



Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer

NEW QUESTION 1

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

- A. Deploy a Dataflow job from the batch template "Datastore to 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 topic
- E. Create an application that sends all medical images to the Pub/Sub topic

Answer: C

Explanation:

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

NEW QUESTION 2

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

- A. Create a signed URL with a four-hour expiration and share the URL with the company.
- B. Set object access to 'public' and use object lifecycle management to remove the object after four hours.
- C. Configure the storage bucket as a static website and furnish the object's URL to the company
- D. Delete the object from the storage bucket after four hours.
- E. Create a new Cloud Storage bucket specifically for the external company to access
- F. Copy the object to that bucket
- 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 3

You received a JSON file that contained a private key of a Service Account in order to get access to several resources in a Google Cloud project. You downloaded and installed the Cloud SDK and want to use this private key for authentication and authorization when performing gcloud commands. What should you do?

- A. Use the command `gcloud auth login` and point it to the private key
- B. Use the command `gcloud auth activate-service-account` and point it to the private key
- C. Place the private key file in the installation directory of the Cloud SDK and rename it to "credentials.json"
- D. Place the private key file in your home directory and rename it to "GOOGLE_APPLICATION_CREDENTIALS".

Answer: B

Explanation:

Authorizing with a service account

`gcloud auth activate-service-account` authorizes access using a service account. As with `gcloud init` and `gcloud auth login`, this command saves the service account credentials to the local system on successful completion and sets the specified account as the active account in your Cloud SDK configuration.

https://cloud.google.com/sdk/docs/authorizing#authorizing_with_a_service_account

NEW QUESTION 4

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 cluster
- 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 5

You are creating a Google Kubernetes Engine (GKE) cluster with a cluster autoscaler feature enabled. You need to make sure that each node of the cluster will run a monitoring pod that sends container metrics to a third-party monitoring solution. What should you do?

- A. Deploy the monitoring pod in a StatefulSet object.
- B. Deploy the monitoring pod in a DaemonSet object.
- C. Reference the monitoring pod in a Deployment object.
- D. Reference the monitoring pod in a cluster initializer at the GKE cluster creation time.

Answer: B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset> https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset#usage_patterns
DaemonSets attempt to adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed.

In GKE, DaemonSets manage groups of replicated Pods and adhere to a one-Pod-per-node model, either across the entire cluster or a subset of nodes. As you add nodes to a node pool, DaemonSets automatically add Pods to the new nodes as needed. So, this is a perfect fit for our monitoring pod.

Ref: <https://cloud.google.com/kubernetes-engine/docs/concepts/daemonset>

DaemonSets are useful for deploying ongoing background tasks that you need to run on all or certain nodes, and which do not require user intervention. Examples of such tasks include storage daemons like ceph, log collection daemons like fluentd, and node monitoring daemons like collectd. For example, you could have DaemonSets for each type of daemon run on all of your nodes. Alternatively, you could run multiple DaemonSets for a single type of daemon, but have them use different configurations for different hardware types and resource needs.

NEW QUESTION 6

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

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

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

Answer: A

Explanation:

When we create subnets in the same VPC with different CIDR ranges, they can communicate automatically within VPC. Resources within a VPC network can communicate with one another by using internal (private) IPv4 addresses, subject to applicable network firewall rules

Ref: <https://cloud.google.com/vpc/docs/vpc>

NEW QUESTION 8

Your company uses BigQuery for data warehousing. Over time, many different business units in your company have created 1000+ datasets across hundreds of projects. Your CIO wants you to examine all datasets to find tables that contain an employee_ssn column. You want to minimize effort in performing this task. What should you do?

- A. Go to Data Catalog and search for employee_ssn in the search box.
- B. Write a shell script that uses the bq command line tool to loop through all the projects in your organization.
- C. Write a script that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find the employee_ssn column.
- D. Write a Cloud Dataflow job that loops through all the projects in your organization and runs a query on INFORMATION_SCHEMA.COLUMNS view to find employee_ssn column.

Answer: A

Explanation:

<https://cloud.google.com/bigquery/docs/quickstarts/quickstart-web-ui?authuser=4>

NEW QUESTION 9

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

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

Answer: A

Explanation:

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it can't be changed. The only way is to create a new project and create an App Engine instance in europe-west3, send all user traffic to this instance and delete the app engine instance in us-central.

Ref: <https://cloud.google.com/appengine/docs/locations>

NEW QUESTION 10

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

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

Answer: B

Explanation:

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

NEW QUESTION 10

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: B

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 11

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 16

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

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

Answer: A

Explanation:

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

NEW QUESTION 18

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

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

Answer: C

Explanation:

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

NEW QUESTION 22

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

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

Answer: C

Explanation:

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

NEW QUESTION 27

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

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

Answer: C

Explanation:

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

NEW QUESTION 28

You need to manage a third-party application that will run on a Compute Engine instance. Other Compute Engine instances are already running with default configuration. Application installation files are hosted on Cloud Storage. You need to access these files from the new instance without allowing other virtual machines (VMs) to access these files. What should you do?

- A. Create the instance with the default Compute Engine service account Grant the service account permissions on Cloud Storage.
- B. Create the instance with the default Compute Engine service account Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.
- C. Create a new service account and assign this service account to the new instance Grant the service account permissions on Cloud Storage.
- D. Create a new service account and assign this service account to the new instance Add metadata to the objects on Cloud Storage that matches the metadata on the new instance.

Answer: C

Explanation:

<https://cloud.google.com/iam/docs/best-practices-for-using-and-managing-service-accounts>

If an application uses third-party or custom identities and needs to access a resource, such as a BigQuery dataset or a Cloud Storage bucket, it must perform a transition between principals. Because Google Cloud APIs don't recognize third-party or custom identities, the application can't propagate the end-user's identity to BigQuery or Cloud Storage. Instead, the application has to perform the access by using a different Google identity.

NEW QUESTION 30

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

```
> $ kubectl get pods
NAME READY STATUS RESTARTS AGE
nginx-84748895c4-nqqmt 1/1 Running 0 9m41s
> $ kubectl delete pod nginx-84748895c4-nqqmt
pod nginx-84748895c4-nqqmt deleted
> $ kubectl get pods
NAME READY STATUS RESTARTS AGE
nginx-84748895c4-k6bzl 1/1 Running 0 25s
```

What should you do to delete the deployment and avoid pod getting recreated?

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

Answer: A

Explanation:

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

```
> $ kubectl delete deployment nginx
deployment.apps/nginx deleted
```

Ref: <https://kubernetes.io/docs/reference/kubectl/cheatsheet/#deleting-resources>

NEW QUESTION 33

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

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

Answer: B

Explanation:

Storage Admin (roles/storage.admin) Grants full control of buckets and objects.

When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.

`firebase.projects.get resourceManager.projects.get resourceManager.projects.list storage.buckets.* storage.objects.*`

<https://cloud.google.com/storage/docs/access-control/iam-roles>

This role grants full control of buckets and objects. When applied to an individual bucket, control applies only to the specified bucket and objects within the bucket.

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

NEW QUESTION 37

You have a large 5-TB AVRO file stored in a Cloud Storage bucket. Your analysts are proficient only in SQL and need access to the data stored in this file. You want to find a cost-effective way to complete their request as soon as possible. What should you do?

- A. Load data in Cloud Datastore and run a SQL query against it.
- B. Create a BigQuery table and load data in BigQuery
- C. Run a SQL query on this table and drop this table after you complete your request.
- D. Create external tables in BigQuery that point to Cloud Storage buckets and run a SQL query on these external tables to complete your request.
- E. Create a Hadoop cluster and copy the AVRO file to NDfs by compressing it
- F. Load the file in a hive table and provide access to your analysts so that they can run SQL queries.

Answer: C

Explanation:

<https://cloud.google.com/bigquery/external-data-sources>

An external data source is a data source that you can query directly from BigQuery, even though the data is not stored in BigQuery storage.

BigQuery supports the following external data sources: Amazon S3

Azure Storage Cloud Bigtable Cloud Spanner Cloud SQL Cloud Storage Drive

NEW QUESTION 41

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

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

Answer: A

Explanation:

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

NEW QUESTION 43

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 46

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 47

You have a Dockerfile that you need to deploy on Kubernetes Engine. What should you do?

- A. Use `kubectl app deploy <dockerfilename>`.
- B. Use `gcloud app deploy <dockerfilename>`.
- C. Create a docker image from the Dockerfile and upload it to Container Registry
- D. Create a Deployment YAML file to point to that image
- E. Use `kubectl` to create the deployment with that file.
- F. Create a docker image from the Dockerfile and upload it to Cloud Storage
- G. Create a Deployment YAML file to point to that image
- H. Use `kubectl` to create the deployment with that file.

Answer: C

NEW QUESTION 52

You are the team lead of a group of 10 developers. You provided each developer with an individual Google Cloud Project that they can use as their personal sandbox to experiment with different Google Cloud solutions. You want to be notified if any of the developers are spending above \$500 per month on their sandbox

environment. What should you do?

- A. Create a single budget for all projects and configure budget alerts on this budget.
- B. Create a separate billing account per sandbox project and enable BigQuery billing export
- C. Create a Data Studio dashboard to plot the spending per billing account.
- D. Create a budget per project and configure budget alerts on all of these budgets.
- E. Create a single billing account for all sandbox projects and enable BigQuery billing export
- F. Create a Data Studio dashboard to plot the spending per project.

Answer: C

Explanation:

Set budgets and budget alerts Overview Avoid surprises on your bill by creating Cloud Billing budgets to monitor all of your Google Cloud charges in one place. A budget enables you to track your actual Google Cloud spend against your planned spend. After you've set a budget amount, you set budget alert threshold rules that are used to trigger email notifications. Budget alert emails help you stay informed about how your spend is tracking against your budget. 2. Set budget scope Set the budget Scope and then click Next. In the Projects field, select one or more projects that you want to apply the budget alert to. To apply the budget alert to all the projects in the Cloud Billing account, choose Select all.

<https://cloud.google.com/billing/docs/how-to/budgets#budget-scope>

NEW QUESTION 54

Your VMs are running in a subnet that has a subnet mask of 255.255.255.240. The current subnet has no more free IP addresses and you require an additional 10 IP addresses for new VMs. The existing and new VMs should all be able to reach each other without additional routes. What should you do?

- A. Use gcloud to expand the IP range of the current subnet.
- B. Delete the subnet, and recreate it using a wider range of IP addresses.
- C. Create a new projec
- D. Use Shared VPC to share the current network with the new project.
- E. Create a new subnet with the same starting IP but a wider range to overwrite the current subnet.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/compute/networks/subnets/expand-ip-range>

gcloud compute networks subnets expand-ip-range - expand the IP range of a Compute Engine subnetwork gcloud compute networks subnets expand-ip-range NAME --prefix-length=PREFIX_LENGTH [--region=REGION] [GCLOUD_WIDE_FLAG ...]

NEW QUESTION 56

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

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

Answer: B

NEW QUESTION 58

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

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

Answer: A

NEW QUESTION 63

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

- A. Cloud Run for Anthos
- B. Cloud Run
- C. App Engine Standard
- D. Cloud Functions.

Answer: D

Explanation:

Cloud Functions is Google Cloud's event-driven serverless compute platform that lets you run your code locally or in the cloud without having to provision servers.

Cloud Functions scales up or down, so you pay only for compute resources you use. Cloud Functions have excellent integration with Cloud Pub/Sub, lets you scale down to zero and is recommended by Google as the ideal serverless platform to use when dependent on Cloud Pub/Sub."If you're building a simple API (a small set of functions to be accessed via HTTP or Cloud Pub/Sub), we recommend using Cloud Functions."Ref: <https://cloud.google.com/serverless-options>

NEW QUESTION 65

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

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

Answer: C

Explanation:

<https://cloud.google.com/logging/docs/audit>

Data Access audit logs Data Access audit logs contain API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data.

<https://cloud.google.com/logging/docs/audit#data-access>

NEW QUESTION 69

You deployed an LDAP server on Compute Engine that is reachable via TLS through port 636 using UDP. You want to make sure it is reachable by clients over that port. What should you do?

- A. Add the network tag allow-udp-636 to the VM instance running the LDAP server.
- B. Create a route called allow-udp-636 and set the next hop to be the VM instance running the LDAP server.
- C. Add a network tag of your choice to the instance
- D. Create a firewall rule to allow ingress on UDP port 636 for that network tag.
- E. Add a network tag of your choice to the instance running the LDAP serve
- F. Create a firewall rule to allow egress on UDP port 636 for that network tag.

Answer: C

Explanation:

A tag is simply a character string added to a tags field in a resource, such as Compute Engine virtual machine (VM) instances or instance templates. A tag is not a separate resource, so you cannot create it separately. All resources with that string are considered to have that tag. Tags enable you to make firewall rules and routes applicable to specific VM instances.

NEW QUESTION 73

Your company's infrastructure is on-premises, but all machines are running at maximum capacity. You want to burst to Google Cloud. The workloads on Google Cloud must be able to directly communicate to the workloads on-premises using a private IP range. What should you do?

- A. In Google Cloud, configure the VPC as a host for Shared VPC.
- B. In Google Cloud, configure the VPC for VPC Network Peering.
- C. Create bastion hosts both in your on-premises environment and on Google Clou
- D. Configure both as proxy servers using their public IP addresses.
- E. Set up Cloud VPN between the infrastructure on-premises and Google Cloud.

Answer: D

Explanation:

"Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization."

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

"Cloud Interconnect provides low latency, high availability connections that enable you to reliably transfer data between your on-premises and Google Cloud Virtual Private Cloud (VPC) networks."

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/overview>

and "HA VPN is a high-availability (HA) Cloud VPN solution that lets you securely connect your on-premises network to your VPC network through an IPsec VPN connection in a single region."

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

NEW QUESTION 75

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

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

Answer: A

Explanation:

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

NEW QUESTION 78

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 81

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

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

Answer: A

Explanation:

<https://cloud.google.com/storage/docs/lifecycle#setstorageclass-cost>

The object's time spent set at the original storage class counts towards any minimum storage duration that applies for the new storage class.

NEW QUESTION 83

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

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

Answer: B

Explanation:

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

NEW QUESTION 84

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

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

Answer: B

Explanation:

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

https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances#best_practices

NEW QUESTION 88

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>



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 92

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 93

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

- A. Using the Cloud SDK, create a new project, enable the Compute Engine API in that project, and then create the instance specifying your new project.
- B. Enable the Compute Engine API in the Cloud Console, use the Cloud SDK to create the instance, and then use the `--project` flag to specify a new project.
- C. Using the Cloud SDK, create the new instance, and use the `--project` flag to specify the new project. Answer yes when prompted by Cloud SDK to enable the Compute Engine API.
- D. Enable the Compute Engine API in the Cloud Console
- E. Go to the Compute Engine section of the Console to create a new instance, and look for the Create In A New Project option in the creation form.

Answer: A

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/projects/create> Quickstart: Creating a New Instance Using the Command Line Before you begin

* 1. In the Cloud Console, on the project selector page, select or create a Cloud project.

* 2. Make sure that billing is enabled for your Google Cloud project. Learn how to confirm billing is enabled for your project.

To use the gcloud command-line tool for this quickstart, you must first install and initialize the Cloud SDK:

* 1. Download and install the Cloud SDK using the instructions given on Installing Google Cloud SDK.

* 2. Initialize the SDK using the instructions given on Initializing Cloud SDK.

To use gcloud in Cloud Shell for this quickstart, first activate Cloud Shell using the instructions given on Starting Cloud Shell.

<https://cloud.google.com/ai-platform/deep-learning-vm/docs/quickstart-cli#before-you-begin>

NEW QUESTION 94

You need to manage a Cloud Spanner Instance for best query performance. Your instance in production runs in a single Google Cloud region. You need to improve performance in the shortest amount of time. You want to follow Google best practices for service configuration. What should you do?

- A. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45% If you exceed this threshold, add nodes to your instance.
- B. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 45% Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage
- C. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65% If you exceed this threshold, add nodes to your instance
- D. Create an alert in Cloud Monitoring to alert when the percentage of high priority CPU utilization reaches 65%. Use database query statistics to identify queries that result in high CPU usage, and then rewrite those queries to optimize their resource usage.

Answer: C

Explanation:

<https://cloud.google.com/spanner/docs/cpu-utilization#recommended-max>

NEW QUESTION 99

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

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

Answer: D

Explanation:

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

NEW QUESTION 100

Your company has embraced a hybrid cloud strategy where some of the applications are deployed on Google Cloud. A Virtual Private Network (VPN) tunnel connects your Virtual Private Cloud (VPC) in Google Cloud with your company's on-premises network. Multiple applications in Google Cloud need to connect to an on-premises database server, and you want to avoid having to change the IP configuration in all of your applications when the IP of the database changes. What should you do?

- A. Configure Cloud NAT for all subnets of your VPC to be used when egressing from the VM instances.
- B. Create a private zone on Cloud DNS, and configure the applications with the DNS name.
- C. Configure the IP of the database as custom metadata for each instance, and query the metadata server.
- D. Query the Compute Engine internal DNS from the applications to retrieve the IP of the database.

Answer: B

Explanation:

Forwarding zones Cloud DNS forwarding zones let you configure target name servers for specific private zones. Using a forwarding zone is one way to implement outbound DNS forwarding from your VPC network. A Cloud DNS forwarding zone is a special type of Cloud DNS private zone. Instead of creating records within the zone, you specify a set of forwarding targets. Each forwarding target is an IP address of a DNS server, located in your VPC network, or in an on-premises network connected to your VPC network by Cloud VPN or Cloud Interconnect.

<https://cloud.google.com/nat/docs/overview>

DNS configuration Your on-premises network must have DNS zones and records configured so that Google domain names resolve to the set of IP addresses for either private.googleapis.com or restricted.googleapis.com. You can create Cloud DNS managed private zones and use a Cloud DNS inbound server policy, or you can configure on-premises name servers. For example, you can use BIND or Microsoft Active Directory DNS.

<https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid#config-domain>

NEW QUESTION 102

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 106

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

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

Answer: C

Explanation:

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

NEW QUESTION 109

You want to find out when users were added to Cloud Spanner Identity Access Management (IAM) roles on your Google Cloud Platform (GCP) project. What should you do in the GCP Console?

- A. Open the Cloud Spanner console to review configurations.
- B. Open the IAM & admin console to review IAM policies for Cloud Spanner roles.
- C. Go to the Stackdriver Monitoring console and review information for Cloud Spanner.
- D. Go to the Stackdriver Logging console, review admin activity logs, and filter them for Cloud Spanner IAM roles.

Answer: D

Explanation:

<https://cloud.google.com/monitoring/audit-logging>

NEW QUESTION 110

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

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

Answer: C

Explanation:

➤ Given that the new instance wants to access the application on the existing compute engine instance, these applications seem to be related so they should be within the same VPC. It is possible to have them in different VPCs and peer the VPCs but this is a lot of additional work and we can simplify this by choosing the option below (which is the answer)

* 1. Create a subnet in the same VPC, in europe-west1.

* 2. Create the new instance in the new subnet and use the first instance subnets private address as the endpoint. is the right answer.

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

Ref: <https://cloud.google.com/vpc>

NEW QUESTION 115

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 117

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 120

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

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

Answer: B

Explanation:

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

NEW QUESTION 124

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 127

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

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

Answer: A

Explanation:

"Run gcloud configurations list to start the Compute Engine instances". How the heck are you expecting to "start" GCE instances doing "configuration list". Each gcloud configuration has a 1 to 1 relationship with the region (if a region is defined). Since we have two different regions, we would need to create two separate configurations using gcloud config configurations createRef: <https://cloud.google.com/sdk/gcloud/reference/config/configurations/create>

Secondly, you can activate each configuration independently by running gcloud config configurations activate [NAME]Ref: <https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate>

<https://cloud.google.com/sdk/gcloud/reference/config/configurations/activate>

Finally, while each configuration is active, you can run the gcloud compute instances start [NAME] command to start the instance in the configurations region.<https://cloud.google.com/sdk/gcloud/reference/compute/instances/start>

NEW QUESTION 128

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

- A. Split the users from business units to multiple projects.
- B. Apply a user- or project-level custom query quota for BigQuery data warehouse.
- C. Create separate copies of your BigQuery data warehouse for each business unit.
- D. Split your BigQuery data warehouse into multiple data warehouses for each business unit.
- E. Change your BigQuery query model from on-demand to flat rate.
- 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 129

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: C

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 132

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 136

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

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

Answer: D

NEW QUESTION 140

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: A

Explanation:

IAP controls access to your App Engine apps and Compute Engine VMs running on Google Cloud. It leverages user identity and the context of a request to

determine if a user should be allowed access. IAP is a building block toward BeyondCorp, an enterprise security model that enables employees to work from untrusted networks without using a VPN.

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

Email/password

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

OIDC

Phone number Custom Anonymous

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

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

NEW QUESTION 144

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

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

Answer: D

Explanation:

GCP App Engine natively offers traffic splitting functionality between versions. You can use traffic splitting to specify a percentage distribution of traffic across two or more of the versions within a service. Splitting traffic allows you to conduct A/B testing between your versions and provides control over the pace when rolling out features.

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

NEW QUESTION 145

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 146

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 150

You are deploying an application to a Compute Engine VM in a managed instance group. The application must be running at all times, but only a single instance of the VM should run per GCP project. How should you configure the instance group?

- A. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- B. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- C. Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 2.
- D. Set autoscaling to Off, set the minimum number of instances to 1, and then set the maximum number of instances to 2.

Answer: A

Explanation:

<https://cloud.google.com/compute/docs/autoscaler#specifications>

Autoscaling works independently from autohealing. If you configure autohealing for your group and an instance fails the health check, the autohealer attempts to recreate the instance. Recreating an instance can cause the number of instances in the group to fall below the autoscaling threshold (minNumReplicas) that you specify.

- Since we need the application running at all times, we need a minimum 1 instance.
 - Only a single instance of the VM should run, we need a maximum 1 instance.
 - We want the application running at all times. If the VM crashes due to any underlying hardware failure, we want another instance to be added to MIG so that application can continue to serve requests. We can achieve this by enabling autoscaling. The only option that satisfies these three is Set autoscaling to On, set the minimum number of instances to 1, and then set the maximum number of instances to 1.
- Ref: <https://cloud.google.com/compute/docs/autoscaler>

NEW QUESTION 152

Your organization uses Active Directory (AD) to manage user identities. Each user uses this identity for federated access to various on-premises systems. Your security team has adopted a policy that requires users to log into Google Cloud with their AD identity instead of their own login. You want to follow the Google-recommended practices to implement this policy. What should you do?

- A. Sync Identities with Cloud Directory Sync, and then enable SAML for single sign-on
- B. Sync Identities in the Google Admin console, and then enable Oauth for single sign-on
- C. Sync identities with 3rd party LDAP sync, and then copy passwords to allow simplified login with (he same credentials
- D. Sync identities with Cloud Directory Sync, and then copy passwords to allow simplified login with the same credentials.

Answer: A

NEW QUESTION 157

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 161

You need to configure optimal data storage for files stored in Cloud Storage for minimal cost. The files are used in a mission-critical analytics pipeline that is used continually. The users are in Boston, MA (United States). What should you do?

- A. Configure regional storage for the region closest to the users Configure a Nearline storage class
- B. Configure regional storage for the region closest to the users Configure a Standard storage class
- C. Configure dual-regional storage for the dual region closest to the users Configure a Nearline storageclass
- D. Configure dual-regional storage for the dual region closest to the users Configure a Standard storage class

Answer: D

Explanation:

Keywords: - continually -> Standard - mission-critical analytics -> dual-regional

NEW QUESTION 162

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 163

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

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

Answer: A

Explanation:

You can deploy to a different project by using `--project` flag.

By default, the service is deployed the current project configured via:

```
$ gcloud config set core/project PROJECT
```

To override this value for a single deployment, use the `--project` flag:

```
$ gcloud app deploy ~/my_app/app.yaml --project=PROJECT
```

 Ref: <https://cloud.google.com/sdk/gcloud/reference/app/deploy>

NEW QUESTION 165

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 project
- 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 project
- 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 166

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

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

Answer: C

Explanation:

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

NEW QUESTION 171

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

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

Answer: B

Explanation:

Predefined roles

The following table describes Identity and Access Management (IAM) roles that are associated with Cloud Storage and lists the permissions that are contained in each role. Unless otherwise noted, these roles can be applied either to entire projects or specific buckets.

Storage Object Creator (roles/storage.objectCreator) Allows users to create objects. Does not give permission to view, delete, or overwrite objects.

<https://cloud.google.com/storage/docs/access-control/iam-roles#standard-roles>

NEW QUESTION 176

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

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

Answer: D

Explanation:

Keyword need to look for

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

NEW QUESTION 177

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 180

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

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

Answer: A

Explanation:

Adding an API as a type provider

This page describes how to add an API to Google Cloud Deployment Manager as a type provider. To learn more about types and type providers, read the Types overview documentation.

A type provider exposes all of the resources of a third-party API to Deployment Manager as base types that you can use in your configurations. These types must be directly served by a RESTful API that supports Create, Read, Update, and Delete (CRUD).

If you want to use an API that is not automatically provided by Google with Deployment Manager, you must add the API as a type provider.

<https://cloud.google.com/deployment-manager/docs/configuration/type-providers/creating-type-provider>

NEW QUESTION 182

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 184

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

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

Answer: B

Explanation:

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

NEW QUESTION 189

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 190

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 195

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 199

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

- A. Create a health check on port 443 and use that when creating the Managed Instance Group.
- B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.

- C. In the Instance Template, add the label 'health-check'.
- D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Answer: A

Explanation:

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

NEW QUESTION 204

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 206

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 209

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

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

Answer: C

NEW QUESTION 213

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 215

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 219

You are running multiple microservices in a Kubernetes Engine cluster. One microservice is rendering images.

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

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

Answer: B

NEW QUESTION 223

You need to create a new billing account and then link it with an existing Google Cloud Platform project. What should you do?

- A. Verify that you are Project Billing Manager for the GCP projec
- B. Update the existing project to link it to the existing billing account.
- C. Verify that you are Project Billing Manager for the GCP projec
- D. Create a new billing account and link the new billing account to the existing project.
- E. Verify that you are Billing Administrator for the billing accoun
- F. Create a new project and link the new project to the existing billing account.
- G. Verify that you are Billing Administrator for the billing accoun
- H. Update the existing project to link it to the existing billing account.

Answer: B

Explanation:

Billing Administrators can not create a new billing account, and the project is presumably already created. Project Billing Manager allows you to link the created billing account to the project. It is vague on how the billing account gets created but by process of elimination

NEW QUESTION 226

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 role stage 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 229

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

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

Answer: C

NEW QUESTION 232

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

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

- K. Use Stackdriver to determine the provisioned and used resource amount
- L. Multiply the 1 week cost to determine the monthly costs.

Answer: A

Explanation:

You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You don't incur any charges for doing so.
Ref: <https://cloud.google.com/products/calculator>

NEW QUESTION 237

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 238

You need to verify that a Google Cloud Platform service account was created at a particular time. What should you do?

- A. Filter the Activity log to view the Configuration category
- B. Filter the Resource type to Service Account.
- C. Filter the Activity log to view the Configuration category
- D. Filter the Resource type to Google Project.
- E. Filter the Activity log to view the Data Access category
- F. Filter the Resource type to Service Account.
- G. Filter the Activity log to view the Data Access category
- H. Filter the Resource type to Google Project.

Answer: A

Explanation:

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

NEW QUESTION 243

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 244

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 health check 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 247

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

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

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