Free Questions for Associate-Cloud-Engineer by braindumpscollection

Shared by Bradshaw on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You want to host your video encoding software on Compute Engine. Your user base is growing rapidly, and users need to be able 3 to encode their videos at any time without interruption or CPU limitations. You must ensure that your encoding solution is highly available, and you want to follow Google-recommended practices to automate operations. What should you do?

Options:

A- Deploy your solution on multiple standalone Compute Engine instances, and increase the number of existing instances wnen CPU utilization on Cloud Monitoring reaches a certain threshold.

B- Deploy your solution on multiple standalone Compute Engine instances, and replace existing instances with high-CPU instances when CPU utilization on Cloud Monitoring reaches a certain threshold.

C- Deploy your solution to an instance group, and increase the number of available instances whenever you see high CPU utilization in Cloud Monitoring.

D- Deploy your solution to an instance group, and set the autoscaling based on CPU utilization.

Answer:

D

Explanation:

Instance groups are collections of virtual machine (VM) instances that you can manage as a single entity. Instance groups can help you simplify the management of multiple instances, reduce operational costs, and improve the availability and performance of your applications. Instance groups support autoscaling, which automatically adds or removes instances from the group based on increases or decreases in load. Autoscaling helps your applications gracefully handle increases in traffic and reduces cost when the need for resources is lower. You can set the autoscaling policy based on CPU utilization, load balancing capacity, Cloud Monitoring metrics, or a queue-based workload. In this case, since the video encoding software is CPU-intensive, setting the autoscaling based on CPU utilization is the best option to ensure high availability and optimal performance.Reference:

Instance groups

Autoscaling groups of instances

Question 2

Question Type: MultipleChoice

You are deploying a web application using Compute Engine. You created a managed instance group (MIG) to host the application. You want to follow Google-recommended practices to implement a secure and highly available solution. What should you do?

Options:

A- Use SSL proxy load balancing for the MIG and an A record in your DNS private zone with the load balancer's IP address.

B- Use SSL proxy load balancing for the MIG and a CNAME record in your DNS public zone with the load balancer's IP address.

C- Use HTTP(S) load balancing for the MIG and a CNAME record in your DNS private zone with the load balancer's IP address.

D- Use HTTP(S) load balancing for the MIG and an A record in your DNS public zone with the load balancer's IP address.

Answer:

D

Explanation:

HTTP(S) load balancing is a Google-recommended practice for distributing web traffic across multiple regions and zones, and providing high availability, scalability, and security for web applications. It supports both IPv4 and IPv6 addresses, and can handle SSL/TLS termination and encryption. It also integrates with Cloud CDN, Cloud Armor, and Cloud Identity-Aware Proxy for enhanced performance and protection. A MIG can be used as a backend service for HTTP(S) load balancing, and can automatically scale and heal the VM instances that host the web application.

To configure DNS for HTTP(S) load balancing, you need to create an A record in your DNS public zone with the load balancer's IP address. This will map your domain name to the load balancer's IP address, and allow users to access your web application using the domain name. A CNAME record is not recommended, as it can cause latency and DNS resolution issues. A private zone is not suitable, as it is only visible within your VPC network, and not to the public internet.

HTTP(S) Load Balancing documentation

Setting up DNS records for HTTP(S) load balancing

Choosing a load balancer

Question 3

Question Type: MultipleChoice

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

Options:

A- In the Google Cloud console, validate which SSH keys have been stored as project-wide keys.

B- Navigate to Identity-Aware Proxy and check the permissions for these resources.

C- Enable Audit logs on the 1AM & admin page for all resources, and validate the results.

D- Use the gcloud projects get-iam-policy command to view the current role assignments.

Answer:

D

Explanation:

The gcloud projects get-iam-policy command displays the IAM policy for a project, which includes the roles and members assigned to those roles. The Project Owner role grants full access to all resources and actions in the project. By using this command, you can review who has been granted this role and make any necessary changes.Reference:

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

2: gcloud projects get-iam-policy | Cloud SDK Documentation | Google Cloud

3: Understanding roles | Cloud IAM Documentation | Google Cloud

Question 4

Question Type: MultipleChoice

You are responsible for a web application on Compute Engine. You want your support team to be notified automatically if users experience high latency for at least 5 minutes. You need a Google-recommended solution with no development cost. What should you do?

Options:

A- Create an alert policy to send a notification when the HTTP response latency exceeds the specified threshold.

B- Implement an App Engine service which invokes the Cloud Monitoring API and sends a notification in case of anomalies.

C- Use the Cloud Monitoring dashboard to observe latency and take the necessary actions when the response latency exceeds the specified threshold.

D- Export Cloud Monitoring metrics to BigQuery and use a Looker Studio dashboard to monitor your web applications latency.

Answer:

A

Explanation:

https://cloud.google.com/monitoring/alerts#alerting-example

Question 5

Question Type: MultipleChoice

All development (dev) teams in your organization are located in the United States. Each dev team has its own Google Cloud project. You want to restrict access so that each dev team can only create cloud resources in the United States (US). What should you do?

Options:

A- Create a folder to contain all the dev projects Create an organization policy to limit resources in US locations.

B- Create an organization to contain all the dev projects. Create an Identity and Access Management (IAM) policy to limit the resources in US regions.

C- Create an Identity and Access Management <IAM) policy to restrict the resources locations in the US. Apply the policy to all dev projects.

D- Create an Identity and Access Management (IAM)policy to restrict the resources locations in all dev projects. Apply the policy to all dev roles.

Answer:

C

Question 6

Question Type: **MultipleChoice**

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

Options:

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 a Private Catalog tor 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.

D- Create an infrastructure-as-code template tor all resources in the project by using Terraform. and deploy that template to a new project in your organization. Delete the protect from the startup's Google Cloud organization.

Answer:

A

Question 7

Question Type: MultipleChoice

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

Options:

A- Enable Private Service Access on the Cloud Storage Bucket.

B- Add slorage.googleapis.com to the list of restricted services in a VPC Service Controls perimeter and add your project to the list to protected projects.

C- Enable Private Google Access on the subnet within the custom VPC.

D- Deploy a Cloud NAT instance and route the traffic to the dedicated IP address of the Cloud Storage bucket.

Answer:

A