort and cost. What should you do?

A. Upload the code to Cloud Function
B. Use Cloud Scheduler to start the application.
C. Create a container for the set of binarie
D. Use Cloud Scheduler to start a Cloud Run job for the container.
E. Create a container for the set of binaries Deploy the container to Google Kubernetes Engine (GKE) and use the Kubernetes scheduler to start the application.
F. Lift and shift to a VM on Compute Engin
G. Use an instance schedule to start and stop the instance.

Answer: B

NEW QUESTION 6
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. 1.00.0.0/0
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 7
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 8
You have a development project with appropriate IAM roles defined. You are creating a production project and want to have the same IAM roles on the new project, using the fewest possible steps. What should you do?

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

Answer: A

NEW QUESTION 9
Your company has multiple projects linked to a single billing account in Google Cloud. You need to visualize the costs with specific metrics that should be dynamically calculated based on company-specific criteria. You want to automate the process. What should you do?

A. In the Google Cloud console, visualize the costs related to the projects in the Reports section.
B. In the Google Cloud console, visualize the costs related to the projects in the Cost breakdown section.
C. In the Google Cloud console, use the export functionality of the Cost tabl
D. Create a Looker Studiodashboard on top of the CSV export.
E. Configure Cloud Billing data export to BigOuery for the billing accoun
F. Create a Looker Studio dashboard on top of the BigOuery export.

Answer: D

NEW QUESTION 10
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 10
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 15
Your company's security vulnerability management policy wonts 3 member of the security team to have visibility into vulnerabilities and other OS metadata for a specific Compute Engine instance This Compute Engine instance hosts a critical application in your Goggle Cloud project. You need to implement your company's security vulnerability management policy. What should you dc?

A. • Ensure that the Ops Agent Is Installed on the Compute Engine instance.• Create a custom metric in the Cloud Monitoring dashboard.• Provide the security team member with access to this dashboard.
B. • Ensure that the Ops Agent is installed on tie Compute Engine instance.• Provide the security team member roles/configure.inventoryViewer permission.
C. • Ensure that the OS Config agent Is Installed on the Compute Engine instance.• Provide the security team member roles/configure.vulnerabilityViewer permission.
D. • Ensure that the OS Config agent is installed on the Compute Engine instance• Create a log sink Co a BigQuery dataset.• Provide the security team member with access to this dataset.

Answer: C

NEW QUESTION 19
You have developed an application that consists of multiple microservices, with each microservice packaged in its own Docker container image. You want to deploy the entire application on Google Kubernetes Engine so
that each microservice can be scaled individually. What should you do?

A. Create and deploy a Custom Resource Definition per microservice.
B. Create and deploy a Docker Compose File.
C. Create and deploy a Job per microservice.
D. Create and deploy a Deployment per microservice.

Answer: A

NEW QUESTION 24
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 26
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 28
You are hosting an application from Compute Engine virtual machines (VMs) in us–central1–a. You want to adjust your design to support the failure of a single Compute Engine zone, eliminate downtime, and minimize cost. What should you do?

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

Answer: A

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

NEW QUESTION 33
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 38
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 42
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 45
A company wants to build an application that stores images in a Cloud Storage bucket and wants to generate thumbnails as well as resize the images. They want to use a google managed service that can scale up and scale down to zero automatically with minimal effort. You have been asked to recommend a service. Which GCP service would you suggest?

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

Answer: C

Explanation:
Text Description automatically generated with low confidence
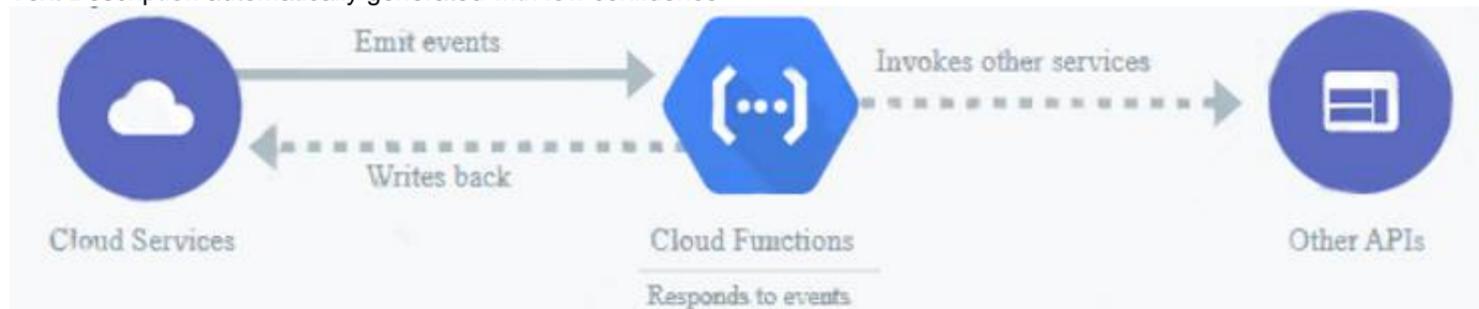
Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.
Ref: https://cloud.google.com/functions
While all other options i.e. Google Compute Engine, Google Kubernetes Engine, Google App Engine support autoscaling, it needs to be configured explicitly based on the load and is not as trivial as the scale up or scale down offered by Google's cloud functions.

NEW QUESTION 47
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 49
You are in charge of provisioning access for all Google Cloud users in your organization. Your company recently acquired a startup company that has their own Google Cloud organization. You need to ensure that your Site Reliability Engineers (SREs) have the same project permissions in the startup company's organization as in your own organization. What should you do?

A. In the Google Cloud console for your organization, select Create role from selection, and choose destination as the startup company's organization
B. In the Google Cloud console for the startup company, select Create role from selection and choose source as the startup company's Google Cloud organization.
C. Use the gcloud iam roles copy command, and provide the Organization ID of the startup company's Google Cloud Organization as the destination.
D. Use the gcloud iam roles copy command, and provide the project IDs of all projects in the startup company s organization as the destination.

Answer: C

Explanation:

https://cloud.google.com/architecture/best-practices-vpc-design#shared-service Cloud VPN is another alternative. Because Cloud VPN establishes reachability through managed IPsec tunnels, it doesn't have the aggregate limits of VPC Network Peering. Cloud VPN uses a VPN Gateway for connectivity and doesn't consider the aggregate resource use of the IPsec peer. The drawbacks of Cloud VPN include increased costs (VPN tunnels and traffic egress), management overhead required to maintain tunnels, and the performance overhead of IPsec.

NEW QUESTION 50
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 51
You have deployed an application on a Compute Engine instance. An external consultant needs to access the Linux-based instance. The consultant is connected to your corporate network through a VPN connection, but the consultant has no Google account. What should you do?

A. Instruct the external consultant to use the gcloud compute ssh command line tool by using Identity-Aware Proxy to access the instance.
B. Instruct the external consultant to use the gcloud compute ssh command line tool by using the public IP address of the instance to access it.
C. Instruct the external consultant to generate an SSH key pair, and request the public key from the consultant.Add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key.
D. Instruct the external consultant to generate an SSH key pair, and request the private key from the consultant.Add the private key to the instance yourself, and have the consultant access the instance through SSH with their public key.

Answer: C

Explanation:
The best option is to instruct the external consultant to generate an SSH key pair, and request the public key from the consultant. Then, add the public key to the instance yourself, and have the consultant access the instance through SSH with their private key. This way, you can grant the consultant access to the instance without requiring a Google account or exposing the instance's public IP address. This option also follows the best practice of using user-managed SSH keys instead of service account keys for SSH access1.
Option A is not feasible because the external consultant does not have a Google account, and therefore cannot use Identity-Aware Proxy (IAP) to access the instance. IAP requires the user to authenticate with a Google account and have the appropriate IAM permissions to access the instance2. Option B is not secure because it exposes the instance's public IP address, which can increase the risk of unauthorized access or attacks. Option D is not correct because it reverses the roles of the public and private keys. The public key should be added to the instance, and the private key should be kept by the consultant. Sharing the private key with anyone else can compromise the security of the SSH connection3.
References:
⪼ 1: https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys
⪼ 2: https://cloud.google.com/iap/docs/using-tcp-forwarding
⪼ 3: https://cloud.google.com/compute/docs/instances/connecting-advanced#sshbetweeninstances

NEW QUESTION 55
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 57
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 61
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 62
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 67
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 69
Your company completed the acquisition of a startup and is now merging the IT systems of both companies. The startup had a production Google Cloud project in their organization. You need to move this project into your organization and ensure that the project is billed lo your organization. You want to accomplish this task with minimal effort. What should you do?

A. Use the project
B. move method to move the project to your organizatio

C. Update the billing account of the project to that of your organization.
D. Ensure that you have an Organization Administrator Identity and Access Management (IAM) role assigned to you in both organization
E. Navigate to the Resource Manager in the startup's Google Cloud organization, and drag the project to your company's organization.
F. Create a Private Catalog tor the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalo
G. Share the Catalog with your organization, and deploy the resources in your company's project.
H. Create an infrastructure-as-code template tor all resources in the project by using Terrafor
I. and deploy that template to a new project in your organizatio
J. Delete the protect from the startup's Google Cloud organization.

Answer: A

NEW QUESTION 73
You are configuring Cloud DNS. You want !to create DNS records to point home.mydomain.com, mydomain.com. and www.mydomain.com to the IP address of your Google Cloud load balancer. What should you do?

A. Create one CNAME record to point mydomain.com to the load balancer, and create two A records to point WWW and HOME lo mydomain.com respectively.
B. Create one CNAME record to point mydomain.com to the load balancer, and create two AAAA records to point WWW and HOME to mydomain.com respectively.
C. Create one A record to point mydomain.com to the load balancer, and create two CNAME records to point WWW and HOME to mydomain.com respectively.
D. Create one A record to point mydomain.com lo the load balancer, and create two NS records to point WWW and HOME to mydomain.com respectively.

Answer: C

NEW QUESTION 75
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 79
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 81
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 84
You used the gcloud container clusters command to create two Google Cloud Kubernetes (GKE) clusters prod-cluster and dev-cluster.
• prod-cluster is a standard cluster.
• dev-cluster is an auto-pilot duster.
When you run the Kubect1 get nodes command, you only see the nodes from prod-cluster Which commands should you run to check the node status for dev-cluster?

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

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

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

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

Answer: C

NEW QUESTION 88
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 89
You manage three Google Cloud projects with the Cloud Monitoring API enabled. You want to follow Google-recommended practices to visualize CPU and network metrics for all three projects together. What should you do?

A. * 1. Create a Cloud Monitoring Dashboard* 2. Collect metrics and publish them into the Pub/Sub topics 3. Add CPU and network Charts (or each of (he three projects
B. * 1. Create a Cloud Monitoring Dashboard.* 2. Select the CPU and Network metrics from the three projects.* 3. Add CPU and network Charts lot each of the three protects.
C. * 1 Create a Service Account and apply roles/viewer on the three projects* 2. Collect metrics and publish them lo the Cloud Monitoring API* 3. Add CPU and network Charts for each of the three projects.
D. * 1. Create a fourth Google Cloud project* 2 Create a Cloud Workspace from the fourth project and add the other three projects

Answer: B

NEW QUESTION 91
You are configuring service accounts for an application that spans multiple projects. Virtual machines (VMs) running in the web-applications project need access to BigQuery datasets in crm-databases-proj. You want to follow Google-recommended practices to give access to the service account in the web-applications project. What should you do?

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

Answer: C

NEW QUESTION 93
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 94
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 99
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 101
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 102
You are managing a project for the Business Intelligence (BI) department in your company. A data pipeline ingests data into BigQuery via streaming. You want the users in the BI department to be able to run the custom SQL queries against the latest data in BigQuery. What should you do?

A. Create a Data Studio dashboard that uses the related BigQuery tables as a source and give the BI team view access to the Data Studio dashboard.
B. Create a Service Account for the BI team and distribute a new private key to each member of the BI team.
C. Use Cloud Scheduler to schedule a batch Dataflow job to copy the data from BigQuery to the BI team's internal data warehouse.
D. Assign the IAM role of BigQuery User to a Google Group that contains the members of the BI team.

Answer: D

Explanation:
When applied to a dataset, this role provides the ability to read the dataset's metadata and list tables in the dataset. When applied to a project, this role also provides the ability to run jobs, including queries, within the project. A member with this role can enumerate their own jobs, cancel their own jobs, and enumerate datasets within a project. Additionally, allows the creation of new datasets within the project; the creator is granted the BigQuery Data Owner role (roles/bigquery.dataOwner) on these new datasets.
https://cloud.google.com/bigquery/docs/access-control

NEW QUESTION 104
You are developing a financial trading application that will be used globally. Data is stored and queried using a relational structure, and clients from all over the world should get the exact identical state of the data. The application will be deployed in multiple regions to provide the lowest latency to end users. You need to select a storage option for the application data while minimizing latency. What should you do?

A. Use Cloud Bigtable for data storage.
B. Use Cloud SQL for data storage.
C. Use Cloud Spanner for data storage.
D. Use Firestore for data storage.

Answer: C

Explanation:
Keywords, Financial data (large data) used globally, data stored and queried using relational structure (SQL), clients should get exact identical copies(Strong Consistency), Multiple region, low latency to end user, select storage option to minimize latency.

NEW QUESTION 107
You are working with a user to set up an application in a new VPC behind a firewall. The user is concerned about data egress. You want to configure the fewest open egress ports. What should you do?

A. Set up a low-priority (65534) rule that blocks all egress and a high-priority rule (1000) that allows only the appropriate ports.
B. Set up a high-priority (1000) rule that pairs both ingress and egress ports.
C. Set up a high-priority (1000) rule that blocks all egress and a low-priority (65534) rule that allows only the appropriate ports.
D. Set up a high-priority (1000) rule to allow the appropriate ports.

Answer: A

Explanation:
Implied rules Every VPC network has two implied firewall rules. These rules exist, but are not shown in the Cloud Console: Implied allow egress rule. An egress rule whose action is allow, destination is 0.0.0.0/0, and priority is the lowest possible (65535) lets any instance send traffic to any destination, except for traffic blocked by Google Cloud. A higher priority firewall rule may restrict outbound access. Internet access is allowed if no other firewall rules deny outbound traffic and if the instance has an external IP address or uses a Cloud NAT instance. For more information, see Internet access requirements. Implied deny ingress rule. An ingress rule whose action is deny, source is 0.0.0.0/0, and priority is the lowest possible (65535) protects all instances by blocking incoming connections to them. A higher priority rule might allow incoming access. The default network includes some additional rules that override this one, allowing certain types of incoming connections. https://cloud.google.com/vpc/docs/firewalls#default_firewall_rules

NEW QUESTION 110
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 114
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 116
Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises workloads consist of:
• A Flask web application
• AbackendAPI
• A scheduled long-running background job for ETL and reporting.
You need to keep operational costs low You want to follow Google-recommended practices to migrate these workloads to serverless solutions on Google Cloud.
What should you do?

A. Migrate the web application to App Engine and the backend API to Cloud Run Use Cloud Tasks to run your background job on Compute Engine
B. Migrate the web application to App Engine and the backend API to Cloud Ru
C. Use Cloud Tasks to run your background job on Cloud Run.
D. Run the web application on a Cloud Storage bucket and the backend API on Cloud Run Use Cloud Tasks to run your background job on Cloud Run.
E. Run the web application on a Cloud Storage bucket and the backend API on Cloud Ru
F. Use Cloud Tasks to run your background job on Compute Engine

Answer: B

NEW QUESTION 117
You have an instance group that you want to load balance. You want the load balancer to terminate the client SSL session. The instance group is used to serve a public web application over HTTPS. You want to follow Google-recommended practices. What should you do?

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

Answer: A

NEW QUESTION 119
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 123
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 128
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 132

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 133

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

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

Answer: D

Explanation:
https://cloud.google.com/appengine/docs/flexible/managing-projects-apps-billing#:~:text=Each%20Cloud%20p Two App engine can't be running on the same project: you can check this easy diagram for more info:
https://cloud.google.com/appengine/docs/standard/an-overview-of-app-engine#components_of_an_application
And you can't change location after setting it for your app Engine. https://cloud.google.com/appengine/docs/standard/locations
App Engine is regional and you cannot change an apps region after you set it. Therefore, the only way to have an app run in another region is by creating a new project and targeting the app engine to run in the required region (asia-northeast1 in our case).
Ref: https://cloud.google.com/appengine/docs/locations

NEW QUESTION 138

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 thecurrent 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 139

The storage costs for your application logs have far exceeded the project budget. The logs are currently being retained indefinitely in the Cloud Storage bucket myapp-gcp-ace-logs. You have been asked to remove logs older than 90 days from your Cloud Storage bucket. You want to optimize ongoing Cloud Storage spend. What should you do?

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

Answer: B

Explanation:
You write a lifecycle management rule in XML and push it to the bucket with gsutil lifecycle set config-xml-file. is not right.
gsutil lifecycle set enables you to set the lifecycle configuration on one or more buckets based on the configuration file provided. However, XML is not a valid supported type for the configuration file.
Ref: https://cloud.google.com/storage/docs/gsutil/commands/lifecycle

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

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

> Write a lifecycle management rule in JSON and push it to the bucket with gsutil lifecycle set config-json-file. is the right answer.

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

NEW QUESTION 140
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 141
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 145
Your organization uses G Suite for communication and collaboration. All users in your organization have a G Suite account. You want to grant some G Suite users access to your Cloud Platform project. What should you do?

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

Answer: B

NEW QUESTION 146
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 149
Your company requires all developers to have the same permissions, regardless of the Google Cloud project they are working on. Your company's security policy also restricts developer permissions to Compute Engine.
Cloud Functions, and Cloud SQL. You want to implement the security policy with minimal effort. What should you do?

A. • Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions in one project within the Google Cloud organization.• Copy the role across all projects created within the organization with the gcloud iam roles copy command.• Assign the role to developers in those projects.
B. • Add all developers to a Google group in Google Groups for Workspace.• Assign the predefined role of Compute Admin to the Google group at the Google Cloud organization level.
C. • Add all developers to a Google group in Cloud Identity.• Assign predefined roles for Compute Engine, Cloud Functions, and Cloud SQL permissions to the Google group for each project in the Google Cloud organization.
D. • Add all developers to a Google group in Cloud Identity.• Create a custom role with Compute Engine, Cloud Functions, and Cloud SQL permissions at the Google Cloud organization level.• Assign the custom role to the Google group.

Answer: D

Explanation:
https://www.cloudskillsboost.google/focuses/1035?parent=catalog#:~:text=custom%20role%20at%20the%20or

NEW QUESTION 151
You need to track and verity modifications to a set of Google Compute Engine instances in your Google Cloud project. In particular, you want to verify OS system

patching events on your virtual machines (VMs). What should you do?

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

Answer: A

NEW QUESTION 156
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 157
Your company implemented BigQuery as an enterprise data warehouse. Users from multiple business units run queries on this data warehouse. However, you notice that query costs for BigQuery are very high, and you need to control costs. Which two methods should you use? (Choose two.)

A. Split the users from business units to multiple projects.
B. Apply a user- or project-level custom query quota for BigQuery data warehouse.
C. Create separate copies of your BigQuery data warehouse for each business unit.
D. Split your BigQuery data warehouse into multiple data warehouses for each business unit.
E. Change your BigQuery query model from on-demand to flat rat
F. Apply the appropriate number of slots to each Project.

Answer: BE

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

NEW QUESTION 161
You have production and test workloads that you want to deploy on Compute Engine. Production VMs need to be in a different subnet than the test VMs. All the VMs must be able to reach each other over internal IP without creating additional routes. You need to set up VPC and the 2 subnets. Which configuration meets these requirements?

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

Answer: A

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

NEW QUESTION 166
You want to verify the IAM users and roles assigned within a GCP project named my-project. What should you do?

A. Run gcloud iam roles lis
B. Review the output section.
C. Run gcloud iam service-accounts lis
D. Review the output section.
E. Navigate to the project and then to the IAM section in the GCP Consol
F. Review the members and roles.
G. Navigate to the project and then to the Roles section in the GCP Consol
H. Review the roles and status.

Answer: C

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

NEW QUESTION 169
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 171
You are working for a startup that was officially registered as a business 6 months ago. As your customer base grows, your use of Google Cloud increases. You want to allow all engineers to create new projects without asking them for their credit card information. What should you do?

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

Answer: A

NEW QUESTION 176
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 have a number of compute instances belonging to an unmanaged instances group. You need to SSH to one of the Compute Engine instances to run an ad hoc script. You've already authenticated gcloud, however, you don't have an SSH key deployed yet. In the fewest steps possible, what's the easiest way to SSH to the instance?

A. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.
B. Use the gcloud compute ssh command.
C. Create a key with the ssh-keygen comman
D. Then use the gcloud compute ssh command.
E. Create a key with the ssh-keygen comman
F. Upload the key to the instanc
G. Run gcloud compute instances list to get the IP address of the instance, then use the ssh command.

Answer: B

Explanation:
gcloud compute ssh ensures that the user's public SSH key is present in the project's metadata. If the user does not have a public SSH key, one is generated using ssh-keygen and added to the project's metadata. This is similar to the other option where we copy the key explicitly to the project's metadata but here it is done automatically for us. There are also security benefits with this approach. When we use gcloud compute ssh to connect to Linux instances, we are adding a layer of security by storing your host keys as guest attributes. Storing SSH host keys as guest attributes improve the security of your connections by helping to protect against vulnerabilities such as man-in-the-middle (MITM) attacks. On the initial boot of a VM instance, if guest attributes are enabled, Compute Engine stores your generated host keys as guest attributes.
Compute Engine then uses these host keys that were stored during the initial boot to verify all subsequent connections to the VM instance.
Ref: https://cloud.google.com/compute/docs/instances/connecting-to-instanceRef: https://cloud.google.com/s

NEW QUESTION 184
You are running a data warehouse on BigQuery. A partner company is offering a recommendation engine based on the data in your data warehouse. The partner company is also running their application on Google Cloud. They manage the resources in their own project, but they need access to the BigQuery dataset in your project. You want to provide the partner company with access to the dataset What should you do?

A. Create a Service Account in your own project, and grant this Service Account access to BigGuery in your project
B. Create a Service Account in your own project, and ask the partner to grant this Service Account access to BigQuery in their project
C. Ask the partner to create a Service Account in their project, and have them give the Service Account access to BigQuery in their project
D. Ask the partner to create a Service Account in their project, and grant their Service Account access to the BigQuery dataset in your project

Answer: D

Explanation:
https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0#:~:text=Go%20to%20t

NEW QUESTION 189
Your company has a Google Cloud Platform project that uses BigQuery for data warehousing. Your data science team changes frequently and has few members. You need to allow members of this team to perform queries. You want to follow Google-recommended practices. What should you do?

A. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery jobUser role to the group.
B. 1. Create an IAM entry for each data scientist's user account.2. Assign the BigQuery dataViewer user role to the group.
C. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery jobUser role to the group.
D. 1. Create a dedicated Google group in Cloud Identity.2. Add each data scientist's user account to the group.3. Assign the BigQuery dataViewer user role to the group.

Answer: C

Explanation:
Read the dataset's metadata and to list tables in the dataset. Read data and metadata from the dataset's tables. When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.
BigQuery Data Viewer (roles/bigquery.dataViewer)
When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.
This role cannot be applied to individual models or routines. When applied to a dataset, this role provides permissions to: Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables.
When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.
Lowest-level resources where you can grant this role: Table
View
BigQuery Job User (roles/bigquery.jobUser)
Provides permissions to run jobs, including queries, within the project.
Lowest-level resources where you can grant this role:
Project
to run jobs https://cloud.google.com/bigquery/docs/access-control#bigquery.jobUser databaseUser needs additional role permission to run jobs
https://cloud.google.com/spanner/docs/iam#spanner.databaseUser

NEW QUESTION 191
Your management has asked an external auditor to review all the resources in a specific project. The security team has enabled the Organization Policy called Domain Restricted Sharing on the organization node by specifying only your Cloud Identity domain. You want the auditor to only be able to view, but not modify, the resources in that project. What should you do?

A. Ask the auditor for their Google account, and give them the Viewer role on the project.
B. Ask the auditor for their Google account, and give them the Security Reviewer role on the project.
C. Create a temporary account for the auditor in Cloud Identity, and give that account the Viewer role on the project.
D. Create a temporary account for the auditor in Cloud Identity, and give that account the Security Reviewer role on the project.

Answer: C

Explanation:
Using primitive roles The following table lists the primitive roles that you can grant to access a project, the description of what the role does, and the permissions bundled within that role. Avoid using primitive roles except when absolutely necessary. These roles are very powerful, and include a large number of permissions across all Google Cloud services. For more details on when you should use primitive roles, see the Identity and Access Management FAQ. IAM predefined roles are much more granular, and allow you to carefully manage the set of permissions that your users have access to. See Understanding Roles for a list of roles that can be granted at the project level. Creating custom roles can further increase the control you have over user permissions. https://cloud.google.com/resource-manager/docs/access-control-proj#using_primitive_roles
https://cloud.google.com/iam/docs/understanding-custom-roles

NEW QUESTION 194
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 196
You need to set up a policy so that videos stored in a specific Cloud Storage Regional bucket are moved to Coldline after 90 days, and then deleted after one year from their creation. How should you set up the policy?

A. Use Cloud Storage Object Lifecycle Management using Age conditions with SetStorageClass and Delete action

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

Answer: A

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

NEW QUESTION 197
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 200
You are migrating a business critical application from your local data center into Google Cloud. As part of your high-availability strategy, you want to ensure that any data used by the application will be immediately available if a zonal failure occurs. What should you do?

A. Store the application data on a zonal persistent dis
B. Create a snapshot schedule for the dis
C. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
D. Store the application data on a zonal persistent dis
E. If an outage occurs, create an instance in another zone with this disk attached.
F. Store the application data on a regional persistent dis
G. Create a snapshot schedule for the dis
H. If an outage occurs, create a new disk from the most recent snapshot and attach it to a new VM in another zone.
I. Store the application data on a regional persistent disk If an outage occurs, create an instance in another zone with this disk attached.

Answer: A

NEW QUESTION 204
You host a static website on Cloud Storage. Recently, you began to include links to PDF files on this site. Currently, when users click on the links to these PDF files, their browsers prompt them to save the file onto their local system. Instead, you want the clicked PDF files to be displayed within the browser window directly, without prompting the user to save the file locally. What should you do?

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

Answer: C

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

NEW QUESTION 208
You want to set up a Google Kubernetes Engine cluster Verifiable node identity and integrity are required for the cluster, and nodes cannot be accessed from the internet. You want to reduce the operational cost of managing your cluster, and you want to follow Google-recommended practices. What should you do?

A. Deploy a private autopilot cluster
B. Deploy a public autopilot cluster.
C. Deploy a standard public cluster and enable shielded nodes.
D. Deploy a standard private cluster and enable shielded nodes.

Answer: D

NEW QUESTION 209
Your organization has three existing Google Cloud projects. You need to bill the Marketing department for only their Google Cloud services for a new initiative within their group. What should you do?

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

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

C. * 1. Verify that you are assigned the Organization Administrator IAM role for your organization's Google Cloud account* 2. Create a new Google Cloud Project for the Marketing department 3. Link the new project to a Marketing Billing Account.

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

Answer: A

NEW QUESTION 212
......

Relate Links

100% Pass Your Associate-Cloud-Engineer Exam with Exambible Prep Materials

https://www.exambible.com/Associate-Cloud-Engineer-exam/

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - https://www.exambible.com/