Free Questions for Professional-Cloud-Architect by go4braindumps

Shared by Whitaker on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: **MultipleChoice**

You team needs to create a Google Kubernetes Engine (GKE) cluster to host a newly built application that requires access to third-party services on the internet. Your company does not allow any Compute Engine instance to have a public IP address on Google Cloud. You need to create a deployment strategy that adheres to these guidelines. What should you do?

Options:

A- Create a Compute Engine instance, and install a NAT Proxy on the instance. Configure all workloads on GKE to pass through this proxy to access third-party services on the Internet

B- Configure the GKE cluster as a private cluster, and configure Cloud NAT Gateway for the cluster subnet

C- Configure the GKE cluster as a route-based cluster. Configure Private Google Access on the Virtual Private Cloud (VPC)

D- Configure the GKE cluster as a private cluster. Configure Private Google Access on the Virtual Private Cloud (VPC)

Answer:

B

Explanation:

A Cloud NAT gateway can perform NAT for nodes and Pods in a private cluster, which is a type of VPC-native cluster. The Cloud NAT gateway must be configured to apply to at least the following subnet IP address ranges for the subnet that your cluster uses:

Subnet primary IP address range (used by nodes)

Subnet secondary IP address range used for Pods in the cluster

Subnet secondary IP address range used for Services in the cluster

The simplest way to provide NAT for an entire private cluster is to configure a Cloud NAT gateway to apply to all of the cluster's subnet's IP address ranges.

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

Question 2

Question Type: **MultipleChoice**

Your company has an application running on a deployment in a GKE cluster. You have a separate cluster for development, staging and production. You have discovered that the team is able to deploy a Docker image to the production cluster without first testing the deployment in development and then staging. You want to allow the team to have autonomy but want to prevent this from happening. You want a Google Cloud solution that can be implemented quickly with minimal effort. What should you do?

Options:

A- Create a Kubernetes admission controller to prevent the container from starting if it is not approved for usage in the given environment

B- Configure a Kubernetes lifecycle hook to prevent the container from starting if it is not approved for usage in the given environment

C- Implement a corporate policy to prevent teams from deploying Docker image to an environment unless the Docker image was tested in an earlier environment

D- Configure the binary authorization policies for the development, staging and production clusters. Create attestations as part of the continuous integration pipeline"

Answer:

D

Explanation:

https://cloud.google.com/architecture/prep-kubernetes-engine-for-prod#binary-authorization

The most common Binary Authorization use cases involve attestations. An attestation certifies that a specific image has completed a previous stage, as described previously. You configure the Binary Authorization policy to verify the attestation before allowing the image to be deployed. At deploy time, instead of redoing activities that were completed in earlier stages, Binary Authorization only needs to verify the attestation. https://cloud.google.com/binary-authorization/docs/overview

Question 3

Question Type: MultipleChoice

Your company is planning to upload several important files to Cloud Storage. After the upload is completed, they want to verify that the upload content is identical to what they have on- premises. You want to minimize the cost and effort of performing this check. What should you do?

Options:

A- 1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil cp to download the uploaded files
3) Use Linux diff to compare the content of the files

B- 1) Use gsutil -m to upload all the files to Cloud Storage.
2) Develop a custom Java application that computes CRC32C hashes
3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4) Compare the hashes

C- 1) Use Linux shasum to compute a digest of files you want to upload
2) Use gsutil -m to upload all the files to the Cloud Storage
3) Use gsutil cp to download the uploaded files
4) Use Linux shasum to compute a digest of the downloaded files 5.Compre the hashes

D- 1) Use gsutil -m to upload all the files to Cloud Storage.
2) Use gsutil hash -c FILE_NAME to generate CRC32C hashes of all on-premises files

3) Use gsutil ls -L gs://[YOUR_BUCKET_NAME] to collect CRC32C hashes of the uploaded files
4) Compare the hashes

Answer:

D

Explanation:

https://cloud.google.com/storage/docs/gsutil/commands/hash

Question 4

Question Type: **MultipleChoice**

Your company has a stateless web API that performs scientific calculations. The web API runs on a single Google Kubernetes Engine (GKE) cluster. The cluster is currently deployed in us-central1. Your company has expanded to offer your API to customers in Asi

a. You want to reduce the latency for the users in Asia. What should you do?

Options:

A- Use a global HTTP(s) load balancer with Cloud CDN enabled

B- Create a second GKE cluster in asia-southeast1, and expose both API's using a Service of type Load Balancer. Add the public Ips to the Cloud DNS zone

C- Increase the memory and CPU allocated to the application in the cluster

D- Create a second GKE cluster in asia-southeast1, and use kubemci to create a global HTTP(s) load balancer

Answer:

D

Explanation:

https://cloud.google.com/kubernetes-engine/docs/concepts/multi-cluster-ingress#how_works

https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress

https://cloud.google.com/blog/products/gcp/how-to-deploy-geographically-distributed-services-on-kubernetes-engine-with-kubemci

Question 5

Question Type: **MultipleChoice**

You are migrating third-party applications from optimized on-premises virtual machines to Google Cloud. You are unsure about the optimum CPU and memory options. The application have a consistent usage patterns across multiple weeks. You want to optimize resource usage for the lowest cost. What should you do?

Options:

A- Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console

B- Create an App Engine flexible environment, and deploy the third party application using a Docker file and a custom runtime. Set CPU and memory options similar to your application's current on-premises virtual machine in the app.yaml file.

C- Create an instance template with the smallest available machine type, and use an image of the third party application taken from the current on-premises virtual machine. Create a managed instance group that uses average CPU to autoscale the number of instances in the group. Modify the average CPU utilization threshold to optimize the number of instances running.

D- Create multiple Compute Engine instances with varying CPU and memory options. Install the cloud monitoring agent and deploy the third-party application on each of them. Run a load test with high traffic levels on the application and use the results to determine the optimal settings.

Answer:

A

Explanation:

Create a Compute engine instance with CPU and Memory options similar to your application's current on-premises virtual machine. Install the cloud monitoring agent, and deploy the third party application. Run a load with normal traffic levels on third party application and follow the Rightsizing Recommendations in the Cloud Console

https://cloud.google.com/migrate/compute-engine/docs/4.9/concepts/planning-a-migration/cloud-instance-rightsizing?hl=en

Question 6

Question Type: MultipleChoice

You have deployed an application on Anthos clusters (formerly Anthos GKE). According to the SRE practices at your company you need to be alerted if the request latency is above a certain threshold for a specified amount of time. What should you do?

Options:

A- Enable the Cloud Trace API on your project and use Cloud Monitoring Alerts to send an alert based on the Cloud Trace metrics

B- Configure Anthos Config Management on your cluster and create a yaml file that defines the SLO and alerting policy you want to deploy in your cluster

C- Use Cloud Profiler to follow up the request latency. Create a custom metric in Cloud Monitoring based on the results of Cloud Profiler, and create an Alerting Policy in case this metric exceeds the threshold

D- Install Anthos Service Mesh on your cluster. Use the Google Cloud Console to define a Service Level Objective (SLO)

Answer:

D

Explanation:

https://cloud.google.com/service-mesh/docs/overview

https://cloud.google.com/service-mesh/docs/observability/slo-overview

Question 7

Question Type: **MultipleChoice**

Your company has a support ticketing solution that uses App Engine Standard. The project that contains the App Engine application already has a Virtual Private Cloud(VPC) network fully connected to the company's on-premises environment through a Cloud VPN tunnel. You want to enable App Engine application to communicate with a database that is running in the company's on-premises environment. What should you do?

Options:

A- Configure private services access

B- Configure private Google access for on-premises hosts only

C- Configure serverless VPC access

D- Configure private Google access

Answer:

C

Explanation:

https://cloud.google.com/appengine/docs/standard/python3/connecting-vpc

https://cloud.google.com/appengine/docs/flexible/python/using-third-party-databases#on_premises

Question 8

Question Type: MultipleChoice

You are designing a Data Warehouse on Google Cloud and want to store sensitive data in BigQuery. Your company requires you to generate encryption keys outside of Google Cloud. You need to implement a solution. What should you do?

Options:

A- Generate a new key in Cloud Key Management Service (Cloud KMS). Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a BigQuery dataset.

B- Generate a new key in Cloud Key Management Service (Cloud KMS). Create a dataset in BigQuery using the customer-managed key option and select the created key

C- Import a key in Cloud KMS. Store all data in Cloud Storage using the customer-managed key option and select the created key. Set up a Dataflow pipeline to decrypt the data and to store it in a new BigQuery dataset.

D- Import a key in Cloud KMS. Create a dataset in BigQuery using the customer-supplied key option and select the created key.

Answer:

D

Explanation:

https://cloud.google.com/bigquery/docs/customer-managed-encryption

Question 9

Question Type: **MultipleChoice**

Your organization has stored sensitive data in a Cloud Storage bucket. For regulatory reasons, your company must be able to rotate the encryption key used to encrypt the data in the bucket. The data will be processed in Dataproc. You want to follow Google-recommended practices for security What should you do?

Options:

A- Create a key with Cloud Key Management Service (KMS) Encrypt the data using the encrypt method of Cloud KMS.

B- Create a key with Cloud Key Management Service (KMS). Set the encryption key on the bucket to the Cloud KMS key.

C- Generate a GPG key pair. Encrypt the data using the GPG key. Upload the encrypted data to the bucket.

D- Generate an AES-256 encryption key. Encrypt the data in the bucket using the customer-supplied encryption keys feature.

Answer:

B

Explanation:

https://cloud.google.com/storage/docs/encryption/using-customer-managed-keys#add-object-key

Question 10

Question Type: **MultipleChoice**

Your company is developing a web-based application. You need to make sure that production deployments are linked to source code commits and are fully auditable. What should you do?

Options:

A- Make sure a developer is tagging the code commit with the date and time of commit

B- Make sure a developer is adding a comment to the commit that links to the deployment.

C- Make the container tag match the source code commit hash.

D- Make sure the developer is tagging the commits with :latest

Answer:

C

20%
DISCOUNT