

Google.Associate-Cloud-Engineer.v2024-01-29.q127

Exam Code:	Associate-Cloud-Engineer
Exam Name:	Google Associate Cloud Engineer Exam
Certification Provider:	Google
Free Question Number:	127
Version:	v2024-01-29
# of views:	1000
# of Questions views:	1270
https://www.freepdfdumps.com/Google.Associate-Cloud-Engineer.v2024-01-29.q127.html	

NEW QUESTION: 1

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 ([LEAVE A REPLY](#))

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: 2

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 (LEAVE A REPLY)

[https://developer.mozilla.org/en-](https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_the_correct_mime_type)

[US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_the_correct_mime_type](https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/MIME_Types#importance_of_setting_the_correct_mime_type)

NEW QUESTION: 3

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

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

Answer: B (LEAVE A REPLY)

You write a lifecycle management rule in XML and push it to the bucket with `gsutil lifecycle set config-xml-file`. is not right.

`gsutil lifecycle set` enables you to set the lifecycle configuration on one or more buckets based on the configuration file provided. However, XML is not a valid supported type for the configuration file.

Ref: <https://cloud.google.com/storage/docs/gsutil/commands/lifecycle>

Write a script that runs `gsutil ls -lr gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days. Repeat this process every morning. is not right.

This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

Write a script that runs `gsutil ls -l gs://myapp-gcp-ace-logs/**` to find and remove items older than 90 days. Schedule the script with cron. is not right.

This manual approach is error-prone, time-consuming and expensive. GCP Cloud Storage provides lifecycle management rules that let you achieve this with minimal effort.

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

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

Ref: <https://cloud.google.com/storage/docs/gsutil/commands/lifecycle>

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

NEW QUESTION: 4

You have a virtual machine that is currently configured with 2 vCPUs and 4 GB of memory. It is running out of memory. You want to upgrade the virtual machine to have 8 GB of memory. What should you do?

- A. Rely on live migration to move the workload to a machine with more memory.
- B. Use `gcloud` to add metadata to the VM. Set the key to `required-memory-size` and the value to 8 GB.
- C. Stop the VM, change the machine type to `n1-standard-8`, and start the VM.
- D. Stop the VM, increase the memory to 8 GB, and start the VM.

Answer: (SHOW ANSWER)

In Google compute engine, if predefined machine types don't meet your needs, you can create an instance with custom virtualized hardware settings. Specifically, you can create an instance with a custom number of vCPUs and custom memory, effectively using a custom machine type. Custom machine types are ideal for the following scenarios: 1. Workloads that aren't a good fit for the predefined machine types that are available to you. 2. Workloads that require more processing power or more memory but don't need all of the upgrades that are provided by the next machine type level. In our scenario, we only need a memory upgrade. Moving to a bigger instance would also bump up the CPU which we don't need so we have to use a custom machine type. It is not possible to change memory while the instance is running so you need to first stop the instance,

change the memory and then start it again. See below a screenshot that shows how CPU/Memory can be customized for an instance that has been stopped. Ref:

<https://cloud.google.com/compute/docs/instances/creating-instance-with-custom-machine-type>

NEW QUESTION: 5

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

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 6

You used the gcloud container clusters command to create two Google Cloud Kubernetes (GKE) clusters prod-cluster and dev-cluster.

* prod-cluster is a standard cluster.

* dev-cluster is an auto-pilot cluster.

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

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

A.

B.

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

C.

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

D.

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 7

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 page. Use the pricing calculator to total the monthly costs for each GCP product.
- B.** For each GCP product in the solution, review the pricing details on the products pricing page. Create a Google Sheet that summarizes the expected monthly costs for each product.
- C.** Provision the solution on GCP. Leave the solution provisioned for 1 week. Navigate to the Billing Report page in the Google Cloud Platform Console. Multiply the 1 week cost to determine the monthly costs.
- D.** Provision the solution on GCP. Leave the solution provisioned for 1 week. Use Stackdriver to determine the provisioned and used resource amounts. Multiply the 1 week cost to determine the monthly costs.

Answer: ([SHOW ANSWER](#))

You can use the Google Cloud Pricing Calculator to total the estimated monthly costs for each GCP product. You don't incur any charges for doing so.

Ref: <https://cloud.google.com/products/calculator>

NEW QUESTION: 8

Your organization needs to grant users access to query datasets in BigQuery but prevent them from accidentally deleting the datasets. You want a solution that follows Google-recommended practices. What should you do?

- A.** Add users to roles/bigquery user role only, instead of roles/bigquery dataOwner.
- B.** Add users to roles/bigquery dataEditor role only, instead of roles/bigquery dataOwner.
- C.** Create a custom role by removing delete permissions, and add users to that role only.
- D.** Create a custom role by removing delete permissions. Add users to the group, and then add the group to the custom role.

Answer: **D** ([LEAVE A REPLY](#))

https://cloud.google.com/bigquery/docs/access-control#custom_roles

Custom roles enable you to enforce the principle of least privilege, ensuring that the user and service accounts in your organization have only the permissions essential to performing their intended functions.

NEW QUESTION: 9

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 (LEAVE A REPLY)

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

NEW QUESTION: 10

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 desktop. Verify that a firewall rule for port 3389 exists.
- B. Install a RDP client in your desktop. Set a Windows username and password in the GCP Console. Use the credentials to log in to the instance.
- C. Set a Windows password in the GCP Console. Verify that a firewall rule for port 22 exists. Click the RDP button in the GCP Console and supply the credentials to log in.
- D. Set a Windows username and password in the GCP Console. Verify that a firewall rule for port 3389 exists. Click the RDP button in the GCP Console, and supply the credentials to log in.

Answer: D (LEAVE A REPLY)

<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: 11

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 projects. Grant the VM Service Account the role Storage Object Creator on corp-iot-insights.
- D. 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 (LEAVE A REPLY)

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: 12

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 ([LEAVE A REPLY](#))

Reference:

Service Account User (roles/iam.serviceAccountUser): Includes permissions to list service accounts, get details about a service account, and impersonate a service account. Service Account Admin (roles/iam.serviceAccountAdmin): Includes permissions to list service accounts and get details about a service account. Also includes permissions to create, update, and delete service accounts, and to view or change the IAM policy on a service account.

NEW QUESTION: 13

Your organization is a financial company that needs to store audit log files for 3 years. Your organization has hundreds of Google Cloud projects. You need to implement a cost-effective approach for log file retention. What should you do?

- A. Create an export to the sink that saves logs from Cloud Audit to BigQuery.
- B. Create an export to the sink that saves logs from Cloud Audit to a Coldline Storage bucket.
- C. Write a custom script that uses logging API to copy the logs from Stackdriver logs to BigQuery.
- D. Export these logs to Cloud Pub/Sub and write a Cloud Dataflow pipeline to store logs to Cloud SQL.

Answer: B ([LEAVE A REPLY](#))

Coldline Storage is the perfect service to store audit logs from all the projects and is very cost-efficient as well. Coldline Storage is a very low-cost, highly durable storage service for storing infrequently accessed data.

NEW QUESTION: 14

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 (LEAVE A REPLY)

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: 15

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. Add the group to the backend service for the load balancer. When all instances in the new managed instance group are healthy, delete the old managed instance group.
- D.** Create a new instance template with the new application version. Update the existing managed instance group with the new instance template. 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 (LEAVE A REPLY)

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

NEW QUESTION: 16

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 (LEAVE A REPLY)

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>

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

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

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

Answer: B (LEAVE A REPLY)

Reference:

On-demand pricing Under on-demand pricing, BigQuery charges for queries by using one metric: the number of bytes processed (also referred to as bytes read). You are charged for the number of bytes processed whether the data is stored in BigQuery or in an external data source such as Cloud Storage, Drive, or Cloud Bigtable. On-demand pricing is based solely on usage.

https://cloud.google.com/bigquery/pricing#on_demand_pricing

NEW QUESTION: 18

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

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

Answer: D ([LEAVE A REPLY](#))

Keyword need to look for

- "High Throughput",
 - "Consistent",
 - "Property based data insert/fetch like engine status, distance traveled, fuel level, and more." which can be designed in column,
 - "Large Scale Customer Base + Each Customer has multiple sensor which send event in seconds"
- This will go for petabytes situation,
- Export data based on the time of the event.
 - Atomic

o BigTable will fit all requirement.

o DataStore is not fully Atomic

o CloudStorage is not an option where we can export data based on time of event. We need another solution to do that o Firestore can be used with MobileSDK.

NEW QUESTION: 19

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

- A. Use the projects.move method to move the project to your organization. Update the billing account of the project to that of your organization.
- B. Ensure that you have an Organization Administrator Identity and Access Management (IAM) role assigned to you in both organizations. Navigate to the Resource Manager in the startup's Google Cloud organization, and drag the project to your company's organization.
- C. Create an infrastructure-as-code template for all resources in the project by using Terraform. and deploy that template to a new project in your organization. Delete the project from the startup's Google Cloud organization.

D. Create a Private Catalog for the Google Cloud Marketplace, and upload the resources of the startup's production project to the Catalog. Share the Catalog with your organization, and deploy the resources in your company's project.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 20

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

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

Answer: ([SHOW ANSWER](#)**)**

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

NEW QUESTION: 21

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 ([LEAVE A REPLY](#))

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

NEW QUESTION: 22

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

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

Answer: D ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/blog/products/gcp/introducing-coldline-and-a-unified-platform-for-data-storage> Cloud Storage Coldline: a low-latency storage class for long-term archiving Coldline is a

new Cloud Storage class designed for long-term archival and disaster recovery. Coldline is perfect for the archival needs of big data or multimedia content, allowing businesses to archive years of data. Coldline provides fast and instant (millisecond) access to data and changes the way that companies think about storing and accessing their cold data.

NEW QUESTION: 23

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 24

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.

Answer: A ([LEAVE A REPLY](#))

- Restore from backups when notified.

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: 25

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

- A. Create a health check on port 443 and use that when creating the Managed Instance Group.
- B. Select Multi-Zone instead of Single-Zone when creating the Managed Instance Group.
- C. In the Instance Template, add the label 'health-check'.
- D. In the Instance Template, add a startup script that sends a heartbeat to the metadata server.

Answer: A (LEAVE A REPLY)

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

NEW QUESTION: 26

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

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

Answer: D (LEAVE A REPLY)

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

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

NEW QUESTION: 27

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

Answer: C (LEAVE A REPLY)

<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: 28

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: (SHOW ANSWER)

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: 29

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

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

Answer: (SHOW ANSWER)

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

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

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

NEW QUESTION: 30

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

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

Answer: C (LEAVE A REPLY)

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0> You create the service account in proj-sa and take note of the service account email, then you go to proj-vm in IAM > ADD and add the service account's email as new member and give it the Compute Storage Admin role.

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

NEW QUESTION: 31

After a recent security incident, your startup company wants better insight into what is happening in the Google Cloud environment. You need to monitor unexpected firewall changes and instance creation. Your company prefers simple solutions. What should you do?

- A. Use Cloud Logging filters to create log-based metrics for firewall and instance actions. Monitor the changes and set up reasonable alerts.
- B. Install Kibana on a compute Instance. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Pub/Sub. Target the Pub/Sub topic to push messages to the Kibana instance. Analyze the logs on Kibana in real time.

C. Turn on Google Cloud firewall rules logging, and set up alerts for any insert, update, or delete events.

D. Create a log sink to forward Cloud Audit Logs filtered for firewalls and compute instances to Cloud Storage. Use BigQuery to periodically analyze log events in the storage bucket.

Answer: A (LEAVE A REPLY)

This answer is the simplest and most effective way to monitor unexpected firewall changes and instance creation in Google Cloud. Cloud Logging filters allow you to specify the criteria for the log entries that you want to view or export. You can use the Logging query language to write filters based on the LogEntry fields, such as resource.type, severity, or protoPayload.methodName. For example, you can filter for firewall-related events by using the following query:

resource.type="gce_subnetwork" logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Ffirewall" You can filter for instance-related events by using the following query:

resource.type="gce_instance" logName="projects/PROJECT_ID/logs/compute.googleapis.com%2Factivity_log" You can create log-based metrics from these filters to measure the rate or count of log entries that match the filter. Log-based metrics can be used to create charts and dashboards in Cloud Monitoring, or to set up alerts based on the metric values. For example, you can create an alert policy that triggers when the log-based metric for firewall changes exceeds a certain threshold in a given time interval. This way, you can get notified of any unexpected or malicious changes to your firewall rules.

Option B is incorrect because it is unnecessarily complex and costly. Installing Kibana on a compute instance requires additional configuration and maintenance. Creating a log sink to forward Cloud Audit Logs to Pub/Sub also incurs additional charges for the Pub/Sub service. Analyzing the logs on Kibana in real time may not be feasible or efficient, as it requires constant monitoring and manual intervention.

Option C is incorrect because Google Cloud firewall rules logging is a different feature from Cloud Audit Logs. Firewall rules logging allows you to audit, verify, and analyze the effects of your firewall rules by creating connection records for each rule that applies to traffic. However, firewall rules logging does not log the insert, update, or delete events for the firewall rules themselves. Those events are logged by Cloud Audit Logs, which record the administrative activities in your Google Cloud project.

Option D is incorrect because it is not a real-time solution. Creating a log sink to forward Cloud Audit Logs to Cloud Storage requires additional storage space and charges. Using BigQuery to periodically analyze log events in the storage bucket also incurs additional costs for the BigQuery service. Moreover, this option does not provide any alerting mechanism to notify you of any unexpected or malicious changes to your firewall rules or instances.

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have**

been updated and answers have been corrected get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 32

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. Once they grant you that role, replicate the setup you have in the development environment in that project.

Answer: A (LEAVE A REPLY)

This aligns with Googles 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: 33

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: (SHOW ANSWER)

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 So the next possible step is to Meet with the cloud enablement team to discuss load balancer options. We don't 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: 34

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

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

Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 35

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 days.
 3. Use Cloud Scheduler to trigger both Cloud Functions daily at 1:00 AM.
- D.** 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 (LEAVE A REPLY)

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: 36

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 (LEAVE A REPLY)

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-containers-filtering-by-pod-label>

NEW QUESTION: 37

You are using Google Kubernetes Engine with autoscaling enabled to host a new application. You want to expose this new application to the public, using HTTPS on a public IP address. What should you do?

- A. Create a Kubernetes Service of type `NodePort` for your application, and a Kubernetes Ingress to expose this Service via a Cloud Load Balancer.
- B. Create a Kubernetes Service of type `ClusterIP` for your application. Configure the public DNS name of your application using the IP of this Service.
- C. Create a Kubernetes Service of type `NodePort` to expose the application on port 443 of each node of the Kubernetes cluster. Configure the public DNS name of your application with the IP of every node of the cluster to achieve load-balancing.

D. Create a HAProxy pod in the cluster to load-balance the traffic to all the pods of the application. Forward the public traffic to HAProxy with an iptable rule. Configure the DNS name of your application using the public IP of the node HAProxy is running on.

Answer: A ([LEAVE A REPLY](#))

Reference:

Create a Kubernetes Service of type ClusterIP for your application. Configure the public DNS name of your application using the IP of this Service. is not right.

Kubernetes Service of type ClusterIP exposes the Service on a cluster-internal IP. Choosing this value makes the Service only reachable from within the cluster so you can not route external traffic to this IP.

Ref: <https://kubernetes.io/docs/concepts/services-networking/service/>

NEW QUESTION: 38

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. 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.

C. Set up Cloud VPN between your Google Cloud VPC and the internal network of the operations partner.

D. Ask the operations partner to generate SSH key pairs, and add the public keys to the VM instances.

Answer: ([SHOW ANSWER](#)**)**

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: 39

Your continuous integration and delivery (CI/CD) server can't execute Google Cloud actions in a specific project because of permission issues. You need to validate whether the used service account has the appropriate roles in the specific project. What should you do?

- A. Open the Google Cloud console, and run a query to determine which resources this service account can access.
- B. Open the Google Cloud console, and run a query of the audit logs to find permission denied errors for this service account.
- C. Open the Google Cloud console, and check the organization policies.
- D. Open the Google Cloud console, and check the Identity and Access Management (IAM) roles assigned to the service account at the project or inherited from the folder or organization levels.

Answer: D ([LEAVE A REPLY](#))

This answer is the most effective way to validate whether the service account used by the CI/CD server has the appropriate roles in the specific project. By checking the IAM roles assigned to the service account, you can see which permissions the service account has and which resources it can access. You can also check if the service account inherits any roles from the folder or organization levels, which may affect its access to the project. You can use the Google Cloud console, the gcloud command-line tool, or the IAM API to view the IAM roles of a service account.

NEW QUESTION: 40

You need to update a deployment in Deployment Manager without any resource downtime in the deployment. Which command should you use?

- A. `gcloud deployment-manager resources create --config <deployment-config-path>`
- B. `gcloud deployment-manager deployments create --config <deployment-config-path>`
- C. `gcloud deployment-manager resources update --config <deployment-config-path>`
- D. `gcloud deployment-manager deployments update --config <deployment-config-path>`

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 41

You are hosting an application on bare-metal servers in your own data center. The application needs access to Cloud Storage. However, security policies prevent the servers hosting the application from having public IP addresses or access to the internet. You want to follow Google-recommended practices to provide the application with access to Cloud Storage. What should you do?

- A. 1. Use nslookup to get the IP address for storage.googleapis.com. 2. Negotiate with the security team to be able to give a public IP address to the servers. 3. Only allow egress traffic from those servers to the IP addresses for storage.googleapis.com.
- B. 1. Using Cloud VPN, create a VPN tunnel to a Virtual Private Cloud (VPC) in Google Cloud Platform (GCP). 2. In this VPC, create a Compute Engine instance and install the Squid proxy

server on this instance.3. Configure your servers to use that instance as a proxy to access Cloud Storage.

C. 1. Use Migrate for Compute Engine (formerly known as Velostrata) to migrate those servers to Compute Engine.2. Create an internal load balancer (ILB) that uses storage.googleapis.com as backend.3. Configure your new instances to use this ILB as proxy.

D. 1. Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP.2. Use Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.3. In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com.

Answer: D (LEAVE A REPLY)

Our requirement is to follow Google recommended practices to achieve the end result. Configuring Private Google Access for On-Premises Hosts is best achieved by VPN/Interconnect + Advertise Routes + Use restricted Google IP Range.

Using Cloud VPN or Interconnect, create a tunnel to a VPC in GCP

Using Cloud Router to create a custom route advertisement for 199.36.153.4/30. Announce that network to your on-premises network through the VPN tunnel.

In your on-premises network, configure your DNS server to resolve *.googleapis.com as a CNAME to restricted.googleapis.com is the right answer right, and it is what Google recommends.

Ref: <https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid> You must configure routes so that Google API traffic is forwarded through your Cloud VPN or Cloud Interconnect connection, firewall rules on your on-premises firewall to allow the outgoing traffic, and DNS so that traffic to Google APIs resolves to the IP range youve added to your routes.

You can use Cloud Router Custom Route Advertisement to announce the Restricted Google APIs IP addresses through Cloud Router to your on-premises network. The Restricted Google APIs IP range is 199.36.153.4/30. While this is technically a public IP range, Google does not announce it publicly. This IP range is only accessible to hosts that can reach your Google Cloud projects through internal IP ranges, such as through a Cloud VPN or Cloud Interconnect connection.

Without having a public IP address or access to the internet, the only way you could connect to cloud storage is if you have an internal route to it.

So Negotiate with the security team to be able to give public IP addresses to the servers is not right. Following Google recommended practices is synonymous with using Googles services (Not quite, but it is at least for the exam !!).

So In this VPC, create a Compute Engine instance and install the Squid proxy server on this instance is not right.

Migrating the VM to Compute Engine is a bit drastic when Google says it is perfectly fine to have Hybrid Connectivity architectures <https://cloud.google.com/hybrid-connectivity>.

So,

Use Migrate for Compute Engine (formerly known as Velostrata) to migrate these servers to Compute Engine is not right.

NEW QUESTION: 42

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

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

Answer: A (LEAVE A REPLY)

Quickstart: using the Google Cloud Console

This page shows you how to perform basic tasks in Pub/Sub using the Google Cloud Console.

Note: If you are new to Pub/Sub, we recommend that you start with the interactive tutorial.

Before you begin

Set up a Cloud Console project.

Set up a project

Click to:

Create or select a project.

Enable the Pub/Sub API for that project.

You can view and manage these resources at any time in the Cloud Console.

Install and initialize the Cloud SDK.

Note: You can run the gcloud tool in the Cloud Console without installing the Cloud SDK. To run the gcloud tool in the Cloud Console, use Cloud Shell .

<https://cloud.google.com/pubsub/docs/quickstart-console>

NEW QUESTION: 43

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

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

Answer: A (LEAVE A REPLY)

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: 44

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



Each tier has a service account that is associated with all instances within it. You need to enable communication on TCP port 8080 between tiers as follows:

- * Instances in tier #1 must communicate with tier #2.
- * Instances in tier #2 must communicate with tier #3.

What should you do?

A. 1. Create an ingress firewall rule with the following settings: * Targets: all instances * Source filter: IP ranges (with the range set to 10.0.2.0/24) * Protocols: allow all 2. Create an ingress firewall

rule with the following settings:* Targets: all instances* Source filter: IP ranges (with the range set to 10.0.1.0/24)* Protocols: allow all

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

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

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

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

Answer: B (LEAVE A REPLY)

1. Create an ingress firewall rule with the following settings:

- * Targets: all instances with tier #2 service account

- * Source filter: all instances with tier #1 service account

- * Protocols: allow TCP:8080 2. Create an ingress firewall rule with the following settings:

- * Targets: all instances with tier #3 service account

- * Source filter: all instances with tier #2 service account

- * Protocols: allow TCP: 8080

NEW QUESTION: 45

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

A. Verify that you are the project billing administrator. Select the associated billing account and create a budget and alert for the appropriate project.

B. Verify that you are the project billing administrator. Select the associated billing account and create a budget and a custom alert.

C. Verify that you are the project administrator. Select the associated billing account and create a budget for the appropriate project.

D. Verify that you are project administrator. Select the associated billing account and create a budget and a custom alert.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 46

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

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

Answer: C (LEAVE A REPLY)

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

https://cloud.google.com/resource-manager/docs/access-control-proj#using_primitive_roles

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

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount:

Freepdfdumps)

NEW QUESTION: 47

You recently discovered that your developers are using many service account keys during their development process. While you work on a long term improvement, you need to quickly implement a process to enforce short-lived service account credentials in your company. You have the following requirements:

- * All service accounts that require a key should be created in a centralized project called pj-sa.
- * Service account keys should only be valid for one day.

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

A. Implement a Cloud Run job to rotate all service account keys periodically in pj-sa. Enforce an org policy to deny service account key creation with an exception to pj-sa.

B. Implement a Kubernetes Cronjob to rotate all service account keys periodically. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

C. Enforce an org policy constraint allowing the lifetime of service account keys to be 24 hours. Enforce an org policy constraint denying service account key creation with an exception on pj-sa.

D. Enforce a DENY org policy constraint over the lifetime of service account keys for 24 hours. Disable attachment of service accounts to resources in all projects with an exception to pj-sa.

Answer: ([SHOW ANSWER](#))

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

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

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

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

Reference:

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

5: Create and delete service account keys - Google Cloud

Organization policy constraints for service accounts

NEW QUESTION: 48

Your existing application running in Google Kubernetes Engine (GKE) consists of multiple pods running on four GKE n1-standard-2 nodes. You need to deploy additional pods requiring n2-highmem-16 nodes without any downtime. What should you do?

A. Use gcloud container clusters upgrade. Deploy the new services.

B. Create a new Node Pool and specify machine type n2-highmem-16. Deploy the new pods.

C. Create a new cluster with n2-highmem-16 nodes. Redeploy the pods and delete the old cluster.

D. Create a new cluster with both n1-standard-2 and n2-highmem-16 nodes. Redeploy the pods and delete the old cluster.

Answer: **B** ([LEAVE A REPLY](#))

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

NEW QUESTION: 49

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 ([LEAVE A REPLY](#))

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: 50

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: (SHOW ANSWER)

<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: 51

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

- A. In the Google Cloud console, use the export functionality of the Cost table. Create a Looker Studio dashboard on top of the CSV export.
- B. In the Google Cloud console, visualize the costs related to the projects in the Cost breakdown section.
- C. In the Google Cloud console, visualize the costs related to the projects in the Reports section.
- D. Configure Cloud Billing data export to BigQuery for the billing account. Create a Looker Studio dashboard on top of the BigQuery export.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 52

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. Create a P12 key for this service account and use it as an imagePullSecrets in Kubernetes.
- D. Configure the ACLs on each image in Cloud Storage to give read-only access to the default Compute Engine service account.

Answer: A ([LEAVE A REPLY](#))

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: 53

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 54

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: ([SHOW ANSWER](#)**)**

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: 55

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 (LEAVE A REPLY)

"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 create` Ref:

<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: 56

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

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

Answer: C (LEAVE A REPLY)

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

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

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

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

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

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

NEW QUESTION: 57

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

- A.** Create a single custom VPC with 2 subnets. Create each subnet in a different region and with a different CIDR range.
- B.** Create a single custom VPC with 2 subnets. Create each subnet in the same region and with the same CIDR range.
- C.** Create 2 custom VPCs, each with a single subnet. Create each subnet in a different region and with a different CIDR range.
- D.** Create 2 custom VPCs, each with a single subnet. Create each subnet in the same region and with the same CIDR range.

Answer: A (LEAVE A REPLY)

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

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

NEW QUESTION: 58

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

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

Answer: (SHOW ANSWER)

The Datastore emulator provides local emulation of the production Datastore environment. You can use the emulator to develop and test your application locally Ref:

<https://cloud.google.com/datastore/docs/tools/datastore-emulator>

NEW QUESTION: 59

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 (LEAVE A REPLY)

"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: 60

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

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

Answer: B (LEAVE A REPLY)

Reference:

Default behavior does not grant access to the "your GCP Project" Default behavior allow only create billing account and project - When the organization is created, all users in your domain are automatically granted Project Creator and Billing Account Creator IAM roles at the organization level. This enables users in your domain to continue creating projects with no disruption.

NEW QUESTION: 61

You are developing a new application and are looking for a Jenkins installation to build and deploy your source code. You want to automate the installation as quickly and easily as possible. What should you do?

- A. Deploy Jenkins through the Google Cloud Marketplace.
- B. Create a new Compute Engine instance. Run the Jenkins executable.
- C. Create a new Kubernetes Engine cluster. Create a deployment for the Jenkins image.
- D. Create an instance template with the Jenkins executable. Create a managed instance group with this template.

Answer: (SHOW ANSWER)

Installing Jenkins

In this section, you use Cloud Marketplace to provision a Jenkins instance. You customize this instance to use the agent image you created in the previous section.

Go to the Cloud Marketplace solution for Jenkins.

Click Launch on Compute Engine.

Change the Machine Type field to 4 vCPUs 15 GB Memory, n1-standard-4.

Machine type selection for Jenkins deployment.

Click Deploy and wait for your Jenkins instance to finish being provisioned. When it is finished, you will see:

Jenkins has been deployed.

https://cloud.google.com/solutions/using-jenkins-for-distributed-builds-on-compute-engine#installing_jenkins

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 62

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

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 63

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

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

Answer: ([SHOW ANSWER](#)**)**

Reference:

According to this guide for setting up an HTTP (S) load balancer in GCP: The client SSL session terminates at the load balancer. Sessions between the load balancer and the instance can either be HTTPS (recommended) or HTTP.

<https://cloud.google.com/load-balancing/docs/ssl>

NEW QUESTION: 64

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

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

Answer: B,E (LEAVE A REPLY)

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

https://cloud.google.com/bigquery/pricing#flat_rate_pricing

NEW QUESTION: 65

Your managed instance group raised an alert stating that new instance creation has failed to create new instances. You need to maintain the number of running instances specified by the template to be able to process expected application traffic. What should you do?

- A. Create an instance template that contains valid syntax which will be used by the instance group. Delete any persistent disks with the same name as instance names.
- B. Create an instance template that contains valid syntax that will be used by the instance group. Verify that the instance name and persistent disk name values are not the same in the template.
- C. Verify that the instance template being used by the instance group contains valid syntax. Delete any persistent disks with the same name as instance names. Set the disks.autoDelete property to true in the instance template.
- D. Delete the current instance template and replace it with a new instance template. Verify that the instance name and persistent disk name values are not the same in the template. Set the disks.autoDelete property to true in the instance template.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/compute/docs/troubleshooting/troubleshooting-migs>

https://cloud.google.com/compute/docs/instance-templates#how_to_update_instance_templates

NEW QUESTION: 66

You have a single binary application that you want to run on Google Cloud Platform. You decided to automatically scale the application based on underlying infrastructure CPU usage. Your

organizational policies require you to use virtual machines directly. You need to ensure that the application scaling is operationally efficient and completed as quickly as possible. What should you do?

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

Answer: B ([LEAVE A REPLY](#))

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

NEW QUESTION: 67

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

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

Answer: A ([LEAVE A REPLY](#))

<https://cloud.google.com/storage/docs/audit-logs>

https://cloud.google.com/compute/docs/logging/audit-logging#audited_operations

NEW QUESTION: 68

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

- A.** Modify the existing subnet range to 172.16.20.0/24.
- B.** Create a new Secondary IP Range in the VPC and configure the VMs to use that range.

C. Create a new VPC network for the VMs. Enable VPC Peering between the VMs' VPC network and the Dataproc cluster VPC network.

D. Create a new VPC network for the VMs with a subnet of 172.32.0.0/16. Enable VPC network Peering between the Dataproc VPC network and the VMs VPC network. Configure a custom Route exchange.

Answer: A ([LEAVE A REPLY](#))

/25:

CIDR to IP Range

Result

CIDR Range 172.16.20.128/25

Netmask 255.255.255.128

Wildcard Bits 0.0.0.127

First IP 172.16.20.128

First IP (Decimal) 2886734976

Last IP 172.16.20.255

Last IP (Decimal) 2886735103

Total Host 128

CIDR

172.16.20.128/25

/24:

CIDR to IP Range

Result

CIDR Range 172.16.20.128/24

Netmask 255.255.255.0

Wildcard Bits 0.0.0.255

First IP 172.16.20.0

First IP (Decimal) 2886734848

Last IP 172.16.20.255

Last IP (Decimal) 2886735103

Total Host 256

CIDR

172.16.20.128/24

NEW QUESTION: 69

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

A. .00.0.0/0

B. 10.0.0.0/8

C. 172.16.0.0/12

D. 192.168.0.0/16

Answer: B ([LEAVE A REPLY](#))

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

NEW QUESTION: 70

Your company developed a mobile game that is deployed on Google Cloud. Gamers are connecting to the game with their personal phones over the Internet. The game sends UDP packets to update the servers about the gamers' actions while they are playing in multiplayer mode. Your game backend can scale over multiple virtual machines (VMs), and you want to expose the VMs over a single IP address. What should you do?

- A. Configure an SSL Proxy load balancer in front of the application servers.
- B. Configure an Internal UDP load balancer in front of the application servers.
- C. Configure an External HTTP(s) load balancer in front of the application servers.
- D. Configure an External Network load balancer in front of the application servers.

Answer: D ([LEAVE A REPLY](#))

cell phones are sending UDP packets and the only that can receive that type of traffic is a External Network TCP/UDP <https://cloud.google.com/load-balancing/docs/network>

<https://cloud.google.com/load-balancing/docs/choosing-load-balancer#lb-decision-tree>

NEW QUESTION: 71

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

Answer: A ([LEAVE A REPLY](#))

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: 72

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. On the Project Metadata, add that JSON as the value for the key compute-engine-service-account.
- C.** Download a JSON Private Key for the service account. On the Custom Metadata of the VM, add that JSON as the value for the key compute-engine-service-account.
- D.** Download a JSON Private Key for the service account. After creating the VM, ssh into the VM and save the JSON under `~/gcloud/compute-engine-service-account.json`.

Answer: ([SHOW ANSWER](#))

Reference:

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

Changing the service account and access scopes for an instance If you want to run the VM as a different identity, or you determine that the instance needs a different set of scopes to call the required APIs, you can change the service account and the access scopes of an existing instance. For example, you can change access scopes to grant access to a new API, or change an instance so that it runs as a service account that you created, instead of the Compute Engine default service account. However, Google recommends that you use the fine-grained IAM policies instead of relying on access scopes to control resource access for the service account. To change an instance's service account and access scopes, the instance must be temporarily stopped. To stop your instance, read the documentation for Stopping an instance. After changing the service account or access scopes, remember to restart the instance. Use one of the following methods to the change service account or access scopes of the stopped instance.

NEW QUESTION: 73

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 ([LEAVE A REPLY](#))

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: 74

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. Try deploying your group to another zone.
- B.** You have hit your instance quota for the region.
- C.** Your managed instance group's VM's are toggled to only last 1 minute in preemptible settings.
- D.** 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 (LEAVE A REPLY)

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: 75

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

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

D. Use the gsutil to rewrite the storage class for the bucket Set up Object Lifecycle management on the bucket

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 76

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. Add the parameter runtimeClassName: gvisor to the specification of your customers' Pods.

D. Use the cos_containerd image for your GKE nodes. Add a nodeSelector with the value cloud.google.com/gke-os-distribution: cos_containerd to the specification of your customers' Pods.

Answer: **C** ([LEAVE A REPLY](#))

Reference:

GKE Sandbox provides an extra layer of security to prevent untrusted code from affecting the host kernel on your cluster nodes when containers in the Pod execute unknown or untrusted code. Multi-tenant clusters and clusters whose containers run untrusted workloads are more exposed to security vulnerabilities than other clusters. Examples include SaaS providers, web-hosting providers, or other organizations that allow their users to upload and run code. When you enable GKE Sandbox on a node pool, a sandbox is created for each Pod running on a node in that node pool. In addition, nodes running sandboxed Pods are prevented from accessing other Google Cloud services or cluster metadata. Each sandbox uses its own userspace kernel. With this in mind, you can make decisions about how to group your containers into Pods, based on the level of isolation you require and the characteristics of your applications.

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

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 77

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 78

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

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

Answer: **D** ([LEAVE A REPLY](#))

NEW QUESTION: 79

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

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

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 80

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

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

Answer: C ([LEAVE A REPLY](#))

Reference:

bigquery.dataViewer role provides permissions to read the datasets metadata and list tables in the dataset as well as Read data and metadata from the datasets tables. This is exactly what we need to fulfil this requirement and follows the least privilege principle.

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

NEW QUESTION: 81

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 object. Modify the YAML file to populate the DB_PASSWORD environment variable from the Secret.
- C.** Store the database password inside a ConfigMap object. Modify the YAML file to populate the DB_PASSWORD environment variable from the ConfigMap.
- D.** 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: ([SHOW ANSWER](#))

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

NEW QUESTION: 82

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

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

Answer: A (LEAVE A REPLY)

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

NEW QUESTION: 83

You have just created a new project which will be used to deploy a globally distributed application. You will use Cloud Spanner for data storage. You want to create a Cloud Spanner instance. You want to perform the first step in preparation of creating the instance. What should you do?

- A. Grant yourself the IAM role of Cloud Spanner Admin
- B. Create a new VPC network with subnetworks in all desired regions
- C. Configure your Cloud Spanner instance to be multi-regional
- D. Enable the Cloud Spanner API

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/spanner/docs/getting-started/set-up>

NEW QUESTION: 84

You create a Deployment with 2 replicas in a Google Kubernetes Engine cluster that has a single preemptible node pool. After a few minutes, you use kubectl to examine the status of your Pod and observe that one of them is still in Pending status:

<pre>\$ kubectl get pods -l app=myapp</pre>				
NAME	READY	STATUS	RESTART	AGE
myapp-deployment-58ddb995-lp86m	0/1	Pending	0	9m
myapp-deployment-58ddb995-qjpkq	1/1	Running	0	9m

What is the most likely cause?

- A. The pending Pod's resource requests are too large to fit on a single node of the cluster.
- B. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod.
- C. The node pool is configured with a service account that does not have permission to pull the container image used by the pending Pod.
- D. The pending Pod was originally scheduled on a node that has been preempted between the creation of the Deployment and your verification of the Pods' status. It is currently being rescheduled on a new node.

Answer: (SHOW ANSWER)

The pending Pods resource requests are too large to fit on a single node of the cluster. Too many Pods are already running in the cluster, and there are not enough resources left to schedule the pending Pod. is the right answer.

When you have a deployment with some pods in running and other pods in the pending state, more often than not it is a problem with resources on the nodes. Heres a sample output of this use case. We see that the problem is with insufficient CPU on the Kubernetes nodes so we have to either enable auto-scaling or manually scale up the nodes.

NEW QUESTION: 85

Your company wants to migrate their on-premises workloads to Google Cloud. The current on-premises workloads consist of:

- * A Flask web application
- * A backend API
- * A scheduled long-running background job for ETL and reporting.

You need to keep operational costs low You want to follow Google-recommended practices to migrate these workloads to serverless solutions on Google Cloud. What should you do?

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

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 86

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: ([SHOW ANSWER](#))

```
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: 87

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

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

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 88

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

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

Answer: A ([LEAVE A REPLY](#))

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

<https://cloud.google.com/sdk/gcloud/reference/services/list#--available> Run the following command to list the enabled APIs and services in your current project:

```
gcloud services list
```

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

```
gcloud services list -available
```

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

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

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

```
gcloud services list --enabled
```

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

```
gcloud services list -available
```

NEW QUESTION: 89

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. Create an application that sends all medical images to the Pub/Sub topic

Answer: C (LEAVE A REPLY)

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: 90

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 (LEAVE A REPLY)

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: 91

You created several resources in multiple Google Cloud projects. All projects are linked to different billing accounts. To better estimate future charges, you want to have a single visual representation of all costs incurred. You want to include new cost data as soon as possible. What should you do?

- A. Configure Billing Data Export to BigQuery and visualize the data in Data Studio.
- B. Visit the Cost Table page to get a CSV export and visualize it using Data Studio.
- C. Fill all resources in the Pricing Calculator to get an estimate of the monthly cost.
- D. Use the Reports view in the Cloud Billing Console to view the desired cost information.

Answer: A (LEAVE A REPLY)

<https://cloud.google.com/billing/docs/how-to/export-data-bigquery> "Cloud Billing export to BigQuery enables you to export detailed Google Cloud billing data (such as usage, cost estimates, and pricing data) automatically throughout the day to a BigQuery dataset that you specify."

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount:

Freepdfdumps)

NEW QUESTION: 92

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

- A. Link the acquired company's projects to your company's billing account.
- B. Configure the acquired company's billing account and your company's billing account to export the billing data into the same BigQuery dataset.
- C. Migrate the acquired company's projects into your company's GCP organization. Link the migrated projects to your company's billing account.

D. Create a new GCP organization and a new billing account. Migrate the acquired company's projects and your company's projects into the new GCP organization and link the projects to the new billing account.

Answer: ([SHOW ANSWER](#))

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

<https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION: 93

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

A. Contact cloud-billing@google.com with your bank account details and request a corporate billing account for your company.

B. Create a ticket with Google Support and wait for their call to share your credit card details over the phone.

C. In the Google Platform Console, go to the Resource Manager and move all projects to the root Organization.

D. In the Google Cloud Platform Console, create a new billing account and set up a payment method.

Answer: ([SHOW ANSWER](#))

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

<https://cloud.google.com/billing/docs/concepts>

<https://cloud.google.com/resource-manager/docs/project-migration>

NEW QUESTION: 94

Your company wants to standardize the creation and management of multiple Google Cloud resources using Infrastructure as Code. You want to minimize the amount of repetitive code needed to manage the environment. What should you do?

A. Create a bash script that contains all requirement steps as `gcloud` commands

B. Develop templates for the environment using Cloud Deployment Manager

C. Use `curl` in a terminal to send a REST request to the relevant Google API for each individual resource.

D. Use the Cloud Console interface to provision and manage all related resources

Answer: **B** ([LEAVE A REPLY](#))

You can use Google Cloud Deployment Manager to create a set of Google Cloud resources and manage them as a unit, called a deployment. For example, if your team's development environment needs two virtual machines (VMs) and a BigQuery database, you can define these resources in a configuration file, and use Deployment Manager to create, change, or delete these resources. You can make the configuration file part of your team's code repository, so that anyone can create the same environment with consistent results. <https://cloud.google.com/deployment-manager/docs/quickstart>

NEW QUESTION: 95

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: (SHOW ANSWER)

<https://cloud.google.com/architecture/identity/migrating-consumer-accounts>

NEW QUESTION: 96

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

- A. Deploy the application to Cloud functions. Suffix the version number in the functions name.
- B. Deploy the application to Cloud Run. Use gradual rollouts for traffic splitting.
- C. Deploy the application to App Engine. For each new version, create a new service.
- D. Deploy the application to Google Kubernetes Engine. Use Anthos Service Mesh for traffic splitting.

Answer: B (LEAVE A REPLY)

NEW QUESTION: 97

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

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

Answer: C (LEAVE A REPLY)

NEW QUESTION: 98

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

- A. Deploy the container on Cloud Run for Anthos, and set the minimum number of instances to zero
- B. Deploy the container on Cloud Run (fully managed), and set the minimum number of instances to zero.

- C.** Deploy the container on App Engine flexible environment with autoscaling. and set the value min_instances to zero in the app yaml
- D.** Deploy the container on App Engine flexible environment with manual scaling, and set the value instances to zero in the app yaml

Answer: B ([LEAVE A REPLY](#))

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

NEW QUESTION: 99

You are running a web application on Cloud Run for a few hundred users. Some of your users complain that the initial web page of the application takes much longer to load than the following pages. You want to follow Google's recommendations to mitigate the issue. What should you do?

- A.** Set the maximum number of instances for your Cloud Run service to 100.
- B.** Set the minimum number of instances for your Cloud Run service to 3.
- C.** Set the concurrency number to 1 for your Cloud Run service.
- D.** Update your web application to use the protocol HTTP/2 instead of HTTP/1.1

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 100

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

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

Answer: D ([LEAVE A REPLY](#))

Reference:

<https://cloud.google.com/appengine/docs/standard/go/config/appref>

"App Engine calculates the number of instances necessary to serve your current application traffic based on scaling settings such as target_cpu_utilization and target_throughput_utilization. Setting min_idle_instances specifies the number of instances to run in addition to this calculated number. For example, if App Engine calculates that 5 instances are necessary to serve traffic, and min_idle_instances is set to 2, App Engine will run 7 instances (5, calculated based on traffic, plus 2 additional per min_idle_instances)." Automatic scaling creates dynamic instances based on request rate, response latencies, and other application metrics. However, if you specify the number of minimum idle instances, that specified number of instances run as resident instances while any additional instances are dynamic.

Ref: <https://cloud.google.com/appengine/docs/standard/python/how-instances-are-managed>

NEW QUESTION: 101

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

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 102

Your team is running an on-premises ecommerce application. The application contains a complex set of microservices written in Python, and each microservice is running on Docker containers. Configurations are injected by using environment variables. You need to deploy your current application to a serverless Google Cloud cloud solution. What should you do?

- A. Use your existing codebase and deploy each service as a separate Cloud Run Use the same configurations as on-premises.
- B. Use your existing CI/CD pipeline Use the generated Docker images and deploy them to Cloud Run. Update the configurations and the required endpoints.
- C. Use your existing continuous integration and delivery (CI/CD) pipeline. Use the generated Docker images and deploy them to Cloud Function. Use the same configuration as on-premises.
- D. Use the existing codebase and deploy each service as a separate Cloud Function Update the configurations and the required endpoints.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 103

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

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

Answer: A ([LEAVE A REPLY](#))

App engine is a regional service, which means the infrastructure that runs your app(s) is located in a specific region and is managed by Google to be redundantly available across all the zones within that region. Once an app engine deployment is created in a region, it cant be changed. The only way is to create a new project and create an App Engine instance in europe-west3, send all user traffic to this instance and delete the app engine instance in us-central.

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

NEW QUESTION: 104

For analysis purposes, you need to send all the logs from all of your Compute Engine instances to a BigQuery dataset called platform-logs. You have already installed the Stackdriver Logging agent on all the instances. You want to minimize cost. What should you do?

- A.** 1. Give the BigQuery Data Editor role on the platform-logs dataset to the service accounts used by your instances.2. Update your instances' metadata to add the following value: logs-destination: bq://platform-logs.
- B.** 1. In Stackdriver Logging, create a logs export with a Cloud Pub/Sub topic called logs as a sink.2. Create a Cloud Function that is triggered by messages in the logs topic.3. Configure that Cloud Function to drop logs that are not from Compute Engine and to insert Compute Engine logs in the platform-logs dataset.
- C.** 1. In Stackdriver Logging, create a filter to view only Compute Engine logs.2. Click Create Export.3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.
- D.** 1. Create a Cloud Function that has the BigQuery User role on the platform-logs dataset.2. Configure this Cloud Function to create a BigQuery Job that executes this query:INSERT INTO dataset.platform-logs (timestamp, log)SELECT timestamp, log FROM compute.logsWHERE timestamp > DATE_SUB(CURRENT_DATE(), INTERVAL 1 DAY)3. Use Cloud Scheduler to trigger this Cloud Function once a day.

Answer: C (LEAVE A REPLY)

- 1. In Stackdriver Logging, create a filter to view only Compute Engine logs. 2. Click Create Export.
- 3. Choose BigQuery as Sink Service, and the platform-logs dataset as Sink Destination.

NEW QUESTION: 105

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

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

Answer: B (LEAVE A REPLY)

Navigate to Stackdriver Logging and select resource.labels.project_id=*. is not right.

Log entries are held in Stackdriver Logging for a limited time known as the retention period which is 30 days (default configuration). After that, the entries are deleted. To keep log entries longer, you need to export them outside of Stackdriver Logging by configuring log sinks.

Ref: <https://cloud.google.com/blog/products/gcp/best-practices-for-working-with-google-cloud-audit-logging> Configure a Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery. Configure the table expiration to 60 days. is not right.

While this works, it makes no sense to use Cloud Scheduler job to read from Stackdriver and store the logs in BigQuery when Google provides a feature (export sinks) that does exactly the same thing and works out of the box. Ref:

https://cloud.google.com/logging/docs/export/configure_export_v2 Create a Stackdriver Logging Export with a Sink destination to Cloud Storage. Create a lifecycle rule to delete objects after 60 days. is not right.

You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2 Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud organization. Ref:

https://cloud.google.com/logging/docs/export/aggregated_sinks Either way, we now have the data in Cloud Storage, but querying logs information from Cloud Storage is harder than Querying information from BigQuery dataset. For this reason, we should prefer Big Query over Cloud Storage.

Create a Stackdriver Logging Export with a Sink destination to a BigQuery dataset. Configure the table expiration to 60 days. is the right answer.

You can export logs by creating one or more sinks that include a logs query and an export destination. Supported destinations for exported log entries are Cloud Storage, BigQuery, and Pub/Sub. Ref: https://cloud.google.com/logging/docs/export/configure_export_v2 Sinks are limited to exporting log entries from the exact resource in which the sink was created: a Google Cloud project, organization, folder, or billing account. If it makes it easier to exporting from all projects of an organization, you can create an aggregated sink that can export log entries from all the projects, folders, and billing accounts of a Google Cloud organization. Ref:

https://cloud.google.com/logging/docs/export/aggregated_sinks Either way, we now have the data in a BigQuery Dataset. Querying information from a Big Query dataset is easier and quicker than analyzing contents in Cloud Storage bucket. As our requirement is to Quickly analyze the log contents, we should prefer Big Query over Cloud Storage.

Also, You can control storage costs and optimize storage usage by setting the default table expiration for newly created tables in a dataset. If you set the property when the dataset is created, any table created in the dataset is deleted after the expiration period. If you set the property after the dataset is created, only new tables are deleted after the expiration period. For example, if you set the default table expiration to 7 days, older data is automatically deleted after 1 week. Ref:

<https://cloud.google.com/bigquery/docs/best-practices-storage>

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: A,B (LEAVE A REPLY)

(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>

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 107

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

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


Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 108

You are building a pipeline to process time-series data

a. 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 (LEAVE A REPLY)

Reference:

<https://cloud.google.com/blog/products/data-analytics/handling-duplicate-data-in-streaming-pipeline-using-pubsub-dataflow>

<https://cloud.google.com/bigtable/docs/schema-design-time-series>

NEW QUESTION: 109

Your company has a single sign-on (SSO) identity provider that supports Security Assertion Markup Language (SAML) integration with service providers. Your company has users in Cloud Identity. You would like users to authenticate using your company's SSO provider. What should you do?

- A. In Cloud Identity, set up SSO with Google as an identity provider to access custom SAML apps.
- B. In Cloud Identity, set up SSO with a third-party identity provider with Google as a service provider.
- C. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Mobile & Desktop Apps.
- D. Obtain OAuth 2.0 credentials, configure the user consent screen, and set up OAuth 2.0 for Web Server Applications.

Answer: B (LEAVE A REPLY)

https://support.google.com/cloudidentity/answer/6262987?hl=en&ref_topic=7558767

NEW QUESTION: 110

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

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

Answer: A (LEAVE A REPLY)

performs a peering between the two VPC's (the statement makes sure that this option is feasible since it clearly specifies that there is no overlapping between the ip ranges of both vpc's), deploy the LoadBalancer as internal with the annotation, and configure the endpoint so that the compute engine instance can access the application internally, that is, without the need to have a public ip

at any time and therefore, without the need to go outside the google network. The traffic, therefore, never crosses the public internet.

<https://medium.com/pablo-perez/k8s-externaltrafficpolicy-local-or-cluster-40b259a19404>

<https://cloud.google.com/kubernetes-engine/docs/how-to/internal-load-balancing> clients in a VPC network connected to the LoadBalancer network using VPC Network Peering can also access the Service

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

NEW QUESTION: 111

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 (LEAVE A REPLY)

<https://gtseres.medium.com/using-service-accounts-across-projects-in-gcp-cf9473fef8f0#:~:text=Go%20to%20the%20destination%20project,Voila!>

NEW QUESTION: 112

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 (LEAVE A REPLY)

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: 113

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

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

Answer: A ([LEAVE A REPLY](#))

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

NEW QUESTION: 114

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

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

Answer: C ([LEAVE A REPLY](#))

`gcloud config configurations list` can help check for the existing configurations and `activate` can help switch to the configuration.

`gcloud config configurations list` lists existing named configurations

`gcloud config configurations activate` activates an existing named configuration Obtains access credentials for your user account via a web-based authorization flow. When this command completes successfully, it sets the active account in the current configuration to the account specified. If no configuration exists, it creates a configuration named `default`.

NEW QUESTION: 115

You have been asked to set up Object Lifecycle Management for objects stored in storage buckets. The objects are written once and accessed frequently for 30 days. After 30 days, the objects are not read again unless there is a special need. The object should be kept for three years, and you need to minimize cost. What should you do?

- A.** Set up a policy that uses Nearline storage for 30 days and then moves to Archive storage for three years.
- B.** Set up a policy that uses Standard storage for 30 days and then moves to Archive storage for three years.

C. Set up a policy that uses Nearline storage for 30 days, then moves the Coldline for one year, and then moves to Archive storage for two years.

D. Set up a policy that uses Standard storage for 30 days, then moves to Coldline for one year, and then moves to Archive storage for two years.

Answer: ([SHOW ANSWER](#))

The key to understand the requirement is : "The objects are written once and accessed frequently for 30 days" Standard Storage Standard Storage is best for data that is frequently accessed ("hot" data) and/or stored for only brief periods of time.

Archive Storage

Archive Storage is the lowest-cost, highly durable storage service for data archiving, online backup, and disaster recovery. Unlike the "coldest" storage services offered by other Cloud providers, your data is available within milliseconds, not hours or days. Archive Storage is the best choice for data that you plan to access less than once a year.

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

NEW QUESTION: 116

An external member of your team needs list access to compute images and disks in one of your projects. You want to follow Google-recommended practices when you grant the required permissions to this user. What should you do?

A. Create a custom role, and add all the required compute.disks.list and compute, images.list permissions as includedPermissions. Grant the custom role to the user at the project level.

B. Create a custom role based on the Compute Storage Admin role. Exclude unnecessary permissions from the custom role. Grant the custom role to the user at the project level.

C. Grant the Compute Storage Admin role at the project level.

D. Create a custom role based on the Compute Image User role Add the compute.disks, list to the includedPermissions field Grant the custom role to the user at the project level

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 117

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

A. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Coldline Storage.

B. Add a bucket lifecycle rule that archives data with newer versions after 30 days to Nearline Storage.

C. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Coldline Storage.

D. Add a bucket lifecycle rule that archives data from regional storage after 30 days to Nearline Storage.

Answer: ([**SHOW ANSWER**](#)**)**

Reference:

Nearline Storage is ideal for data you plan to read or modify on average once per month or less. And this option archives just the noncurrent versions which is what we want to do.

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

NEW QUESTION: 118

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 ([**LEAVE A REPLY**](#)**)**

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

NEW QUESTION: 119

You are using Data Studio to visualize a table from your data warehouse that is built on top of BigQuery. Data is appended to the data warehouse during the day. At night, the daily summary is recalculated by overwriting the table. You just noticed that the charts in Data Studio are broken, and you want to analyze the problem. What should you do?

- A. Use the BigQuery interface to review the nightly Job and look for any errors
- B. Review the Error Reporting page in the Cloud Console to find any errors.
- C. In Cloud Logging create a filter for your Data Studio report
- D. Use the open source CLI tool. Snapshot Debugger, to find out why the data was not refreshed correctly.

Answer: D ([**LEAVE A REPLY**](#)**)**

Cloud Debugger helps inspect the state of an application, at any code location, without stopping or slowing down the running app // <https://cloud.google.com/debugger/docs>

NEW QUESTION: 120

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: ([SHOW ANSWER](#))



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: 121

Your company runs one batch process in an on-premises server that takes around 30 hours to complete. The task runs monthly, can be performed offline, and must be restarted if interrupted. You want to migrate this workload to the cloud while minimizing cost. What should you do?

- A. Migrate the workload to a Compute Engine Preemptible VM.
- B. Migrate the workload to a Google Kubernetes Engine cluster with Preemptible nodes.
- C. Migrate the workload to a Compute Engine VM. Start and stop the instance as needed.
- D. Create an Instance Template with Preemptible VMs On. Create a Managed Instance Group from the template and adjust Target CPU Utilization. Migrate the workload.

Answer: D ([LEAVE A REPLY](#))

Install the workload in a compute engine VM, start and stop the instance as needed, because as per the question the VM runs for 30 hours, process can be performed offline and should not be interrupted, if interrupted we need to restart the batch process again. Preemptible VMs are cheaper, but they will not be available beyond 24hrs, and if the process gets interrupted the preemptible VM will restart.

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: <https://www.actual4test.com/Associate-Cloud-Engineer>

NEW QUESTION: 122

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

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

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 123

The DevOps group in your organization needs full control of Compute Engine resources in your development project. However, they should not have permission to create or update any other resources in the project. You want to follow Google's recommendations for setting permissions for the DevOps group. What should you do?

- A. Grant the basic role roles/viewer and the predefined role roles/compute.admin to the DevOps group.
- B. Create a custom role at the folder level and grant all compute. instanceAdmin. * permissions to the role Grant the custom role to the DevOps group.
- C. Create an IAM policy and grant all compute. instanceAdmin." permissions to the policy Attach the policy to the DevOps group.
- D. Grant the basic role roles/editor to the DevOps group.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 124

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

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

Answer: C (LEAVE A REPLY)

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

NEW QUESTION: 125

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

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

Answer: B (LEAVE A REPLY)

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

NEW QUESTION: 126

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

Answer: B (LEAVE A REPLY)

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: 127

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

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

Answer: B (LEAVE A REPLY)

You can generate Windows passwords using either the Google Cloud Console or the `gcloud` command-line tool. This option uses the right syntax to reset the windows password.

`gcloud compute reset-windows-password windows-instance`

Ref: <https://cloud.google.com/compute/docs/instances/windows/creating-passwords-for-windows-instances#gcloud>

Valid Associate-Cloud-Engineer Dumps shared by Actual4test.com for Helping Passing Associate-Cloud-Engineer Exam! Actual4test.com now offer the **newest Associate-Cloud-Engineer exam dumps**, the Actual4test.com Associate-Cloud-Engineer exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com Associate-Cloud-Engineer dumps with Test Engine here: https://www.actual4test.com/Associate-Cloud-Engineer_examcollection.html (335 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)