Free Questions for Professional-Cloud-Network-Engineer by certscare

Shared by Savage on 20-10-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You have the following private Google Kubernetes Engine (GKE) cluster deployment:

```
gcloud container clusters describe customer-1-cluster --zone us-central1-c

...

clusterIpv4Cidr: 192.168.36.0/24
endpoint: 192.168.38.2
ipAllocationPolicy:
  clusterIpv4Cidr: 192.168.36.0/24
  clusterIpv4CidrBlock: 192.168.36.0/24
  clusterSecondaryRangeName: customer-1-pods
  servicesIpv4Cidr: 192.168.37.0/24
  servicesIp4CidrBlock: 192.168.37.0/24
  servicesSecondaryRangeName: customer-1-svc
  useIpAliases: true

...

masterAuthorizedNetworksConfig:

...

privateClusterConfig:
  enablePrivateEndpoint: true
  enablePrivateNodes: true
  masterIpv4CidrBlock: 192.168.38.0/28
  privateEndpoint: 192.168.38.2
  publicEndpoint: 35.224.37.17

...

servicesIpv4Cidr: 192.162.37.0/24

...
```

You have a virtual machine (VM) deployed in the same VPC in the subnetwork kubernetes-management with internal IP address 192.168.40 2/24 and no external IP address assigned. You need to communicate with the cluster master using kubectl. What should you do?

Options:

A- Add the network 192.168.40.0/24 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2.

B- Add the network 192.168.38.0/28 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2

C- Add the network 192.168.36.0/24 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2

D- Add an external IP address to the VM, and add this IP address in the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 35.224.37.17.

Answer:

A

Question 2

Question Type: **MultipleChoice**

Your company has a Virtual Private Cloud (VPC) with two Dedicated Interconnect connections in two different regions: us-west1 and us-east1. Each Dedicated Interconnect connection is attached to a Cloud Router in its respective region by a VLAN attachment. You need to configure a high availability failover path. By default, all ingress traffic from the on-premises environment should flow to the VPC using the us-west1 connection. If us-west1 is unavailable, you want traffic to be rerouted to us-east1. How should you configure the multi-exit discriminator (MED) values to enable this failover path?

Options:

A- Use regional routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1

B- Use global routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1

C- Use regional routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1

D- Use global routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1

Answer:

A

Question 3

Question Type: **MultipleChoice**

One instance in your VPC is configured to run with a private IP address only. You want to ensure that even if this instance is deleted, its current private IP address will not be automatically assigned to a different instance.

In the GCP Console, what should you do?

Options:

A- Assign a public IP address to the instance.

B- Assign a new reserved internal IP address to the instance.

C- Change the instance's current internal IP address to static.

D- Add custom metadata to the instance with key internal-address and value reserved.

Answer:

C

Explanation:

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#reservenewip Since here https://cloud.google.com/compute/docs/ip-addresses/reserve-static-internal-ip-address#reservenewip it is written that 'automatically allocated or an unused address from an existing subnet'.

Question 4

Question Type: MultipleChoice

You are using the gcloud command line tool to create a new custom role in a project by coping a predefined role. You receive this error message:

INVALID_ARGUMENT: Permission resourcemanager.projects.list is not valid

What should you do?

Options:

A- Add the resourcemanager.projects.get permission, and try again.

B- Try again with a different role with a new name but the same permissions.

C- Remove the resourcemanager.projects.list permission, and try again.

D- Add the resourcemanager.projects.setIamPolicy permission, and try again.

Answer:

C

Question 5

Question Type: MultipleChoice

You need to centralize the Identity and Access Management permissions and email distribution for the WebServices Team as efficiently as possible.

What should you do?

Options:

A- Create a Google Group for the WebServices Team.

B- Create a G Suite Domain for the WebServices Team.

C- Create a new Cloud Identity Domain for the WebServices Team.

D- Create a new Custom Role for all members of the WebServices Team.

Answer:

A

Question 6

Question Type: MultipleChoice

You want to use Partner Interconnect to connect your on-premises network with your VPC. You already have an Interconnect partner.

What should you first?

Options:

A- Log in to your partner's portal and request the VLAN attachment there.

B- Ask your Interconnect partner to provision a physical connection to Google.

C- Create a Partner Interconnect type VLAN attachment in the GCP Console and retrieve the pairing key.

D- Run gcloud compute interconnect attachments partner update / -- region <region> --admin-enabled.

Answer:

B

Explanation:

https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview?hl=En#provisioning 'To provision a Partner Interconnect connection with a service provider, you start by connecting your on-premises network to a supported service provider. Work with the service provider to establish connectivity.

Question 7

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 8

Question Type: MultipleChoice

In your company, two departments with separate GCP projects (code-dev and data-dev) in the same organization need to allow full cross-communication between all of their virtual machines in GCP. Each department has one VPC in its project and wants full control over their network. Neither department intends to recreate its existing computing resources. You want to implement a solution that minimizes cost.

Which two steps should you take? (Choose two.)

Options:

A- Connect both projects using Cloud VPN.

B- Connect the VPCs in project code-dev and data-dev using VPC Network Peering.

C- Enable Shared VPC in one project (e. g., code-dev), and make the second project (e. g., data-dev) a service project.

D- Enable firewall rules to allow all ingress traffic from all subnets of project code-dev to all instances in project data-dev, and vice versa.

E- Create a route in the code-dev project to the destination prefixes in project data-dev and use nexthop as the default gateway, and vice versa.

Answer:

B, D

Question 9

Question Type: MultipleChoice

You have created a firewall with rules that only allow traffic over HTTP, HTTPS, and SSH ports. While testing, you specifically try to reach the server over multiple ports and protocols; however, you do not see any denied connections in the firewall logs. You want to resolve the issue.

What should you do?

Options:

A- Enable logging on the default Deny Any Firewall Rule.

B- Enable logging on the VM Instances that receive traffic.

C- Create a logging sink forwarding all firewall logs with no filters.

D- Create an explicit Deny Any rule and enable logging on the new rule.

Answer:

D

Explanation:

https://cloud.google.com/vpc/docs/firewall-rules-logging#egress_deny_example

You can only enable Firewall Rules Logging for rules in a Virtual Private Cloud (VPC) network. Legacy networks are not supported. Firewall Rules Logging only records TCP and UDP connections. Although you can create a firewall rule applicable to other protocols, you cannot log their connections. You cannot enable Firewall Rules Logging for the implied deny ingress and implied allow egress rules. Log entries are written from the perspective of virtual machine (VM) instances. Log entries are only created if a firewall rule has logging enabled and if the rule applies to traffic sent to or from the VM. Entries are created according to the connection logging limits on a best effort basis. The number of connections that can be logged in a given interval is based on the machine type. Changes to firewall rules can be viewed in VPC audit logs. https://cloud.google.com/vpc/docs/firewall-rules-logging#specifications

Question 10

Question Type: MultipleChoice

You need to define an address plan for a future new GKE cluster in your VPC. This will be a VPC native cluster, and the default Pod IP range allocation will be used. You must pre-provision all the needed VPC subnets and their respective IP address ranges before cluster

creation. The cluster will initially have a single node, but it will be scaled to a maximum of three nodes if necessary. You want to allocate the minimum number of Pod IP addresses.

Which subnet mask should you use for the Pod IP address range?

Options:

A- /21

B- /22

C- /23

D- /25

Answer:

B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/alias-ips#cluster_sizing_secondary_range_pods

https://cloud.google.com/kubernetes-engine/docs/how-to/flexible-pod-cidr

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#defaults_limits

Question 11

Question Type: MultipleChoice

Your on-premises data center has 2 routers connected to your GCP through a VPN on each router. All applications are working correctly; however, all of the traffic is passing across a single VPN instead of being load-balanced across the 2 connections as desired.

During troubleshooting you find:

* Each on-premises router is configured with the same ASN.

* Each on-premises router is configured with the same routes and priorities.

* Both on-premises routers are configured with a VPN connected to a single Cloud Router.

* The VPN logs have no-proposal-chosen lines when the VPNs are connecting.

* BGP session is not established between one on-premises router and the Cloud Router.

What is the most likely cause of this problem?

Options:

A- One of the VPN sessions is configured incorrectly.

B- A firewall is blocking the traffic across the second VPN connection.

C- You do not have a load balancer to load-balance the network traffic.

D- BGP sessions are not established between both on-premises routers and the Cloud Router.

Answer:

A

Explanation:

If the VPN logs show a no-proposal-chosen error, this error indicates that Cloud VPN and your peer VPN gateway were unable to agree on a set of ciphers. For IKEv1, the set of ciphers must match exactly. For IKEv2, there must be at least one common cipher proposed by each gateway. Make sure that you use supported ciphers to configure your peer VPN gateway. https://cloud.google.com/network-connectivity/docs/vpn/support/troubleshooting#:~:text=If%20the%20VPN%20logs%20show,of%20ciphers%20must%20match%20exactly.&text=Mak