



Free Questions for Professional-Cloud-Architect by certsinside

Shared by Pruitt on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You need to design a solution for global load balancing based on the URL path being requested. You need to ensure operations reliability and end-to-end in-transit encryption based on Google best practices.

What should you do?

Options:

- A- Create a cross-region load balancer with URL Maps.
- B- Create an HTTPS load balancer with URL maps.
- C- Create appropriate instance groups and instances. Configure SSL proxy load balancing.
- D- Create a global forwarding rule. Configure SSL proxy balancing.

Answer:

B

Explanation:

Reference <https://cloud.google.com/load-balancing/docs/https/url-map>

Question 2

Question Type: MultipleChoice

A development team at your company has created a dockerized HTTPS web application. You need to deploy the application on Google Kubernetes Engine (GKE) and make sure that the application scales automatically.

How should you deploy to GKE?

Options:

- A-** Use the Horizontal Pod Autoscaler and enable cluster autoscaling. Use an Ingress resource to loadbalance the HTTPS traffic.
- B-** Use the Horizontal Pod Autoscaler and enable cluster autoscaling on the Kubernetes cluster. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.
- C-** Enable autoscaling on the Compute Engine instance group. Use an Ingress resource to load balance the HTTPS traffic.
- D-** Enable autoscaling on the Compute Engine instance group. Use a Service resource of type LoadBalancer to load-balance the HTTPS traffic.

Answer:

B

Explanation:

<https://cloud.google.com/kubernetes-engine/docs/tutorials/http-balancer>

<https://cloud.google.com/kubernetes-engine/docs/concepts/network-overview#ext-lb>

Question 3

Question Type: MultipleChoice

Your customer runs a web service used by e-commerce sites to offer product recommendations to users. The company has begun experimenting with a machine learning model on Google Cloud Platform to improve the quality of results.

What should the customer do to improve their model's results over time?

Options:

- A-** Export Cloud Machine Learning Engine performance metrics from Stackdriver to BigQuery, to be used to analyze the efficiency of the model.
- B-** Build a roadmap to move the machine learning model training from Cloud GPUs to Cloud TPUs, which offer better results.
- C-** Monitor Compute Engine announcements for availability of newer CPU architectures, and deploy the model to them as soon as they are available for additional performance.
- D-** Save a history of recommendations and results of the recommendations in BigQuery, to be used as training data.

Answer:

D

Explanation:

<https://cloud.google.com/solutions/building-a-serverless-ml-model>

Question 4

Question Type: MultipleChoice

All compute Engine instances in your VPC should be able to connect to an Active Directory server on specific ports. Any other traffic emerging from your instances is not allowed. You want to enforce this using VPC firewall rules.

How should you configure the firewall rules?

Options:

- A-** Create an egress rule with priority 1000 to deny all traffic for all instances. Create another egress rule with priority 100 to allow the Active Directory traffic for all instances.
- B-** Create an egress rule with priority 100 to deny all traffic for all instances. Create another egress rule with priority 1000 to allow the Active Directory traffic for all instances.
- C-** Create an egress rule with priority 1000 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 100 to block all traffic for all instances.
- D-** Create an egress rule with priority 100 to allow the Active Directory traffic. Rely on the implied deny egress rule with priority 1000 to block all traffic for all instances.

Answer:

B

Explanation:

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

Question 5

Question Type: MultipleChoice

You have an App Engine application that needs to be updated. You want to test the update with production traffic before replacing the current application version.

What should you do?

Options:

- A-** Deploy the update using the Instance Group Updater to create a partial rollout, which allows for canary testing.
- B-** Deploy the update as a new version in the App Engine application, and split traffic between the new and current versions.
- C-** Deploy the update in a new VPC, and use Google's global HTTP load balancing to split traffic between the update and current applications.
- D-** Deploy the update as a new App Engine application, and use Google's global HTTP load balancing to split traffic between the new and current applications.

Answer:

B

Explanation:

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

Question 6

Question Type: MultipleChoice

Your architecture calls for the centralized collection of all admin activity and VM system logs within your project.

How should you collect these logs from both VMs and services?

Options:

- A-** All admin and VM system logs are automatically collected by Stackdriver.
- B-** Stackdriver automatically collects admin activity logs for most services. The Stackdriver Logging agent must be installed on each instance to collect system logs.
- C-** Launch a custom syslogd compute instance and configure your GCP project and VMs to forward all logs to it.
- D-** Install the Stackdriver Logging agent on a single compute instance and let it collect all audit and access logs for your environment.

Answer:

B

Explanation:

<https://cloud.google.com/logging/docs/agent/default-logs>

Question 7

Question Type: MultipleChoice

You are using Cloud CDN to deliver static HTTP(S) website content hosted on a Compute Engine instance group. You want to improve the cache hit ratio.

What should you do?

Options:

- A-** Customize the cache keys to omit the protocol from the key.
- B-** Shorten the expiration time of the cached objects.
- C-** Make sure the HTTP(S) header "Cache-Region" points to the closest region of your users.
- D-** Replicate the static content in a Cloud Storage bucket. Point CloudCDN toward a load balancer on that bucket.

Answer:

A

Explanation:

Reference [https://cloud.google.com/cdn/docs/bestpractices#](https://cloud.google.com/cdn/docs/bestpractices#using_custom_cache_keys_to_improve_cache_hit_ratio)

[using_custom_cache_keys_to_improve_cache_hit_ratio](https://cloud.google.com/cdn/docs/bestpractices#using_custom_cache_keys_to_improve_cache_hit_ratio)

To Get Premium Files for Professional-Cloud-Architect Visit

<https://www.p2pexams.com/products/professional-cloud-architect>

For More Free Questions Visit

<https://www.p2pexams.com/google/pdf/professional-cloud-architect>

