

Google

Exam Questions Associate-Cloud-Engineer

Google Cloud Certified - Associate Cloud Engineer



NEW QUESTION 1

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

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

Answer: B

Explanation:

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

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

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

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

NEW QUESTION 2

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

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

Answer: D

NEW QUESTION 3

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 images 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 4

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 5

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 6

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

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

Answer: D

Explanation:

This is the most optimal solution. Rather than recreating all nodes, you create a new node pool with GPU enabled. You then modify the pod specification to target particular GPU types by adding node selector to your workloads Pod specification. YOu still have a single cluster so you pay Kubernetes cluster management fee for just one cluster thus minimizing the

cost.Ref: <https://cloud.google.com/kubernetes-engine/docs/how-to/gpus>Ref: <https://cloud.google.com/kubern>

Example:

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

NEW QUESTION 7

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 8

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

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

Answer: C

Explanation:

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

NEW QUESTION 9

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 10

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 10

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 15

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

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

Answer: D

Explanation:

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

NEW QUESTION 17

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 20

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

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

Answer: D

Explanation:

"...we recommend that you enable Cloud Billing data export to BigQuery at the same time that you create a Cloud Billing account. "

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery>

<https://medium.com/google-cloud/analyzing-google-cloud-billing-data-with-big-query-30bae1c2aae4>

NEW QUESTION 23

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 25

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

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

Answer: C

Explanation:

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

NEW QUESTION 28

You are building an application that stores relational data from users. Users across the globe will use this application. Your CTO is concerned about the scaling requirements because the size of the user base is unknown. You need to implement a database solution that can scale with your user growth with minimum configuration changes. Which storage solution should you use?

- A. Cloud SQL
- B. Cloud Spanner
- C. Cloud Firestore
- D. Cloud Datastore

Answer: B

Explanation:

Cloud Spanner is a relational database and is highly scalable. Cloud Spanner is a highly scalable, enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with a non-relational horizontal scale. This combination delivers high-performance transactions and strong consistency across rows, regions, and continents with an industry-leading 99.999% availability SLA, no planned downtime, and enterprise-grade security

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

Graphical user interface, application, Teams Description automatically generated

	CLOUD SPANNER	TRADITIONAL RELATIONAL	TRADITIONAL NON-RELATIONAL
Schema	✓ Yes	✓ Yes	✗ No
SQL	✓ Yes	✓ Yes	✗ No
Consistency	✓ Strong	✓ Strong	✗ Eventual
Availability	✓ High	✗ Failover	✓ High
Scalability	✓ Horizontal	✗ Vertical	✓ Horizontal
Replication	✓ Automatic	⚙️ Configurable	⚙️ Configurable

NEW QUESTION 32

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

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

Answer: D

Explanation:

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

NEW QUESTION 33

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

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

Answer: C

Explanation:

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

NEW QUESTION 36

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 41

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 are assigned the Billing Administrator IAM role for 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. 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. Set the default key value project labels to department marketing for all services in this project

Answer: A

NEW QUESTION 43

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 47

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

- A. Enable node auto-provisioning on the GKE cluster.
- B. Create a VerticalPodAutscaler for those workloads.
- C. Create a node pool with preemptible VMs and GPUs attached to those VMs.

D. Create a node pool of instances with GPUs, and enable autoscaling on this node pool with a minimum size of 1.

Answer: A

Explanation:

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

NEW QUESTION 48

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

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

Answer: A

Explanation:

> Google HTTP(S) Load Balancing has native support for the WebSocket protocol when you use HTTP or HTTPS, not HTTP/2, as the protocol to the backend.
Ref: https://cloud.google.com/load-balancing/docs/https#websocket_proxy_support

> We dont need to convert WebSocket code to use HTTP streaming or Redesign the application, as WebSocket support is offered by Google HTTP(S) Load Balancing. Reviewing the encryption requirements is a good idea but it has nothing to do with WebSockets.

NEW QUESTION 51

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

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

Answer: C

Explanation:

This command correctly lists pods that have the label `app=prod`. When creating the deployment, we used the label `app=prod` so listing pods that have this label retrieve the pods belonging to nginx deployments. You can list pods by using Kubernetes CLI `kubectl get pods`.

Ref: <https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/>

Ref: <https://kubernetes.io/docs/tasks/access-application-cluster/list-all-running-container-images/#list-containe>

NEW QUESTION 54

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 57

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

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

Answer: C

NEW QUESTION 60

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

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

Answer: B

Explanation:

<https://cloud.google.com/sql/docs/postgres>

NEW QUESTION 65

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

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

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

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

Answer: AB

Explanation:

(A) Object Lifecycle Management Delete

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

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

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

(B) Signed URLs

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

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

NEW QUESTION 68

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

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

Answer: D

Explanation:

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

NEW QUESTION 70

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 72

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 77

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 82

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

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

Answer: D

Explanation:

C:\GCP\appeng>gcloud config list [core]

account = xxx@gmail.com disable_usage_reporting = False project = my-first-demo-xxxx

<https://cloud.google.com/endpoints/docs/openapi/troubleshoot-gce-deployment>

NEW QUESTION 85

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

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

Answer: B

Explanation:

The resource that you need to get access is in the other project. roles/bigquery.dataViewer BigQuery Data Viewer

When applied to a table or view, this role provides permissions to: Read data and metadata from the table or view.

This role cannot be applied to individual models or routines.

When applied to a dataset, this role provides permissions to: Read the dataset's metadata and list tables in the dataset. Read data and metadata from the dataset's tables.

When applied at the project or organization level, this role can also enumerate all datasets in the project. Additional roles, however, are necessary to allow the running of jobs.

NEW QUESTION 87

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 90

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 91

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 95

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

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

Answer: A

Explanation:

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

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

Google Cloud Storage Classes in the Organization

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

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

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

NEW QUESTION 99

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 101

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 106

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

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

Answer: B

Explanation:

Shared VPC allows an organization to connect resources from multiple projects to a common Virtual Private Cloud (VPC) network, so that they can communicate with each other securely and efficiently using internal IPs from that network. When you use Shared VPC, you designate a project as a host project and attach one or more other service projects to it. The VPC networks in the host project are called Shared VPC networks. Eligible resources from service projects can use subnets in the Shared VPC network

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

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

NEW QUESTION 109

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

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

Answer: C

NEW QUESTION 114

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

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

Answer: C

NEW QUESTION 116

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

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

Answer: A

Explanation:

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

NEW QUESTION 118

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

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

Answer: D

Explanation:

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

NEW QUESTION 121

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

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

Answer: A

Explanation:

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

NEW QUESTION 123

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

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

Answer: B

NEW QUESTION 128

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

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

Answer: D

NEW QUESTION 130

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

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

Answer: A

Explanation:

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

NEW QUESTION 134

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

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

Answer: BE

Explanation:

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

NEW QUESTION 138

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

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp-deployment
spec:
  selector:
    matchLabels:
      app: myapp
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp
    spec:
      containers:
        - name: myapp
          image: myapp:1.1
          ports:
            - containerPort: 80

```

```

apiVersion: v1
kind: Service
metadata:
  name: myapp-service
spec:
  ports:
    - port: 8000
      targetPort: 80
      protocol: TCP
  selector:
    app: myapp

```

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

```

kubectl get pods -l app=myapp

```

NAME	READY	STATUS	RESTART	AGE
myapp-deployment-58ddb995-lp86m	0/1	Pending	0	9m
myapp-deployment-58ddb995-qjpkg	1/1	Running	0	9m

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

- A. Review details of the myapp-service Service object and check for error messages.
- B. Review details of the myapp-deployment Deployment object and check for error messages.
- C. Review details of myapp-deployment-58ddb995-lp86m Pod and check for warning messages.
- D. View logs of the container in myapp-deployment-58ddb995-lp86m pod and check for warning messages.

Answer: C

Explanation:

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

NEW QUESTION 139

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 144

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 147

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

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

Answer: D

Explanation:

A simple approach would be to use the command flags available when listing all the IAM policy for a given project. For instance, the following command: `gcloud projects get-iam-policy $PROJECT_ID`

`--flatten="bindings[].members" --format="table(bindings.members)" --filter="bindings.role:roles/owner"` outputs all the users and service accounts associated with the role 'roles/owner' in the project in question. <https://groups.google.com/g/google-cloud-dev/c/Z6sZs7TvygQ?pli=1>

NEW QUESTION 150

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 154

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 155

You are performing a monthly security check of your Google Cloud environment and want to know who has access to view data stored in your Google Cloud Project. What should you do?

- A. Enable Audit Logs for all APIs that are related to data storage.
- B. Review the IAM permissions for any role that allows for data access
- C. Most Voted
- D. Review the Identity-Aware Proxy settings for each resource.
- E. Create a Data Loss Prevention job.

Answer: B

Explanation:

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

NEW QUESTION 160

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

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

Answer: B

Explanation:

Google Cloud Storage Triggers

Cloud Functions can respond to change notifications emerging from Google Cloud Storage. These notifications can be configured to trigger in response to various events inside a bucket—object creation, deletion, archiving and metadata updates.

Note: Cloud Functions can only be triggered by Cloud Storage buckets in the same Google Cloud Platform project.

Event types

Cloud Storage events used by Cloud Functions are based on Cloud Pub/Sub Notifications for Google Cloud Storage and can be configured in a similar way.

Supported trigger type values are: google.storage.object.finalize google.storage.object.delete google.storage.object.archive google.storage.object.metadataUpdate

Object Finalize

Trigger type value: google.storage.object.finalize

This event is sent when a new object is created (or an existing object is overwritten, and a new generation of that object is created) in the bucket.

https://cloud.google.com/functions/docs/calling/storage#event_types

NEW QUESTION 161


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

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

Answer: C

Explanation:

Text Description automatically generated with low confidence



Cloud Functions is Google Cloud's event-driven serverless compute platform. It automatically scales based on the load and requires no additional configuration. You pay only for the resources used.

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

While all other options i.e. Google Compute Engine, Google Kubernetes Engine, Google App Engine support autoscaling, it needs to be configured explicitly based on the load and is not as trivial as the scale up or scale down offered by Google's cloud functions.

NEW QUESTION 163

You have 32 GB of data in a single file that you need to upload to a Nearline Storage bucket. The WAN connection you are using is rated at 1 Gbps, and you are the only one on the connection. You want to use as much of the rated 1 Gbps as possible to transfer the file rapidly. How should you upload the file?

- A. Use the GCP Console to transfer the file instead of gsutil.
- B. Enable parallel composite uploads using gsutil on the file transfer.
- C. Decrease the TCP window size on the machine initiating the transfer.
- D. Change the storage class of the bucket from Nearline to Multi-Regional.

Answer: B

Explanation:

<https://cloud.google.com/storage/docs/parallel-composite-uploads> <https://cloud.google.com/storage/docs/uploads-downloads#parallel-composite-uploads>

NEW QUESTION 167

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

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

Answer: D

Explanation:

Google Cloud Coldline is a new cold-tier storage for archival data with access frequency of less than once per year. Unlike other cold storage options, Nearline has no delays prior to data access, so now it is the leading solution among competitors.

The Real description is about Coldline storage Class: Coldline Storage

Coldline Storage is a very-low-cost, highly durable storage service for storing infrequently accessed data. Coldline Storage is a better choice than Standard Storage or Nearline Storage in scenarios where slightly lower availability, a 90-day minimum storage duration, and higher costs for data access are acceptable trade-offs for lowered at-rest storage costs.

Coldline Storage is ideal for data you plan to read or modify at most once a quarter. Note, however, that for data being kept entirely for backup or archiving purposes, Archive Storage is more cost-effective, as it offers the lowest storage costs.

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

NEW QUESTION 169

You are using Container Registry to centrally store your company's container images in a separate project. In another project, you want to create a Google Kubernetes Engine (GKE) cluster. You want to ensure that Kubernetes can download images from Container Registry. What should you do?

- A. In the project where the images are stored, grant the Storage Object Viewer IAM role to the service account used by the Kubernetes nodes.
- B. When you create the GKE cluster, choose the Allow full access to all Cloud APIs option under 'Access scopes'.
- C. Create a service account, and give it access to Cloud Storage
- D. Create a P12 key for this service account and use it as an imagePullSecrets in Kubernetes.
- E. Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account.

Answer: A

Explanation:

Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account. is not right. As mentioned above, Container Registry ignores permissions set on individual objects within the storage bucket so this isn't going to work.

Ref: <https://cloud.google.com/container-registry/docs/access-control>

NEW QUESTION 170

You have successfully created a development environment in a project for an application. This application uses Compute Engine and Cloud SQL. Now, you need to create a production environment for this application.

The security team has forbidden the existence of network routes between these 2 environments, and asks you to follow Google-recommended practices. What should you do?

- A. Create a new project, enable the Compute Engine and Cloud SQL APIs in that project, and replicate the setup you have created in the development environment.
- B. Create a new production subnet in the existing VPC and a new production Cloud SQL instance in your existing project, and deploy your application using those resources.
- C. Create a new project, modify your existing VPC to be a Shared VPC, share that VPC with your new project, and replicate the setup you have in the development environment in that new project, in the Shared VPC.
- D. Ask the security team to grant you the Project Editor role in an existing production project used by another division of your company.
- E. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A

Explanation:

This aligns with Google's recommended practices. By creating a new project, we achieve complete isolation between development and production environments; as well as isolate this production application from production applications of other departments.

Ref: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#define-hierarchy>

NEW QUESTION 171

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 to 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 176

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 179

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

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

Answer: A

Explanation:

Choosing a region and zone You choose which region or zone hosts your resources, which controls where your data is stored and used. Choosing a region and zone is important for several reasons:

Handling failures

Distribute your resources across multiple zones and regions to tolerate outages. Google designs zones to be independent from each other: a zone usually has power, cooling, networking, and control planes that are isolated from other zones, and most single failure events will affect only a single zone. Thus, if a zone becomes unavailable, you can transfer traffic to another zone in the same region to keep your services running. Similarly, if a region experiences any disturbances, you should have backup services running in a different region. For more information about distributing your resources and designing a robust system, see [Designing Robust Systems](#). Decreased network latency To decrease network latency, you might want to choose a region or zone that is close to your point of service.

https://cloud.google.com/compute/docs/regions-zones#choosing_a_region_and_zone

NEW QUESTION 182

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

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

Answer: A

Explanation:

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

NEW QUESTION 187

You've deployed a microservice called `myapp1` to a Google Kubernetes Engine cluster using the YAML file specified below:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: myapp1-deployment
spec:
  selector:
    matchLabels:
      app: myapp1
  replicas: 2
  template:
    metadata:
      labels:
        app: myapp1
    spec:
      containers:
        - name: main-container
          image: gcr.io/my-company-repo/myapp1:1.4
          env:
            - name: DB_PASSWORD
              value: "t0ugh2guess!"
          ports:
            - containerPort: 8080

```

You need to refactor this configuration so that the database password is not stored in plain text. You want to follow Google-recommended practices. What should you do?

- A. Store the database password inside the Docker image of the container, not in the YAML file.
- B. Store the database password inside a Secret objec
- C. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
- D. Store the database password inside a ConfigMap objec
- E. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
- F. Store the database password in a file inside a Kubernetes persistent volume, and use a persistent volume claim to mount the volume to the container.

Answer: B

Explanation:

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

NEW QUESTION 191

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

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

Answer: C

Explanation:

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

NEW QUESTION 192

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 196

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

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

Answer: C

Explanation:

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

NEW QUESTION 199

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 204

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 206

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 208

You are building a pipeline to process time-series data. Which Google Cloud Platform services should you put in boxes 1,2,3, and 4?



- A. Cloud Pub/Sub, Cloud Dataflow, Cloud Datastore, BigQuery
- B. Firebase Messages, Cloud Pub/Sub, Cloud Spanner, BigQuery
- C. Cloud Pub/Sub, Cloud Storage, BigQuery, Cloud Bigtable
- D. Cloud Pub/Sub, Cloud Dataflow, Cloud Bigtable, BigQuery

Answer: D

NEW QUESTION 210

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

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

Answer: A

NEW QUESTION 212

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

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

Answer: A

NEW QUESTION 217

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

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

Answer: A

Explanation:

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

NEW QUESTION 218

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 do?

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

Answer: A

Explanation:

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

NEW QUESTION 222

Your organization has user identities in Active Directory. Your organization wants to use Active Directory as their source of truth for identities. Your organization wants to have full control over the Google accounts used by employees for all Google services, including your Google Cloud Platform (GCP) organization. What should you do?

- A. Use Google Cloud Directory Sync (GCDS) to synchronize users into Cloud Identity.
- B. Use the cloud Identity APIs and write a script to synchronize users to Cloud Identity.
- C. Export users from Active Directory as a CSV and import them to Cloud Identity via the Admin Console.
- D. Ask each employee to create a Google account using self signu
- E. Require that each employee use their company email address and password.

Answer: A

NEW QUESTION 227

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

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

Answer: D

NEW QUESTION 230

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

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

Answer: B

Explanation:

Best practices for persistent disk snapshots

You can create persistent disk snapshots at any time, but you can create snapshots more quickly and with greater reliability if you use the following best practices.

Creating frequent snapshots efficiently

Use snapshots to manage your data efficiently.

Create a snapshot of your data on a regular schedule to minimize data loss due to unexpected failure. Improve performance by eliminating excessive snapshot downloads and by creating an image and reusing it. Set your snapshot schedule to off-peak hours to reduce snapshot time.

Snapshot frequency limits

Creating snapshots from persistent disks

You can snapshot your disks at most once every 10 minutes. If you want to issue a burst of requests to snapshot your disks, you can issue at most 6 requests in 60 minutes.

If the limit is exceeded, the operation fails and returns the following error: <https://cloud.google.com/compute/docs/disks/snapshot-best-practices>

NEW QUESTION 233

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

- A. Add the auditors group to the 'logging.viewer' and 'bigQuery.dataViewer' predefined IAM roles.

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 236

You want to configure 10 Compute Engine instances for availability when maintenance occurs. Your requirements state that these instances should attempt to automatically restart if they crash. Also, the instances should be highly available including during system maintenance. What should you do?

- A. Create an instance template for the instance
- B. Set the 'Automatic Restart' to o
- C. Set the 'On-host maintenance' to Migrate VM instanc
- D. Add the instance template to an instance group.
- E. Create an instance template for the instance
- F. Set 'Automatic Restart' to of
- G. Set 'On-host maintenance' to Terminate VM instance
- H. Add the instance template to an instance group.
- I. Create an instance group for the instance
- J. Set the 'Autohealing' health check to healthy (HTTP).
- K. Create an instance group for the instanc
- L. Verify that the 'Advanced creation options' setting for 'do not retry machine creation' is set to off.

Answer: A

Explanation:

Create an instance template for the instances so VMs have same specs. Set the "'Automatic Restart' to on to VM automatically restarts upon crash. Set the "'On-host maintenance' to Migrate VM instance. This will take care of VM during maintenance window. It will migrate VM instance making it highly available Add the instance template to an instance group so instances can be managed.

- onHostMaintenance: Determines the behavior when a maintenance event occurs that might cause your instance to reboot.
- [Default] MIGRATE, which causes Compute Engine to live migrate an instance when there is a maintenance event.
- TERMINATE, which stops an instance instead of migrating it.
- automaticRestart: Determines the behavior when an instance crashes or is stopped by the system.
- [Default] true, so Compute Engine restarts an instance if the instance crashes or is stopped.
- false, so Compute Engine does not restart an instance if the instance crashes or is stopped.

Enabling automatic restart ensures that compute engine instances are automatically restarted when they crash. And Enabling Migrate VM Instance enables live migrates i.e. compute instances are migrated during system maintenance and remain running during the migration.

Automatic Restart If your instance is set to terminate when there is a maintenance event, or if your instance crashes because of an underlying hardware issue, you can set up Compute Engine to automatically restart the instance by setting the automaticRestart field to true. This setting does not apply if the instance is taken offline through a user action, such as calling sudo shutdown, or during a zone

outage.Ref: <https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#autorestart>

Enabling the Migrate VM Instance option migrates your instance away from an infrastructure maintenance event, and your instance remains running during the migration. Your instance might experience a short period of decreased performance, although generally, most instances should not notice any difference. This is ideal for instances that require constant uptime and can tolerate a short period of decreased

performance.Ref: https://cloud.google.com/compute/docs/instances/setting-instance-scheduling-options#live_

NEW QUESTION 240

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 244

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

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

Answer: A

Explanation:

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

NEW QUESTION 245

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 247

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 project
- B. Update the existing project to link it to the existing billing account.
- C. Verify that you are Project Billing Manager for the GCP project
- 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 account
- 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 account
- 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 249

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

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

Answer: A

Explanation:

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

NEW QUESTION 251

You have files in a Cloud Storage bucket that you need to share with your suppliers. You want to restrict the time that the files are available to your suppliers to 1 hour. You want to follow Google recommended practices. What should you do?

- A. Create a service account with just the permissions to access files in the bucket
- B. Create a JSON key for the service account
- C. Execute the command `gsutil signurl -m 1h gs://**`.
- D. Create a service account with just the permissions to access files in the bucket
- E. Create a JSON key for the service account
- F. Execute the command `gsutil signurl -d 1h gs://**`.
- G. Create a service account with just the permissions to access files in the bucket
- H. Create a JSON key for the service account
- I. Execute the command `gsutil signurl -p 60m gs://`.
- J. Create a JSON key for the Default Compute Engine Service Account
- K. Execute the command `gsutil signurl -t 60m gs://***`

Answer: B

Explanation:

This command correctly specifies the duration that the signed url should be valid for by using the -d flag. The default is 1 hour so omitting the -d flag would have also resulted in the same outcome. Times may be specified with no suffix (default hours), or with s = seconds, m = minutes, h = hours, d = days. The max duration allowed is 7d. Ref: <https://cloud.google.com/storage/docs/gsutil/commands/signurl>

NEW QUESTION 255

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

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

Answer: A

Explanation:

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

NEW QUESTION 256

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

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

Answer: C

Explanation:

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

NEW QUESTION 257

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 261

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

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

Answer: A

Explanation:

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

NEW QUESTION 263

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

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

Answer: D

NEW QUESTION 265

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

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

Answer: D

Explanation:

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

<https://cloud.google.com/bigquery/docs/access-control>

NEW QUESTION 267

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

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

Answer: A

Explanation:

<https://cloud.google.com/run/docs/tutorials/gcloud> <https://cloud.google.com/resource-manager/docs/creating-managing-projects>

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

You can shut down projects using the Cloud Console. When you shut down a project, this immediately happens: All billing and traffic serving stops, You lose access to the project, The owners of the project will be notified and can stop the deletion within 30 days, The project will be scheduled to be deleted after 30 days. However, some resources may be deleted much earlier.

NEW QUESTION 272

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

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

Answer: D

Explanation:

as the instances (normal or preemptible) would be terminated and relaunched if the health check fails either due to application not configured properly or the instances firewall do not allow health check to happen.

GCP provides health check systems that connect to virtual machine (VM) instances on a configurable, periodic basis. Each connection attempt is called a probe. GCP records the success or failure of each probe.

Health checks and load balancers work together. Based on a configurable number of sequential successful or failed probes, GCP computes an overall health state for each VM in the load balancer. VMs that respond successfully for the configured number of times are considered healthy. VMs that fail to respond successfully for a separate number of times are unhealthy.

GCP uses the overall health state of each VM to determine its eligibility for receiving new requests. In addition to being able to configure probe frequency and health state thresholds, you can configure the criteria that define a successful probe.

NEW QUESTION 273

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

Associate-Cloud-Engineer Practice Exam Features:

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

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