Free Questions for Professional-Cloud-Network-Engineer

Shared by Aguirre on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You built a web application with several containerized microservices. You want to run those microservices on Cloud Run. You must also ensure that the services are highly available to your customers with low latency. What should you do?

Options:

A- Deploy the Cloud Run services to multiple availability zones. Create a global TCP load balancer. Add the Cloud Run endpoints to its backend service.
B- Deploy the Cloud Run services to multiple regions. Create serverless network endpoint groups (NEGs) that point to the services. Create a global HTTPS load balancer, and attach the serverless NEGs as backend services of the load balancer.
C- Deploy the Cloud Run services to multiple availability zones. Create Cloud Endpoints that point to the services. Create a global HTTPS load balancer, and attach the Cloud Endpoints to its backend
D- Deploy the Cloud Run services to multiple regions. Configure a round-robin A record in Cloud DNS.

Answer:

B

Question 2

Question Type: MultipleChoice

You need to create a GKE cluster in an existing VPC that is accessible from on-premises. You must meet the following requirements:

IP ranges for pods and services must be as small as possible.

The nodes and the master must not be reachable from the internet.

You must be able to use kubectl commands from on-premises subnets to manage the cluster.

How should you create the GKE cluster?

Options:

A- * Create a private cluster that uses VPC advanced routes.
* Set the pod and service ranges as /24.
* Set up a network proxy to access the master.
B- * Create a VPC-native GKE cluster using GKE-managed IP ranges.
* Set the pod IP range as /21 and service IP range as /24.
* Set up a network proxy to access the master.
C- * Create a VPC-native GKE cluster using user-managed IP ranges.
* Enable a GKE cluster network policy, set the pod and service ranges as /24.
* Set up a network proxy to access the master.
* Enable master authorized networks.
D- * Create a VPC-native GKE cluster using user-managed IP ranges.
* Enable privateEndpoint on the cluster master.
* Set the pod and service ranges as /24.
* Set up a network proxy to access the master.
* Enable master authorized networks.

Answer:

D

Explanation:

Creating GKE private clusters with network proxies for controller access When you create a GKE private cluster with a private cluster controller endpoint, the cluster's controller node is inaccessible from the public internet, but it needs to be accessible for administration. By default, clusters can access the controller through its private endpoint, and authorized networks can be defined within the VPC network. To access the controller from on-premises or another VPC network, however, requires additional steps. This is because the VPC network that hosts the controller is owned by Google and cannot be accessed from resources connected through another VPC network peering connection, Cloud VPN or Cloud Interconnect.
https://cloud.google.com/solutions/creating-kubernetes-engine-private-clusters-with-net-proxies

Question 3

Question Type: MultipleChoice

You are designing a hub-and-spoke network architecture for your company's cloud-based environment. You need to make sure that all spokes are peered with the hub. The spokes must use the hub's virtual appliance for internet access.

The virtual appliance is configured in high-availability mode with two instances using an internal load balancer with IP address 10.0.0.5. What should you do?

Options:

A- Create a default route in the hub VPC that points to IP address 10.0.0.5.
Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway.
Export the custom routes in the hub.
Import the custom routes in the spokes.
B- Create a default route in the hub VPC that points to IP address 10.0.0.5.
Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway.
Export the custom routes in the hub. Import the custom routes in the spokes.
Delete the default internet gateway route of the spokes.
C- Create two default routes in the hub VPC that point to the next hop instances of the virtual appliances.
Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway.
Export the custom routes in the hub. Import the custom routes in the spokes.
D- Create a default route in the hub VPC that points to IP address 10.0.0.5.
Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway.
Create a new route in the spoke VPC that points to IP address 10.0.0.5.

Answer:

B

Question 4

Question Type: MultipleChoice

You have enabled HTTP(S) load balancing for your application, and your application developers have reported that HTTP(S) requests are not being distributed correctly to your Compute Engine Virtual Machine instances. You want to find data about how the request are being distributed.

Which two methods can accomplish this? (Choose two.)

Options:

A- On the Load Balancer details page of the GCP Console, click on the Monitoring tab, select your backend service, and look at the graphs.
B- In Stackdriver Error Reporting, look for any unacknowledged errors for the Cloud Load

Balancers service.
C- In Stackdriver Monitoring, select Resources > Metrics Explorer and search for https/request_bytes_count metric.
D- In Stackdriver Monitoring, select Resources > Google Cloud Load Balancers and review the Key Metrics graphs in the dashboard.
E- In Stackdriver Monitoring, create a new dashboard and track the https/backend_request_count metric for the load balancer.

Answer:

A, E

Question 5

Question Type: MultipleChoice

Your company's security team wants to limit the type of inbound traffic that can reach your web servers to protect against security threats. You need to configure the firewall rules on the web servers within your Virtual Private Cloud (VPC) to handle HTTP and HTTPS web traffic for TCP only. What should you do?

Options:

A- Create an allow on match ingress firewall rule with the target tag ''web-server'' to allow all IP addresses for TCP port 80.
B- Create an allow on match egress firewall rule with the target tag ''web-server'' to allow all IP addresses for TCP port 80.
C- Create an allow on match ingress firewall rule with the target tag ''web-server'' to allow all IP addresses for TCP ports 80 and 443.
D- Create an allow on match egress firewall rule with the target tag ''web-server' to allow web server IP addresses for TCP ports 60 and 443.

Answer:

C

Question 6

Question Type: MultipleChoice

You are designing a new global application using Compute Engine instances that will be exposed by a global HTTP(S) load balancer. You need to secure your application from distributed denial-of-

service and application layer (layer 7) attacks. What should you do?

Options:

A- Configure VPC Service Controls and create a secure perimeter. Define fine-grained perimeter controls and enforce that security posture across your Google Cloud services and projects.
B- Configure a Google Cloud Armor security policy in your project, and attach it to the backend service to secure the application.
C- Configure VPC firewall rules to protect the Compute Engine instances against distributed denial-of-service attacks.
D- Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at the organization level.

Answer:

C

Question 7

Question Type: MultipleChoice

You have an HA VPN connection with two tunnels running in active/passive mode between your Virtual Private Cloud (VPC) and on-premises network. Traffic over the connection has recently increased from 1 gigabit per second (Gbps) to 4 Gbps, and you notice that packets are being dropped. You need to configure your VPN connection to Google Cloud to support 4 Gbps. What should you do?

Options:

A- Configure the remote autonomous system number (ASN) to 4096.
B- Configure a second Cloud Router to scale bandwidth in and out of the VPC.
C- Configure the maximum transmission unit (MTU) to its highest supported value.
D- Configure a second set of active/passive VPN tunnels.

Answer:

D

To Get Premium Files for Professional-Cloud-Network-Engineer Visit
https://www.p2pexams.com/products/professional-cloud-network-engineer

For More Free Questions Visit
https://www.p2pexams.com/google/pdf/professional-cloud-network-engineer

