

Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer



NEW QUESTION 1

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 2

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

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

Answer: C

Explanation:

When you initially click on Monitoring(Stackdriver Monitoring) it creates a workspace(a stackdriver account) linked to the ACTIVE(CURRENT) Project from which it was clicked.

Now if you change the project and again click onto Monitoring it would create an another workspace(a stackdriver account) linked to the changed ACTIVE(CURRENT) Project, we don't want this as this would not consolidate our result into a single dashboard(workspace/stackdriver account).

If you have accidentally created two diff workspaces merge them under Monitoring > Settings > Merge Workspaces > MERGE.

If we have only one workspace and two projects we can simply add other GCP Project under Monitoring > Settings > GCP Projects > Add GCP Projects.

<https://cloud.google.com/monitoring/settings/multiple-projects>

Nothing about groups <https://cloud.google.com/monitoring/settings?hl=en>

NEW QUESTION 3

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

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

Answer: D

Explanation:

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

NEW QUESTION 4

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 5

Your company has a large quantity of unstructured data in different file formats. You want to perform ETL transformations on the data. You need to make the data accessible on Google Cloud so it can be processed by a Dataflow job. What should you do?

- A. Upload the data to BigQuery using the bq command line tool.
- B. Upload the data to Cloud Storage using the gsutil command line tool.
- C. Upload the data into Cloud SQL using the import function in the console.
- D. Upload the data into Cloud Spanner using the import function in the console.

Answer: B

Explanation:

"large quantity" : Cloud Storage or BigQuery "files" a file is nothing but an Object

NEW QUESTION 6

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

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

Answer: C

Explanation:

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0>

You create the service account in proj-sa and take note of the service account email, then you go to proj-vm in IAM > ADD and add the service account's email as new member and give it the Compute Storage Admin role.

<https://cloud.google.com/compute/docs/access/iam#compute.storageAdmin>

NEW QUESTION 7

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 8

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 instance
- C. Run the Jenkins executable.
- D. Create a new Kubernetes Engine cluster
- E. Create a deployment for the Jenkins image.
- F. Create an instance template with the Jenkins executable
- 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 9

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

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

Answer: C

Explanation:

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

NEW QUESTION 10

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 10

Your team is using Linux instances on Google Cloud. You need to ensure that your team logs in to these instances in the most secure and cost efficient way. What should you do?

- A. Attach a public IP to the instances and allow incoming connections from the internet on port 22 for SSH.
- B. Use a third party tool to provide remote access to the instances.
- C. Use the gcloud compute ssh command with the --tunnel-through-iap flag
- D. Allow ingress traffic from the IP range 35.235.240.0/20 on port 22.
- E. Create a bastion host with public internet access
- F. Create the SSH tunnel to the instance through the bastion host.

Answer: D

NEW QUESTION 11

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 tab
- D. Create a Looker Studio dashboard on top of the CSV export.
- E. Configure Cloud Billing data export to BigQuery for the billing account
- F. Create a Looker Studio dashboard on top of the BigQuery export.

Answer: D

NEW QUESTION 13

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

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

Answer: D

Explanation:

IAP controls access to your App Engine apps and Compute Engine VMs running on Google Cloud. It leverages user identity and the context of a request to determine if a user should be allowed access. IAP is a building block toward BeyondCorp, an enterprise security model that enables employees to work from untrusted networks without using a VPN.

By default, IAP uses Google identities and IAM. By leveraging Identity Platform instead, you can authenticate users with a wide range of external identity providers, such as:

Email/password

OAuth (Google, Facebook, Twitter, GitHub, Microsoft, etc.) SAML

OIDC

Phone number Custom Anonymous

This is useful if your application is already using an external authentication system, and migrating your users to Google accounts is impractical.

<https://cloud.google.com/iap/docs/using-tcp-forwarding#grant-permission>

NEW QUESTION 17

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

- A. Enable the Cloud Pub/Sub API in the API Library on the GCP Console.
- B. Rely on the automatic enablement of the Cloud Pub/Sub API when the Service Account accesses it.
- C. Use Deployment Manager to deploy your application
- 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 Admin
- 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 18

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

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

Answer: A

Explanation:

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

NEW QUESTION 21

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

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

Answer: B

Explanation:

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

<https://cloud.google.com/resource-manager/docs/creating-managing-folders>

NEW QUESTION 23

Your company is moving its continuous integration and delivery (CI/CD) pipeline to Compute Engine instances. The pipeline will manage the entire cloud infrastructure through code. How can you ensure that the pipeline has appropriate permissions while your system is following security best practices?

- A. • Add a step for human approval to the CI/CD pipeline before the execution of the infrastructure provisioning. • Use the human approvals IAM account for the provisioning.
- B. • Attach a single service account to the compute instances. • Add minimal rights to the service account. • Allow the service account to impersonate a Cloud Identity user with elevated permissions to create, update, or delete resources.
- C. • Attach a single service account to the compute instances. • Add all required Identity and Access Management (IAM) permissions to this service account to create, update, or delete resources
- D. • Create multiple service accounts, one for each pipeline with the appropriate minimal Identity and Access Management (IAM) permissions. • Use a secret manager service to store the key files of the service accounts. • Allow the CI/CD pipeline to request the appropriate secrets during the execution of the pipeline.

Answer: B

Explanation:

The best option is to attach a single service account to the compute instances and add minimal rights to the service account. Then, allow the service account to impersonate a Cloud Identity user with elevated

permissions to create, update, or delete resources. This way, the service account can use short-lived access tokens to authenticate to Google Cloud APIs without needing to manage service account keys. This option follows the principle of least privilege and reduces the risk of credential leakage and misuse.

Option A is not recommended because it requires human intervention, which can slow down the CI/CD pipeline and introduce human errors. Option C is not secure because it grants all required IAM permissions to a single service account, which can increase the impact of a compromised key. Option D is not cost-effective because it requires creating and managing multiple service accounts and keys, as well as using a secret manager service.

References:

- 1: <https://cloud.google.com/iam/docs/impersonating-service-accounts>
- 2: <https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys>
- 3: <https://cloud.google.com/iam/docs/understanding-service-accounts>

NEW QUESTION 26

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 27

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

- A. Stream data to Pub/Sub, and use Dataflow to send data to Cloud Storage
- B. Stream data to Pub/Su
- C. and use Storage Transfer Service to send data to BigQuery.
- D. Stream data to Dataflow, and use Storage Transfer Service to send data to BigQuery.
- E. Stream data to Dataflow, and use Dataprep by Trifacta to send data to Bigtable.

Answer: B

NEW QUESTION 29

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 33

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

- A. Create a signed URL with a four-hour expiration and share the URL with the company.
- B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.
- C. Configure the storage bucket as a static website and furnish the object's URL to the compan
- D. Delete the object from the storage bucket after four hours.
- E. Create a new Cloud Storage bucket specifically for the external company to acces
- F. Copy the object to that bucke
- G. Delete the bucket after four hours have passed.

Answer: A

Explanation:

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

NEW QUESTION 36

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

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

Answer: D

Explanation:

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

NEW QUESTION 37

A team of data scientists infrequently needs to use a Google Kubernetes Engine (GKE) cluster that you manage. They require GPUs for some long-running, non-restartable jobs. You want to minimize cost. What should you do?

- A. Enable node auto-provisioning on the GKE cluster.
- B. Create a VerticalPodAutscaler for those workloads.
- C. Create a node pool with preemptible VMs and GPUs attached to those VMs.
- D. Create a node pool of instances with GPUs, and enable autoscaling on this node pool with a minimum size of 1.

Answer: A

Explanation:

auto-provisioning = Attaches and deletes node pools to cluster based on the requirements. Hence creating a GPU node pool, and auto-scaling would be better
<https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-provisioning>

NEW QUESTION 40

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 43

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

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

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

Answer: D

NEW QUESTION 50

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

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

Answer: A

Explanation:

VMs cannot communicate over the internet without a public IP address. Private Google Access permits access to Google APIs and services in Google's production infrastructure.

<https://cloud.google.com/vpc/docs/private-google-access>

NEW QUESTION 54

You are running multiple VPC-native Google Kubernetes Engine clusters in the same subnet. The IPs available for the nodes are exhausted, and you want to ensure that the clusters can grow in nodes when needed. What should you do?

- A. Create a new subnet in the same region as the subnet being used.
- B. Add an alias IP range to the subnet used by the GKE clusters.
- C. Create a new VPC, and set up VPC peering with the existing VPC.
- D. Expand the CIDR range of the relevant subnet for the cluster.

Answer: D

Explanation:

gcloud compute networks subnets expand-ip-range NAME gcloud compute networks subnets expand-ip-range

- expand the IP range of a Compute Engine subnetwork <https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

NEW QUESTION 58

You have been asked to create robust Virtual Private Network (VPN) connectivity between a new Virtual Private Cloud (VPC) and a remote site. Key requirements include dynamic routing, a shared address space of 10.19.0.1/22, and no overprovisioning of tunnels during a failover event. You want to follow Google-recommended practices to set up a high availability Cloud VPN. What should you do?

- A. Use a custom mode VPC network, configure static routes, and use active/passive routing
- B. Use an automatic mode VPC network, configure static routes, and use active/active routing
- C. Use a custom mode VPC network use Cloud Router border gateway protocol (86P) routes, and use active/passive routing
- D. Use an automatic mode VPC network, use Cloud Router border gateway protocol (BGP) routes and configure policy-based routing

Answer: C

Explanation:

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/best-practices>

NEW QUESTION 59

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

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

Answer: C

Explanation:

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

NEW QUESTION 64

You need to run an important query in BigQuery but expect it to return a lot of records. You want to find out how much it will cost to run the query. You are using on-demand pricing. What should you do?

- A. Arrange to switch to Flat-Rate pricing for this query, then move back to on-demand.
- B. Use the command line to run a dry run query to estimate the number of bytes read
- C. Then convert that bytes estimate to dollars using the Pricing Calculator.
- D. Use the command line to run a dry run query to estimate the number of bytes returned
- E. Then convert that bytes estimate to dollars using the Pricing Calculator.
- F. Run a select count (*) to get an idea of how many records your query will look through
- G. Then convert that number of rows to dollars using the Pricing Calculator.

Answer: B

NEW QUESTION 69

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

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

Answer: B

NEW QUESTION 70

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

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

Answer: D

Explanation:

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

NEW QUESTION 75

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

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

Answer: D

NEW QUESTION 78

You create a new Google Kubernetes Engine (GKE) cluster and want to make sure that it always runs a supported and stable version of Kubernetes. What should you do?

- A. Enable the Node Auto-Repair feature for your GKE cluster.
- B. Enable the Node Auto-Upgrades feature for your GKE cluster.
- C. Select the latest available cluster version for your GKE cluster.
- D. Select "Container-Optimized OS (cos)" as a node image for your GKE cluster.

Answer: B

Explanation:

Creating or upgrading a cluster by specifying the version as latest does not provide automatic upgrades. Enable node auto-upgrades to ensure that the nodes in your cluster are up-to-date with the latest stable version.

<https://cloud.google.com/kubernetes-engine/versioning-and-upgrades>

Node auto-upgrades help you keep the nodes in your cluster up to date with the cluster master version when your master is updated on your behalf. When you create a new cluster or node pool with Google Cloud Console or the gcloud command, node auto-upgrade is enabled by default.

Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades>

NEW QUESTION 80

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

- A. After the VM has been created, use your Google Account credentials to log in into the VM.
- B. After the VM has been created, use gcloud compute reset-windows-password to retrieve the login credentials for the VM.
- C. When creating the VM, add metadata to the instance using 'windows-password' as the key and a password as the value.
- D. After the VM has been created, download the JSON private key for the default Compute Engine service account
- 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 82

You are building an application that processes data files uploaded from thousands of suppliers. Your primary goals for the application are data security and the expiration of aged data. You need to design the application to:

- Restrict access so that suppliers can access only their own data.
- Give suppliers write access to data only for 30 minutes.
- Delete data that is over 45 days old.

You have a very short development cycle, and you need to make sure that the application requires minimal maintenance. Which two strategies should you use? (Choose two.)

- A. Build a lifecycle policy to delete Cloud Storage objects after 45 days.
- B. Use signed URLs to allow suppliers limited time access to store their objects.
- C. Set up an SFTP server for your application, and create a separate user for each supplier.
- D. Build a Cloud function that triggers a timer of 45 days to delete objects that have expired.
- E. Develop a script that loops through all Cloud Storage buckets and deletes any buckets that are older than 45 days.

Answer: AB

Explanation:

(A) Object Lifecycle Management Delete

The Delete action deletes an object when the object meets all conditions specified in the lifecycle rule.

Exception: In buckets with Object Versioning enabled, deleting the live version of an object causes it to become a noncurrent version, while deleting a noncurrent version deletes that version permanently.

<https://cloud.google.com/storage/docs/lifecycle#delete>

(B) Signed URLs

This page provides an overview of signed URLs, which you use to give time-limited resource access to anyone in possession of the URL, regardless of whether they have a Google account

<https://cloud.google.com/storage/docs/access-control/signed-urls>

NEW QUESTION 86

Your application development team has created Docker images for an application that will be deployed on Google Cloud. Your team does not want to manage the infrastructure associated with this application. You need to ensure that the application can scale automatically as it gains popularity. What should you do?

- A. Create an Instance template with the container image, and deploy a Managed Instance Group with Autoscaling.

- B. Upload Docker images to Artifact Registry, and deploy the application on Google Kubernetes Engine using Standard mode.
- C. Upload Docker images to the Cloud Storage, and deploy the application on Google Kubernetes Engine using Standard mode.
- D. Upload Docker images to Artifact Registry, and deploy the application on Cloud Run.

Answer: D

NEW QUESTION 88

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

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

Answer: A

Explanation:

https://cloud.google.com/iam/docs/job-functions/auditing#scenario_external_auditors

Because if you directly add users to the IAM roles, then if any users left the organization then you have to remove the users from multiple places and need to revoke his/her access from multiple places. But, if you put a user into a group then its very easy to manage these type of situations. Now, if any user left then you just need to remove the user from the group and all the access got revoked

The organization creates a Google group for these external auditors and adds the current auditor to the group. This group is monitored and is typically granted access to the dashboard application. During normal access, the auditors' Google group is only granted access to view the historic logs stored in BigQuery. If any anomalies are discovered, the group is granted permission to view the actual Cloud Logging Admin Activity logs via the dashboard's elevated access mode. At the end of each audit period, the group's access is then revoked. Data is redacted using Cloud DLP before being made accessible for viewing via the dashboard application. The table below explains IAM logging roles that an Organization Administrator can grant to the service account used by the dashboard, as well as the resource level at which the role is granted.

NEW QUESTION 92

You have downloaded and installed the gcloud command line interface (CLI) and have authenticated with your Google Account. Most of your Compute Engine instances in your project run in the europe-west1-d zone. You want to avoid having to specify this zone with each CLI command when managing these instances. What should you do?

- A. Set the europe-west1-d zone as the default zone using the gcloud config subcommand.
- B. In the Settings page for Compute Engine under Default location, set the zone to europe-west1-d.
- C. In the CLI installation directory, create a file called default.conf containing zone=europe-west1-d.
- D. Create a Metadata entry on the Compute Engine page with key compute/zone and value europe-west1-d.

Answer: A

Explanation:

Change your default zone and region in the metadata server Note: This only applies to the default configuration. You can change the default zone and region in your metadata server by making a request to the metadata server. For example: `gcloud compute project-info add-metadata --metadata`

`google-compute-default-region=europe-west1,google-compute-default-zone=europe-west1-b` The gcloud command-line tool only picks up on new default zone and region changes after you rerun the gcloud init command. After updating your default metadata, run gcloud init to reinitialize your default configuration.

https://cloud.google.com/compute/docs/gcloud-compute#change_your_default_zone_and_region_in_the_metad

NEW QUESTION 97

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 101

You need to create a custom IAM role for use with a GCP service. All permissions in the role must be suitable for production use. You also want to clearly share with your organization the status of the custom role. This will be the first version of the custom role. What should you do?

- A. Use permissions in your role that use the 'supported' support level for role permission
- B. Set the rolestage to ALPHA while testing the role permissions.
- C. Use permissions in your role that use the 'supported' support level for role permission
- D. Set the role stage to BETA while testing the role permissions.
- E. Use permissions in your role that use the 'testing' support level for role permission
- F. Set the role stage to ALPHA while testing the role permissions.
- G. Use permissions in your role that use the 'testing' support level for role permission
- H. Set the role stage to BETA while testing the role permissions.

Answer: A

Explanation:

When setting support levels for permissions in custom roles, you can set to one of SUPPORTED, TESTING or NOT_SUPPORTED.

Ref: <https://cloud.google.com/iam/docs/custom-roles-permissions-support>

NEW QUESTION 106

You have an application that runs on Compute Engine VM instances in a custom Virtual Private Cloud (VPC). Your company's security policies only allow the use to internal IP addresses on VM instances and do not let VM instances connect to the internet. You need to ensure that the application can access a file hosted in a Cloud Storage bucket within your project. What should you do?

- A. Enable Private Service Access on the Cloud Storage Bucket.
- B. Add storage.googleapis.com to the list of restricted services in a VPC Service Controls perimeter and add your project to the list to protected projects.
- C. Enable Private Google Access on the subnet within the custom VPC.
- D. Deploy a Cloud NAT instance and route the traffic to the dedicated IP address of the Cloud Storage bucket.

Answer: A

NEW QUESTION 108

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

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

Answer: C

Explanation:

<https://cloud.google.com/run/docs/tutorials/pubsub#integrating-pubsub>

* 1. Create a service account. 2. Give the Cloud Run Invoker role to that service account for your Cloud Run application. 3. Create a Cloud Pub/Sub subscription that uses that service account and uses your Cloud Run application as the push endpoint.

NEW QUESTION 111

Your customer has implemented a solution that uses Cloud Spanner and notices some read latency-related performance issues on one table. This table is accessed only by their users using a primary key. The table schema is shown below.

```
CREATE TABLE Persons (  
    person_id INT64 NOT NULL,    // sequential number based on number of registration  
    account_creation_date DATE, // system date  
    birthdate DATE,             // customer birthdate  
    firstname STRING (255),      // first name  
    lastname STRING (255),       // last name  
    profile_picture BYTES (255)  // profile picture  
) PRIMARY KEY (person_id)
```

You want to resolve the issue. What should you do?

- A. Remove the profile_picture field from the table.
- B. Add a secondary index on the person_id column.
- C. Change the primary key to not have monotonically increasing values.
- D. Create a secondary index using the following Data Definition Language (DDL):

```
CREATE INDEX person_id_ix
ON Persons (
    person_id,
    firstname,
    lastname
) STORING (
    profile_picture
)
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

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

NEW QUESTION 115

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 116

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

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

Answer: B

Explanation:

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

NEW QUESTION 119

You need to enable traffic between multiple groups of Compute Engine instances that are currently running two different GCP projects. Each group of Compute Engine instances is running in its own VPC. What should you do?

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

Answer: B

Explanation:

Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network. When you use Shared VPC, you designate a project as a host project and attach one

or more other service projects to it. The VPC networks in the host project are called Shared VPC networks. Eligible resources from service projects can use subnets in the Shared VPC network

<https://cloud.google.com/vpc/docs/shared-vpc>

"For example, an existing instance in a service project cannot be reconfigured to use a Shared VPC network, but a new instance can be created to use available subnets in a Shared VPC network."

NEW QUESTION 120

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 to 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 organization
- 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 for the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalog
- G. Share the Catalog with your organization, and deploy the resources in your company's project.
- H. Create an infrastructure-as-code template for all resources in the project by using Terraform
- I. and deploy that template to a new project in your organization
- J. Delete the project from the startup's Google Cloud organization.

Answer: A

NEW QUESTION 124

You want to run a single caching HTTP reverse proxy on GCP for a latency-sensitive website. This specific reverse proxy consumes almost no CPU. You want to have a 30-GB in-memory cache, and need an additional 2 GB of memory for the rest of the processes. You want to minimize cost. How should you run this reverse proxy?

- A. Create a Cloud Memorystore for Redis instance with 32-GB capacity.
- B. Run it on Compute Engine, and choose a custom instance type with 6 vCPUs and 32 GB of memory.
- C. Package it in a container image, and run it on Kubernetes Engine, using n1-standard-32 instances as nodes.
- D. Run it on Compute Engine, choose the instance type n1-standard-1, and add an SSD persistent disk of 32 GB.

Answer: A

Explanation:

What is Google Cloud Memorystore?

Overview. Cloud Memorystore for Redis is a fully managed Redis service for Google Cloud Platform. Applications running on Google Cloud Platform can achieve extreme performance by leveraging the highly scalable, highly available, and secure Redis service without the burden of managing complex Redis deployments.

NEW QUESTION 129

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

- A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
- B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.
- C. Deploy the container on App Engine flexible environment with autoscaling
- 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 131

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 cluster.

When you run the kubectl get nodes command, you only see the nodes from prod-cluster. Which commands should you run to check the node status for dev-cluster?

- A.

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

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

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

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

Answer: C

NEW QUESTION 132

You have an application running in Google Kubernetes Engine (GKE) with cluster autoscaling enabled. The application exposes a TCP endpoint. There are several replicas of this application. You have a Compute Engine instance in the same region, but in another Virtual Private Cloud (VPC), called gce-network, that has no overlapping IP ranges with the first VPC. This instance needs to connect to the application on GKE. You want to minimize effort. What should you do?

- A. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Set the service's externalTrafficPolicy to Cluster.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- B. 1. In GKE, create a Service of type NodePort that uses the application's Pods as backend.2. Create a Compute Engine instance called proxy with 2 network interfaces, one in each VPC.3. Use iptables on this instance to forward traffic from gce-network to the GKE nodes.4. Configure the Compute Engine instance to use the address of proxy in gce-network as endpoint.
- C. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add an annotation to this service: cloud.google.com/load-balancer-type: Internal3. Peer the two VPCs together.4. Configure the Compute Engine instance to use the address of the load balancer that has been created.
- D. 1. In GKE, create a Service of type LoadBalancer that uses the application's Pods as backend.2. Add a Cloud Armor Security Policy to the load balancer that whitelists the internal IPs of the MIG's instances.3. Configure the Compute Engine instance to use the address of the load balancer that has been created.

Answer: A

Explanation:

performs a peering between the two VPC's (the statement makes sure that this option is feasible since it clearly specifies that there is no overlapping between the ip ranges of both vpc's), deploy the LoadBalancer as internal with the annotation, and configure the endpoint so that the compute engine instance can access the application internally, that is, without the need to have a public ip at any time and therefore, without the need to go outside the google network. The traffic, therefore, never crosses the public internet.

<https://medium.com/pablo-perez/k8s-externaltrafficpolicy-local-or-cluster-40b259a19404> <https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing>

clients in a VPC network connected to the LoadBalancer network using VPC Network Peering can also access the Service

<https://cloud.google.com/kubernetes-engine/docs/how-to/service-parameters>

NEW QUESTION 137

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 142

You want to configure an SSH connection to a single Compute Engine instance for users in the dev1 group. This instance is the only resource in this particular Google Cloud Platform project that the dev1 users should be able to connect to. What should you do?

- A. Set metadata to enable-oslogin=true for the instanc
- B. Grant the dev1 group the compute.osLogin role. Direct them to use the Cloud Shell to ssh to that instance.
- C. Set metadata to enable-oslogin=true for the instanc
- D. Set the service account to no service account for that instanc
- E. Direct them to use the Cloud Shell to ssh to that instance.
- F. Enable block project wide keys for the instanc
- G. Generate an SSH key for each user in the dev1 group. Distribute the keys to dev1 users and direct them to use their third-party tools to connect.
- H. Enable block project wide keys for the instanc
- I. Generate an SSH key and associate the key with that instanc
- J. Distribute the key to dev1 users and direct them to use their third-party tools to connect.

Answer: A

NEW QUESTION 145

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

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

Answer: A

Explanation:

Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Since we have a requirement to access data once a quarter and want to go with the most cost-efficient option, we should select Coldline Storage.

Ref: <https://cloud.google.com/storage/docs/storage-classes#coldline>

Google Cloud Storage Classes in the Organization

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

Storage Class	Characteristics	Use Cases	Price (Per Gb Per Month)*
Multi-Regional Storage	<ul style="list-style-type: none"> 99.95% availability Geo-redundant 	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	\$0.026 per GB/Month
Regional Storage	<ul style="list-style-type: none"> 99.9% availability Low cost per GB stored Data storage in a small region 	Keeps information that is frequently accessed around the globe, such as videos, gaming, and mobile applications	\$0.02 per GB/Month
Storage Nearline	<ul style="list-style-type: none"> 99.0% availability Very low cost per GB Data fetching costs Higher per-task costs 30-day minimum storage duration 	Keeps data that is not accessed is often ideal for data backups	\$0.01 per GB/Month
Storage Cold line	<ul style="list-style-type: none"> 99.0% availability Lowest cost per GB Data fetching costs Higher per-task costs 90-day minimum storage duration 	Keeps information that is infrequently ideal for disaster recovery or archived data	\$0.007 per GB/Month

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

NEW QUESTION 146

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?

- Ask the other team to grant your default App Engine Service account the role of BigQuery Job User.
- Ask the other team to grant your default App Engine Service account the role of BigQuery Data Viewer.
- In Cloud IAM of your project, ensure that the default App Engine service account has the role of BigQuery Data Viewer.
- 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 148

You recently received a new Google Cloud project with an attached billing account where you will work. You need to create instances, set firewalls, and store data in Cloud Storage. You want to follow Google-recommended practices. What should you do?

- Use the gcloud CLI services enable cloudresourcemanager.googleapis.com command to enable all resources.
- Use the gcloud services enable compute.googleapis.com command to enable Compute Engine and the gcloud services enable storage-api.googleapis.com command to enable the Cloud Storage APIs.
- Open the Google Cloud console and enable all Google Cloud APIs from the API dashboard.
- Open the Google Cloud console and run gcloud init --project <project-id> in a Cloud Shell.

Answer: B

NEW QUESTION 151

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

- Deploy the container on Cloud Run.
- Deploy the container on Cloud Run on GKE.

- C. Deploy the container on App Engine Flexible.
- D. Deploy the container on Google Kubernetes Engine, with cluster autoscaling and horizontal pod autoscaling enabled.

Answer: A

Explanation:

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

<https://cloud.google.com/run>

NEW QUESTION 153

You have two subnets (subnet-a and subnet-b) in the default VPC. Your database servers are running in subnet-a. Your application servers and web servers are running in subnet-b. You want to configure a firewall rule that only allows database traffic from the application servers to the database servers. What should you do?

- A. * Create service accounts sa-app and sa-db. • Associate service account: sa-app with the application servers and the service account sa-db with the database servers. • Create an ingress firewall rule to allow network traffic from source service account sa-app to target service account sa-db.
- B. • Create network tags app-server and db-server. • Add the app-server tag to the application servers and the db-server tag to the database servers. • Create an egress firewall rule to allow network traffic from source network tag app-server to target network tag db-server.
- C. * Create a service account sa-app and a network tag db-server. * Associate the service account sa-app with the application servers and the network tag db-server with the database servers. • Create an ingress firewall rule to allow network traffic from source VPC IP addresses and target the subnet-a IP addresses.
- D. • Create a network tag app-server and service account sa-db. • Add the tag to the application servers and associate the service account with the database servers. • Create an egress firewall rule to allow network traffic from source network tag app-server to target service account sa-db.

Answer: C

NEW QUESTION 156

You just installed the Google Cloud CLI on your new corporate laptop. You need to list the existing instances of your company on Google Cloud. What must you do before you run the gcloud compute instances list command?

Choose 2 answers

- A. Run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI.
- B. Create a Google Cloud service account, and download the service account key.
- C. Place the key file in a folder on your machine where gcloud CLI can find it.
- D. Download your Cloud Identity user account key.
- E. Place the key file in a folder on your machine where gcloud CLI can find it.
- F. Run gcloud config set compute/zone \$my_zone to set the default zone for gcloud CLI.
- G. Run gcloud config set project \$my_project to set the default project for gcloud CLI.

Answer: AE

Explanation:

Before you run the gcloud compute instances list command, you need to do two things: authenticate with your user account and set the default project for gcloud CLI.

To authenticate with your user account, you need to run gcloud auth login, enter your login credentials in the dialog window, and paste the received login token to gcloud CLI. This will authorize the gcloud CLI to access Google Cloud resources on your behalf¹.

To set the default project for gcloud CLI, you need to run gcloud config set project \$my_project, where

\$my_project is the ID of the project that contains the instances you want to list. This will save you from having to specify the project flag for every gcloud command².

Option B is not recommended, because using a service account key increases the risk of credential leakage and misuse. It is also not necessary, because you can use your user account to authenticate to the gcloud CLI³. Option C is not correct, because there is no such thing as a Cloud Identity user account key. Cloud Identity is a service that provides identity and access management for Google Cloud users and groups⁴. Option D is not required, because the gcloud compute instances list command does not depend on the default zone. You can list instances from all zones or filter by a specific zone using the --filter flag.

References:

- 1: <https://cloud.google.com/sdk/docs/authorizing>
- 2: <https://cloud.google.com/sdk/gcloud/reference/config/set>
- 3: <https://cloud.google.com/iam/docs/best-practices-for-managing-service-account-keys>
- 4: <https://cloud.google.com/identity/docs/overview>
- : <https://cloud.google.com/sdk/gcloud/reference/compute/instances/list>

NEW QUESTION 157

You significantly changed a complex Deployment Manager template and want to confirm that the dependencies of all defined resources are properly met before committing it to the project. You want the most rapid feedback on your changes. What should you do?

- A. Use granular logging statements within a Deployment Manager template authored in Python.
- B. Monitor activity of the Deployment Manager execution on the Stackdriver Logging page of the GCP Console.
- C. Execute the Deployment Manager template against a separate project with the same configuration, and monitor for failures.
- D. Execute the Deployment Manager template using the --preview option in the same project, and observe the state of interdependent resources.

Answer: D

NEW QUESTION 160

You need to grant access for three users so that they can view and edit table data on a Cloud Spanner instance. What should you do?

- A. Run gcloud iam roles describe roles/spanner.databaseUser

- B. Add the users to the role.
- C. Run `gcloud iam roles describe roles/spanner.databaseUser`
- D. Add the users to a new group
- E. Add the group to the role.
- F. Run `gcloud iam roles describe roles/spanner.viewer --project my-project`
- G. Add the users to the role.
- H. Run `gcloud iam roles describe roles/spanner.viewer --project my-project`
- I. Add the users to a new group. Add the group to the role.

Answer: B

Explanation:

<https://cloud.google.com/spanner/docs/iam#spanner.databaseUser>

Using the `gcloud` tool, execute the `gcloud iam roles describe roles/spanner.databaseUser` command on Cloud Shell. Attach the users to a newly created Google group and add the group to the role.

NEW QUESTION 163

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

- A. Navigate to Stackdriver Logging and select `resource.labels.project_id=""`
- B. Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset
- C. Configure the table expiration to 60 days.
- D. Create a Stackdriver Logging Export with a Sink destination to Cloud Storage
- 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 BigQuery
- 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 organization, 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 organization, 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 168

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

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

Answer: D

Explanation:

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

NEW QUESTION 173

You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images. The microservice responsible for the image rendering requires a large amount of CPU time compared to the memory it requires. The other microservices are workloads that are optimized for n1-standard machine types. You need to optimize your cluster so that all workloads are using resources as efficiently as possible. What should you do?

- A. Assign the pods of the image rendering microservice a higher pod priority than the older microservices
- B. Create a node pool with compute-optimized machine type nodes for the image rendering microservice Use the node pool with general-purpose machine type nodes for the other microservices
- C. Use the node pool with general-purpose machine type nodes for lite mage rendering microservice Create a nodepool with compute-optimized machine type nodes for the other microservices
- D. Configure the required amount of CPU and memory in the resource requests specification of the imagerendering microservice deployment Keep the resource requests for the other microservices at the default

Answer: B

NEW QUESTION 177

You are deploying a production application on Compute Engine. You want to prevent anyone from accidentally destroying the instance by clicking the wrong button. What should you do?

- A. Disable the flag "Delete boot disk when instance is deleted."
- B. Enable delete protection on the instance.
- C. Disable Automatic restart on the instance.
- D. Enable Preemptibility on the instance.

Answer: D

Explanation:

Preventing Accidental VM Deletion This document describes how to protect specific VM instances from deletion by setting the deletionProtection property on an Instance resource. To learn more about VM instances, read the Instances documentation. As part of your workload, there might be certain VM instances that are critical to running your application or services, such as an instance running a SQL server, a server used as a license manager, and so on. These VM instances might need to stay running indefinitely so you need a way to protect these VMs from being deleted. By setting the deletionProtection flag, a VM instance can be protected from accidental deletion. If a user attempts to delete a VM instance for which you have set the deletionProtection flag, the request fails. Only a user that has been granted a role with compute.instances.create permission can reset the flag to allow the resource to be deleted.
<https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

NEW QUESTION 181

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

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

Answer: D

Explanation:

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

NEW QUESTION 185

You need to deploy an application in Google Cloud using savorless technology. You want to test a new version of the application with a small percentage of production traffic. What should you do?

- A. Deploy the application lo Clou
- B. Ru
- C. Use gradual rollouts for traffic spelling.
- D. Deploy the application lo Google Kubemetes Engin
- E. Use Anthos Service Mesh for traffic splitting.
- F. Deploy the application to Cloud function
- G. Saucily the version number in the functions name.
- H. Deploy the application to App Engin
- I. For each new version, create a new service.

Answer: A

NEW QUESTION 190

Your team maintains the infrastructure for your organization. The current infrastructure requires changes. You need to share your proposed changes with the rest of the team. You want to follow Google's recommended best practices. What should you do?

- A. Use Deployment Manager templates to describe the proposed changes and store them in a Cloud Storage bucket.
- B. Use Deployment Manager templates to describe the proposed changes and store them in Cloud Source Repositories.
- C. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in a shared Storage bucket.
- D. Apply the change in a development environment, run `gcloud compute instances list`, and then save the output in Cloud Source Repositories.

Answer: B

Explanation:

Showing Deployment Manager templates to your team will allow you to define the changes you want to implement in your cloud infrastructure. You can use Cloud Source Repositories to store Deployment Manager templates and collaborate with your team. Cloud Source Repositories are fully-featured, scalable, and private Git repositories you can use to store, manage and track changes to your code.

<https://cloud.google.com/source-repositories/docs/features>

NEW QUESTION 195

You have an application that uses Cloud Spanner as a database backend to keep current state information about users. Cloud Bigtable logs all events triggered by users. You export Cloud Spanner data to Cloud Storage during daily backups. One of your analysts asks you to join data from Cloud Spanner and Cloud Bigtable for specific users. You want to complete this ad hoc request as efficiently as possible. What should you do?

- A. Create a dataflow job that copies data from Cloud Bigtable and Cloud Storage for specific users.
- B. Create a dataflow job that copies data from Cloud Bigtable and Cloud Spanner for specific users.
- C. Create a Cloud Dataproc cluster that runs a Spark job to extract data from Cloud Bigtable and Cloud Storage for specific users.
- D. Create two separate BigQuery external tables on Cloud Storage and Cloud Bigtable.
- E. Use the BigQuery console to join these tables through user fields, and apply appropriate filters.

Answer: D

Explanation:

"The Cloud Spanner to Cloud Storage Text template is a batch pipeline that reads in data from a Cloud Spanner table, optionally transforms the data via a JavaScript User Defined Function (UDF) that you provide, and writes it to Cloud Storage as CSV text files."

<https://cloud.google.com/dataflow/docs/guides/templates/provided-batch#cloudspannertogcstext>

"The Dataflow connector for Cloud Spanner lets you read data from and write data to Cloud Spanner in a Dataflow pipeline"

<https://cloud.google.com/spanner/docs/dataflow-connector> <https://cloud.google.com/bigquery/external-data-sources>

NEW QUESTION 196

You have a Linux VM that must connect to Cloud SQL. You created a service account with the appropriate access rights. You want to make sure that the VM uses this service account instead of the default Compute Engine service account. What should you do?

- A. When creating the VM via the web console, specify the service account under the 'Identity and API Access' section.
- B. Download a JSON Private Key for the service account.
- C. On the Project Metadata, add that JSON as the value for the key `compute-engine-service-account`.
- D. Download a JSON Private Key for the service account.
- E. On the Custom Metadata of the VM, add that JSON as the value for the key `compute-engine-service-account`.
- F. Download a JSON Private Key for the service account.
- G. After creating the VM, ssh into the VM and save the JSON under `~/gcloud/compute-engine-service-account.json`.

Answer: A

NEW QUESTION 198

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

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

Answer: A

Explanation:

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

NEW QUESTION 201

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

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

Answer: B

NEW QUESTION 202

You need a dynamic way of provisioning VMs on Compute Engine. The exact specifications will be in a dedicated configuration file. You want to follow Google's recommended practices. Which method should you use?

- A. Deployment Manager
- B. Cloud Composer
- C. Managed Instance Group
- D. Unmanaged Instance Group

Answer: A

Explanation:

<https://cloud.google.com/deployment-manager/docs/configuration/create-basic-configuration>

NEW QUESTION 207

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 210

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

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

Answer: D

Explanation:

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

NEW QUESTION 215

You are operating a Google Kubernetes Engine (GKE) cluster for your company where different teams can run non-production workloads. Your Machine Learning (ML) team needs access to Nvidia Tesla P100 GPUs to train their models. You want to minimize effort and cost. What should you do?

- A. Ask your ML team to add the “accelerator: gpu” annotation to their pod specification.
- B. Recreate all the nodes of the GKE cluster to enable GPUs on all of them.
- C. Create your own Kubernetes cluster on top of Compute Engine with nodes that have GPU
- D. Dedicate this cluster to your ML team.
- E. Add a new, GPU-enabled, node pool to the GKE cluste
- F. Ask your ML team to add the cloud.google.com/gke -accelerator: nvidia-tesla-p100 nodeSelector to their pod specification.

Answer: D

Explanation:

This is the most optimal solution. Rather than recreating all nodes, you create a new node pool with GPU enabled. You then modify the pod specification to target particular GPU types by adding node selector to your workloads Pod specification. YOu still have a single cluster so you pay Kubernetes cluster management fee for just one cluster thus minimizing the cost. Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/gpus> Ref: <https://cloud.google.com/kubern>

Example:

```
> apiVersion: v1
> kind: Pod
> metadata:
> name: my-gpu-pod
> spec:
> containers:
> name: my-gpu-container
> image: nvidia/cuda:10.0-runtime-ubuntu18.04
> command: [/bin/bash]
> resources:
> limits:
> nvidia.com/gpu: 2
> nodeSelector:
> cloud.google.com/gke-accelerator: nvidia-tesla-k80 # or nvidia-tesla-p100 or nvidia-tesla-p4 or nvidia-tesla-v100 or nvidia-tesla-t4
```

NEW QUESTION 218

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

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

Answer: D

Explanation:

(https://cloud.google.com/resource-manager/docs/project-migration#change_billing_account) <https://cloud.google.com/billing/docs/concepts>
<https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION 220

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 225

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

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

Answer: A

NEW QUESTION 226

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

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

Answer: A

Explanation:

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

NEW QUESTION 231

You want to permanently delete a Pub/Sub topic managed by Config Connector in your Google Cloud project. What should you do?

- A. Use `kubectl` to delete the topic resource.
- B. Use `gcloud` CLI to delete the topic.
- C. Use `kubectl` to create the label `deleted-by-cnrm` and to change its value to `true` for the topic resource.
- D. Use `gcloud` CLI to update the topic label `managed-by-cnrm` to `false`.

Answer: A

NEW QUESTION 236

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

- A. Change the default region property setting in the existing GCP project to `asia-northeast1`.

- B. Change the region property setting in the existing App Engine application from us-central to asia-northeast1.
- C. Create a second App Engine application in the existing GCP project and specify asia-northeast1 as the region to serve your application.
- D. Create a new GCP project and create an App Engine application inside this new project
- 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 240

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

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

Answer: D

Explanation:

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

NEW QUESTION 244

You are developing a new web application that will be deployed on Google Cloud Platform. As part of your release cycle, you want to test updates to your application on a small portion of real user traffic. The majority of the users should still be directed towards a stable version of your application. What should you do?

- A. Deploy me application on App Engine For each update, create a new version of the same service Configure traffic splitting to send a small percentage of traffic to the new version
- B. Deploy the application on App Engine For each update, create a new service Configure traffic splitting to send a small percentage of traffic to the new service.
- C. Deploy the application on Kubernetes Engine For a new release, update the deployment to use the new version
- D. Deploy the application on Kubernetes Engine For a now release, create a new deployment for the new version Update the service e to use the now deployment.

Answer: D

Explanation:

Keyword, Version, traffic splitting, App Engine supports traffic splitting for versions before releasing.

NEW QUESTION 249

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 -l 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 252

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

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

Answer: B

Explanation:

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

NEW QUESTION 256

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 258

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

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

Answer: C

Explanation:

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

NEW QUESTION 261

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 262

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

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

Answer:

B

Explanation:

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

NEW QUESTION 266

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

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

Answer: B

Explanation:

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

NEW QUESTION 268

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 for 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 270

Your Dataproc cluster runs in a single Virtual Private Cloud (VPC) network in a single subnet with range 172.16.20.128/25. There are no private IP addresses available in the VPC network. You want to add new VMs to communicate with your cluster using the minimum number of steps. What should you do?

- A. Modify the existing subnet range to 172.16.20.0/24.
- B. Create a new Secondary IP Range in the VPC and configure the VMs to use that range.
- C. Create a new VPC network for the VM
- D. Enable VPC Peering between the VMs' VPC network and the Dataproc cluster VPC network.
- E. Create a new VPC network for the VMs with a subnet of 172.32.0.0/16. Enable VPC network Peering between the Dataproc VPC network and the VMs VPC network
- F. Configure a custom Route exchange.

Answer: A

Explanation:

/25:
CIDR to IP Range Result
CIDR Range 172.16.20.128/25 Netmask 255.255.255.128
Wildcard Bits 0.0.0.127
First IP 172.16.20.128
First IP (Decimal) 2886734976 Last IP 172.16.20.255
Last IP (Decimal) 2886735103 Total Host 128
CIDR 172.16.20.128/25
/24:
CIDR to IP Range Result
CIDR Range 172.16.20.128/24 Netmask 255.255.255.0
Wildcard Bits 0.0.0.255
First IP 172.16.20.0
First IP (Decimal) 2886734848 Last IP 172.16.20.255
Last IP (Decimal) 2886735103 Total Host 256
CIDR 172.16.20.128/24

NEW QUESTION 273

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

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

Answer: A

Explanation:

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

NEW QUESTION 277

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

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

Answer: A

Explanation:

`gcloud services list --available` returns not only the enabled services in the project but also services that CAN be enabled.

<https://cloud.google.com/sdk/gcloud/reference/services/list#--available>

Run the following command to list the enabled APIs and services in your current project: gcloud services list

whereas, Run the following command to list the APIs and services available to you in your current project: gcloud services list --available

<https://cloud.google.com/sdk/gcloud/reference/services/list#--available>

--available

Return the services available to the project to enable. This list will include any services that the project has already enabled.

To list the services the current project has enabled for consumption, run: gcloud services list --enabled

To list the services the current project can enable for consumption, run: gcloud services list --available

NEW QUESTION 278

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 BigQuery 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 280

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

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

Answer: D

NEW QUESTION 282

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 284

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

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

Answer: D

NEW QUESTION 287

You are assisting a new Google Cloud user who just installed the Google Cloud SDK on their VM. The server needs access to Cloud Storage. The user wants your help to create a new storage bucket. You need to make this change in multiple environments. What should you do?

- A. Use a Deployment Manager script to automate creating storage buckets in an appropriate region
- B. Use a local SSD to improve performance of the VM for the targeted workload
- C. Use the gsutil command to create a storage bucket in the same region as the VM
- D. Use a Persistent Disk SSD in the same zone as the VM to improve performance of the VM

Answer: A

NEW QUESTION 288

You have a managed instance group comprised of preemptible VM's. All of the VM's keep deleting and recreating themselves every minute. What is a possible cause of this behavior?

- A. Your zonal capacity is limited, causing all preemptible VM's to be shutdown to recover capacity
- 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 in preemptible settings.
- E. Your managed instance group's health check is repeatedly failing, either to a misconfigured 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 293

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

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

Answer: C

Explanation:

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

<https://cloud.google.com/compute/docs/instances/preemptible>

NEW QUESTION 298

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 disk
- B. Create a snapshot schedule for the disk
- 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 disk

- E. If an outage occurs, create an instance in another zone with this disk attached.
- F. Store the application data on a regional persistent disk
- G. Create a snapshot schedule for the disk
- 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 300

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

- * Associate-Cloud-Engineer Questions and Answers Updated Frequently
- * Associate-Cloud-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Associate-Cloud-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Associate-Cloud-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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