

Google.Professional Cloud Network Engineer.vJun-2024.by.Enqua.122q

Number: Professional Cloud Network Engineer
Passing Score: 800
Time Limit: 120
File Version: 12.0

Certification: Professional Cloud Network Engineer
Certification Full Name: Professional Cloud Network Engineer



Exam A

QUESTION 1

You recently configured Google Cloud Armor security policies to manage traffic to your application.

You discover that Google Cloud Armor is incorrectly blocking some traffic to your application. You need to identify the web application firewall (WAF) rule that is incorrectly blocking traffic. What should you do?

- A. Enable firewall logs, and view the logs in Firewall Insights.
- B. Enable HTTP(S) Load Balancing logging with sampling rate equal to 1, and view the logs in CloudLogging.
- C. Enable VPC Flow Logs, and view the logs in Cloud Logging.
- D. Enable Google Cloud Armor audit logs, and view the logs on the Activity page in the Google Cloud Console.

Correct Answer: A

Section:

QUESTION 2

You are the Organization Admin for your company. One of your engineers is responsible for setting up multiple host projects across multiple folders and sharing subnets with service projects. You need to enable the engineer's Identity and Access Management (IAM) configuration to complete their task in the fewest number of steps. What should you do?

- A. Set up the engineer with Compute Shared VPC Admin IAM role at the folder level.
- B. Set up the engineer with Compute Shared VPC Admin IAM role at the organization level.
- C. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the folder level.
- D. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the organization level.

Correct Answer: B

Section:

QUESTION 3

You recently deployed Compute Engine instances in regions us-west1 and us-east1 in a Virtual Private Cloud (VPC) with default routing configurations. Your company security policy mandates that virtual machines (VMs) must not have public IP addresses attached to them. You need to allow your instances to fetch updates from the internet while preventing external access. What should you do?

- A. Create a Cloud NAT gateway and Cloud Router in both us-west1 and us-east1.
- B. Create a single global Cloud NAT gateway and global Cloud Router in the VPC.
- C. Change the instances' network interface external IP address from None to Ephemeral.
- D. Create a firewall rule that allows egress to destination 0.0.0.0/0.

Correct Answer: A

Section:

QUESTION 4

Your company has provisioned 2000 virtual machines (VMs) in the private subnet of your Virtual Private Cloud (VPC) in the us-east1 region. You need to configure each VM to have a minimum of 128 TCP connections to a public repository so that users can download software updates and packages over the internet. You need to implement a Cloud NAT gateway so that the VMs are able to perform outbound NAT to the internet. You must ensure that all VMs can simultaneously connect to the public repository and download software updates and packages. Which two methods can you use to accomplish this? (Choose two.)

- A. Configure the NAT gateway in manual allocation mode, allocate 2 NAT IP addresses, and update the minimum number of ports per VM to 256.

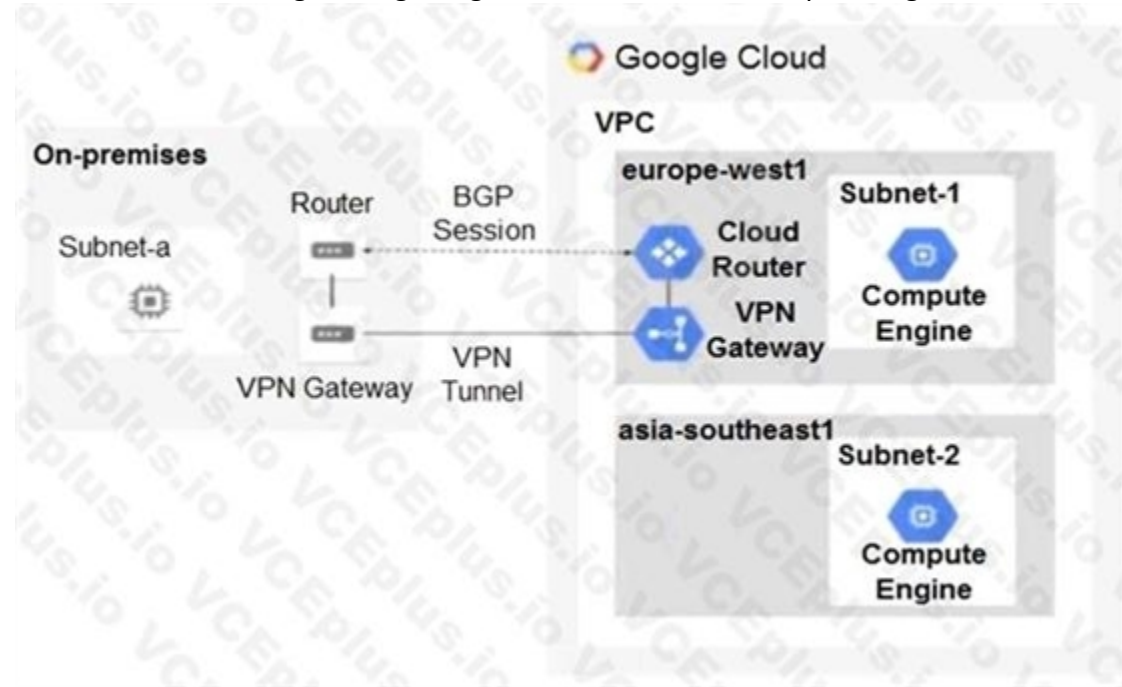
- B. Create a second Cloud NAT gateway with the default minimum number of ports configured per VM to 64.
- C. Use the default Cloud NAT gateway's NAT proxy to dynamically scale using a single NAT IP address.
- D. Use the default Cloud NAT gateway to automatically scale to the required number of NAT IP addresses, and update the minimum number of ports per VM to 128.
- E. Configure the NAT gateway in manual allocation mode, allocate 4 NAT IP addresses, and update the minimum number of ports per VM to 128.

Correct Answer: A, B

Section:

QUESTION 5

You have the following routing design. You discover that Compute Engine instances in Subnet-2 in the asia-southeast1 region cannot communicate with compute resources on-premises. What should you do?



- A. Configure a custom route advertisement on the Cloud Router.
- B. Enable IP forwarding in the asia-southeast1 region.
- C. Change the VPC dynamic routing mode to Global.
- D. Add a second Border Gateway Protocol (BGP) session to the Cloud Router.

Correct Answer: C

Section:

QUESTION 6

You are designing a hybrid cloud environment for your organization. Your Google Cloud environment is interconnected with your on-premises network using Cloud HA VPN and Cloud Router. The Cloud Router is configured with the default settings. Your on-premises DNS server is located at 192.168.20.88 and is protected by a firewall, and your Compute Engine resources are located at 10.204.0.0/24. Your Compute Engine resources need to resolve on-premises private hostnames using the domain corp.altostrat.com while still resolving Google Cloud hostnames. You want to follow Google-recommended practices. What should you do?

- A. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24. Set a custom route advertisement on the Cloud Router for 10.204.0.0/24
- B. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 35.199.192.0/19 Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.
- C. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24.

Modify the /etc/resolv.conf file on your Compute Engine instances to point to 192.168.20.88

- D. Create a private zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com.
Configure DNS Server Policies and create a policy with Alternate DNS servers to 192.168.20.88.
Configure your on-premises firewall to accept traffic from 35.199.192.0/19.
Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Correct Answer: D

Section:

QUESTION 7

Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with on-premises connectivity already in place. You are deploying a new application using Google Kubernetes Engine (GKE), which must be accessible only from the same VPC network and on-premises locations.

You must ensure that the GKE control plane is exposed to a predefined list of on-premises subnets through private connectivity only. What should you do?

- A. Create a GKE private cluster with a private endpoint for the control plane. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Routers. Configure authorized networks to specify the desired on-premises subnets.
- B. Create a GKE private cluster with a public endpoint for the control plane. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Routers.
- C. Create a GKE private cluster with a private endpoint for the control plane. Configure authorized networks to specify the desired on-premises subnets.
- D. Create a GKE public cluster. Configure authorized networks to specify the desired on-premises subnets.

Correct Answer: C

Section:

QUESTION 8

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?

- 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 (NEG) 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.

Correct Answer: B

Section:

QUESTION 9

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?

- 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.

Correct Answer: D

Section:

QUESTION 10

You recently deployed two network virtual appliances in us-central1. Your network appliances provide connectivity to your on-premises network, 10.0.0.0/8. You need to configure the routing for your Virtual Private Cloud (VPC). Your design must meet the following requirements:

All access to your on-premises network must go through the network virtual appliances.

Allow on-premises access in the event of a single network virtual appliance failure.

Both network virtual appliances must be used simultaneously.

Which method should you use to accomplish this?

- A. Configure two routes for 10.0.0.0/8 with different priorities, each pointing to separate network virtual appliances.
- B. Configure an internal HTTP(S) load balancer with the two network virtual appliances as backends.
Configure a route for 10.0.0.0/8 with the internal HTTP(S) load balancer as the next hop.
- C. Configure a network load balancer for the two network virtual appliances. Configure a route for 10.0.0.0/8 with the network load balancer as the next hop.
- D. Configure an internal TCP/UDP load balancer with the two network virtual appliances as backends.
Configure a route for 10.0.0.0/8 with the internal load balancer as the next hop.

Correct Answer: C

Section:

QUESTION 11

You are responsible for enabling Private Google Access for the virtual machine (VM) instances in your Virtual Private Cloud (VPC) to access Google APIs. All VM instances have only a private IP address and need to access Cloud Storage.

You need to ensure that all VM traffic is routed back to your on-premises data center for traffic scrubbing via your existing Cloud Interconnect connection. However, VM traffic to Google APIs should remain in the VPC. What should you do?

- A. Delete the default route in your VPC.
Create a private Cloud DNS zone for googleapis.com, create a CNAME for *.googleapis.com to restricted googleapis.com, and create an A record for restricted googleapis.com that resolves to the addresses in 199.36.153.4/30.
Create a static route in your VPC for the range 199.36.153.4/30 with the default internet gateway as the next hop.
- B. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).
Create a public Cloud DNS zone with a CNAME for *.google.com to private googleapis.com, create a CNAME for * googleapis.com to private googleapis.com, and create an A record for Private googleapis.com that resolves to the addresses in 199.36.153.8/30.
Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.
- C. Configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP) with a lower priority (MED) than the default VPC route.
Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googleapis.com to private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30.
Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.
- D. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).
Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googleapis.com to Private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30.
Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.

Correct Answer: C

Section:

QUESTION 12

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?

- 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.

Correct Answer: C

Section:

QUESTION 13

Your organization's security policy requires that all internet-bound traffic return to your on-premises data center through HA VPN tunnels before egressing to the internet, while allowing virtual machines (VMs) to leverage private Google APIs using private virtual IP addresses 199.36.153.4/30.

You need to configure the routes to enable these traffic flows. What should you do?

- A. Configure a custom route 0.0.0.0/0 with a priority of 500 whose next hop is the default internet gateway. Configure another custom route 199.36.153.4/30 with priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.
- B. Configure a custom route 0.0.0.0/0 with a priority of 1000 whose next hop is the internet gateway. Configure another custom route 199.36.153.4/30 with a priority of 500 whose next hop is the VPN tunnel back to the on-premises data center.
- C. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 1000. Configure a custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the default internet gateway.
- D. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 500. Configure another custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.

Correct Answer: A

Section:

QUESTION 14

Your company has defined a resource hierarchy that includes a parent folder with subfolders for each department. Each department defines their respective project and VPC in the assigned folder and has the appropriate permissions to create Google Cloud firewall rules. The VPCs should not allow traffic to flow between them. You need to block all traffic from any source, including other VPCs, and delegate only the intra-VPC firewall rules to the respective departments.

What should you do?

- A. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 0.
- B. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 1000.
- C. Create two hierarchical firewall policies per department's folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to allow, and another lower-priority rule that blocks traffic from any other source.
- D. Create two hierarchical firewall policies per department's folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to goto_next, and another lower-priority rule that blocks traffic from any other source.

Correct Answer: B

Section:

QUESTION 15

You have two Google Cloud projects in a perimeter to prevent data exfiltration. You need to move a third project inside the perimeter; however, the move could negatively impact the existing environment. You need to validate the impact of the change. What should you do?

- A. Enable Firewall Rules Logging inside the third project.
- B. Modify the existing VPC Service Controls policy to include the new project in dry run mode.
- C. Monitor the Resource Manager audit logs inside the perimeter.



D. Enable VPC Flow Logs inside the third project, and monitor the logs for negative impact.

Correct Answer: B

Section:

QUESTION 16

You are configuring an HA VPN connection between your Virtual Private Cloud (VPC) and onpremises network. The VPN gateway is named VPN_GATEWAY_1. You need to restrict VPN tunnels created in the project to only connect to your on-premises VPN public IP address: 203.0.113.1/32.

What should you do?

- A. Configure a firewall rule accepting 203.0.113.1/32, and set a target tag equal to VPN_GATEWAY_1.
- B. Configure the Resource Manager constraint constraints/compute.restrictVpnPeerIPs to use an allowList consisting of only the 203.0.113.1/32 address.
- C. Configure a Google Cloud Armor security policy, and create a policy rule to allow 203.0.113.1/32.
- D. Configure an access control list on the peer VPN gateway to deny all traffic except 203.0.113.1/32, and attach it to the primary external interface.

Correct Answer: B

Section:

QUESTION 17

Your company has recently installed a Cloud VPN tunnel between your on-premises data center and your Google Cloud Virtual Private Cloud (VPC). You need to configure access to the Cloud Functions API for your on-premises servers.

The configuration must meet the following requirements:

Certain data must stay in the project where it is stored and not be exfiltrated to other projects.

Traffic from servers in your data center with RFC 1918 addresses do not use the internet to access Google Cloud APIs.

All DNS resolution must be done on-premises.

The solution should only provide access to APIs that are compatible with VPC Service Controls.

What should you do?

- A. Create an A record for private.googleapis.com using the 199.36.153.8/30 address range.
Create a CNAME record for *.googleapis.com that points to the A record.
Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.
Remove the default internet gateway from the VPC where your Cloud VPN tunnel terminates.
- B. Create an A record for restricted.googleapis.com using the 199.36.153.4/30 address range.
Create a CNAME record for *.googleapis.com that points to the A record.
Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.
Configure your on-premises firewalls to allow traffic to the restricted.googleapis.com addresses.
- C. Create an A record for restricted.googleapis.com using the 199.36.153.4/30 address range.
Create a CNAME record for *.googleapis.com that points to the A record.
Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.
Remove the default internet gateway from the VPC where your Cloud VPN tunnel terminates.
- D. Create an A record for private.googleapis.com using the 199.36.153.8/30 address range.
Create a CNAME record for *.googleapis.com that points to the A record.
Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.
Configure your on-premises firewalls to allow traffic to the private.googleapis.com addresses.

Correct Answer: C

Section:

QUESTION 18

You need to configure a Google Kubernetes Engine (GKE) cluster. The initial deployment should have 5 nodes with the potential to scale to 10 nodes. The maximum number of Pods per node is 8. The number of services could grow from 100 to up to 1024. How should you design the IP schema to optimally meet this requirement?

- A. Configure a /28 primary IP address range for the node IP addresses. Configure a /25 secondary IP range for the Pods. Configure a /22 secondary IP range for the Services.
- B. Configure a /28 primary IP address range for the node IP addresses. Configure a /25 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.
- C. Configure a /28 primary IP address range for the node IP addresses. Configure a /28 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.
- D. Configure a /28 primary IP address range for the node IP addresses. Configure a /24 secondary IP range for the Pods. Configure a /22 secondary IP range for the Services.

Correct Answer: A

Section:

QUESTION 19

You created a VPC network named Retail in auto mode. You want to create a VPC network named Distribution and peer it with the Retail VPC. How should you configure the Distribution VPC?

- A. Create the Distribution VPC in auto mode. Peer both the VPCs via network peering.
- B. Create the Distribution VPC in custom mode. Use the CIDR range 10.0.0.0/9. Create the necessary subnets, and then peer them via network peering.
- C. Create the Distribution VPC in custom mode. Use the CIDR range 10.128.0.0/9. Create the necessary subnets, and then peer them via network peering.
- D. Rename the default VPC as "Distribution" and peer it via network peering.

Correct Answer: B

Section:

Explanation:

<https://cloud.google.com/vpc/docs/vpc#ip-ranges>



QUESTION 20

You are using a third-party next-generation firewall to inspect traffic. You created a custom route of 0.0.0.0/0 to route egress traffic to the firewall. You want to allow your VPC instances without public IP addresses to access the BigQuery and Cloud Pub/Sub APIs, without sending the traffic through the firewall.

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

- A. Turn on Private Google Access at the subnet level.
- B. Turn on Private Google Access at the VPC level.
- C. Turn on Private Services Access at the VPC level.
- D. Create a set of custom static routes to send traffic to the external IP addresses of Google APIs and services via the default internet gateway.
- E. Create a set of custom static routes to send traffic to the internal IP addresses of Google APIs and services via the default internet gateway.

Correct Answer: A, D

Section:

Explanation:

<https://cloud.google.com/vpc/docs/private-access-options#pga> Private Google Access VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the _external IP addresses_ of Google APIs and services.

QUESTION 21

All the instances in your project are configured with the custom metadata enable-oslogin value set to FALSE and to block project-wide SSH keys. None of the instances are set with any SSH key, and no project-wide SSH keys have been configured. Firewall rules are set up to allow SSH sessions from any IP address range. You want to SSH into one instance.

What should you do?

- A. Open the Cloud Shell SSH into the instance using `gcloud compute ssh`.

- B. Set the custom metadata enable-oslogin to TRUE, and SSH into the instance using a third-party tool like putty or ssh.
- C. Generate a new SSH key pair. Verify the format of the private key and add it to the instance. SSH into the instance using a third-party tool like putty or ssh.
- D. Generate a new SSH key pair. Verify the format of the public key and add it to the project. SSH into the instance using a third-party tool like putty or ssh.

Correct Answer: A

Section:

QUESTION 22

You work for a university that is migrating to GCP.

These are the cloud requirements:

- On-premises connectivity with 10 Gbps
- Lowest latency access to the cloud
- Centralized Networking Administration Team New departments are asking for on-premises connectivity to their projects. You want to deploy the most cost-efficient interconnect solution for connecting the campus to Google Cloud.

What should you do?

- A. Use Shared VPC, and deploy the VLAN attachments and Interconnect in the host project.
- B. Use Shared VPC, and deploy the VLAN attachments in the service projects. Connect the VLAN attachment to the Shared VPC's host project.
- C. Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects' Interconnects.
- D. Use standalone projects and deploy the VLAN attachments and Interconnects in each of the individual projects.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnects-other-projects>

Using Cloud Interconnect with Shared VPC You can use Shared VPC to share your VLAN attachment in a project with other VPC networks. Choosing Shared VPC is preferable if you need to create many projects and would like to prevent individual project owners from managing their connectivity back to your on-premises network. In this scenario, the host project contains a common Shared VPC network usable by VMs in service projects. Because VMs in the service projects use this network, Service Project Admins don't need to create other VLAN attachments or Cloud Routers in the service projects. In this scenario, you must create VLAN attachments and Cloud Routers for a Cloud Interconnect connection only in the Shared VPC host project. The combination of a VLAN attachment and its associated Cloud Router are unique to a given Shared VPC network.

https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiplenetworks-access-same-attachment#using_with

<https://cloud.google.com/vpc/docs/shared-vpc>

QUESTION 23

You have deployed a new internal application that provides HTTP and TFTP services to on-premises hosts. You want to be able to distribute traffic across multiple Compute Engine instances, but need to ensure that clients are sticky to a particular instance across both services.

Which session affinity should you choose?

- A. None
- B. Client IP
- C. Client IP and protocol
- D. Client IP, port and protocol

Correct Answer: B

Section:

QUESTION 24

You created a new VPC network named Dev with a single subnet. You added a firewall rule for the network Dev to allow HTTP traffic only and enabled logging. When you try to log in to an instance in the subnet via Remote Desktop Protocol, the login fails. You look for the Firewall rules logs in Stackdriver Logging, but you do not see any entries for blocked traffic. You want to see the logs for blocked traffic.

What should you do?

- A. Check the VPC flow logs for the instance.
- B. Try connecting to the instance via SSH, and check the logs.
- C. Create a new firewall rule to allow traffic from port 22, and enable logs.
- D. Create a new firewall rule with priority 65500 to deny all traffic, and enable logs.

Correct Answer: D

Section:

Explanation:

Ingress packets in VPC Flow Logs are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs. We want to see the logs for blocked traffic so we have to look for them in firewall logs.

https://cloud.google.com/vpc/docs/flow-logs#key_properties

QUESTION 25

You are trying to update firewall rules in a shared VPC for which you have been assigned only Network Admin permissions. You cannot modify the firewall rules. Your organization requires using the least privilege necessary. Which level of permissions should you request?

- A. Security Admin privileges from the Shared VPC Admin.
- B. Service Project Admin privileges from the Shared VPC Admin.
- C. Shared VPC Admin privileges from the Organization Admin.
- D. Organization Admin privileges from the Organization Admin.

Correct Answer: A

Section:

Explanation:

A Shared VPC Admin can define a Security Admin by granting an IAM member the Security Admin (compute.securityAdmin) role to the host project. Security Admins manage firewall rules and SSL certificates.



QUESTION 26

You want to create a service in GCP using IPv6.

What should you do?

- A. Create the instance with the designated IPv6 address.
- B. Configure a TCP Proxy with the designated IPv6 address.
- C. Configure a global load balancer with the designated IPv6 address.
- D. Configure an internal load balancer with the designated IPv6 address.

Correct Answer: C

Section:

Explanation:

<https://cloud.google.com/load-balancing/docs/load-balancing-overview> mentions to use global loadbalancer for IPv6 termination.

QUESTION 27

You want to deploy a VPN Gateway to connect your on-premises network to GCP. You are using a non BGP-capable on-premises VPN device. You want to minimize downtime and operational overhead when your network grows. The device supports only IKEv2, and you want to follow Google recommended practices.

What should you do?

- A. • Create a Cloud VPN instance. • Create a policy-based VPN tunnel per subnet. • Configure the appropriate local and remote traffic selectors to match your local and remote networks. • Create the appropriate static routes.

- B. • Create a Cloud VPN instance. • Create a policy-based VPN tunnel. • Configure the appropriate local and remote traffic selectors to match your local and remote networks. • Configure the appropriate static routes.
- C. • Create a Cloud VPN instance. • Create a route-based VPN tunnel. • Configure the appropriate local and remote traffic selectors to match your local and remote networks. • Configure the appropriate static routes.
- D. • Create a Cloud VPN instance. • Create a route-based VPN tunnel. • Configure the appropriate local and remote traffic selectors to 0.0.0.0/0. • Configure the appropriate static routes.

Correct Answer: B

Section:

Explanation:

https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-staticvpns#creating_a_gateway_and_tunnel

QUESTION 28

Your company just completed the acquisition of Altostrat (a current GCP customer). Each company has a separate organization in GCP and has implemented a custom DNS solution. Each organization will retain its current domain and host names until after a full transition and architectural review is done in one year. These are the assumptions for both GCP environments.

- Each organization has enabled full connectivity between all of its projects by using Shared VPC.
- Both organizations strictly use the 10.0.0.0/8 address space for their instances, except for bastion hosts (for accessing the instances) and load balancers for serving web traffic.
- There are no prefix overlaps between the two organizations.
- Both organizations already have firewall rules that allow all inbound and outbound traffic from the 10.0.0.0/8 address space.
- Neither organization has Interconnects to their on-premises environment.

You want to integrate networking and DNS infrastructure of both organizations as quickly as possible and with minimal downtime.

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

- A. Provision Cloud Interconnect to connect both organizations together.
- B. Set up some variant of DNS forwarding and zone transfers in each organization.
- C. Connect VPCs in both organizations using Cloud VPN together with Cloud Router.
- D. Use Cloud DNS to create A records of all VMs and resources across all projects in both organizations.
- E. Create a third organization with a new host project, and attach all projects from your company and Altostrat to it using shared VPC.

Correct Answer: B, C

Section:

Explanation:

<https://cloud.google.com/dns/docs/best-practices>

QUESTION 29

Your on-premises data center has 2 routers connected to your Google Cloud environment 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 a unique 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.
- BGP sessions are established between both on-premises routers and the Cloud Router.
- Only 1 of the on-premises router's routes are being added to the routing table.

What is the most likely cause of this problem?

- A. The on-premises routers are configured with the same routes.
- 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. The ASNs being used on the on-premises routers are different.

Correct Answer: D

Section:

Explanation:

<https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp>

QUESTION 30

You have ordered Dedicated Interconnect in the GCP Console and need to give the Letter of Authorization/Connecting Facility Assignment (LOA-CFA) to your cross-connect provider to complete the physical connection. Which two actions can accomplish this? (Choose two.)

- A. Open a Cloud Support ticket under the Cloud Interconnect category.
- B. Download the LOA-CFA from the Hybrid Connectivity section of the GCP Console.
- C. Run `gcloud compute interconnects describe <interconnect>`.
- D. Check the email for the account of the NOC contact that you specified during the ordering process.
- E. Contact your cross-connect provider and inform them that Google automatically sent the LOA/CFA to them via email, and to complete the connection.

Correct Answer: D, E

Section:

Explanation:

<https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrievingloas>

QUESTION 31

Your company offers a popular gaming service. Your instances are deployed with private IP addresses, and external access is granted through a global load balancer. You believe you have identified a potential malicious actor, but aren't certain you have the correct client IP address. You want to identify this actor while minimizing disruption to your legitimate users. What should you do?

- A. Create a Cloud Armor Policy rule that denies traffic and review necessary logs.
- B. Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs.
- C. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to disabled, and review necessary logs.
- D. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to enabled, and review necessary logs.



Correct Answer: B

Section:

Explanation:

https://cloud.google.com/armor/docs/security-policy-concepts#preview_mode

QUESTION 32

Your company's web server administrator is migrating on-premises backend servers for an application to GCP. Libraries and configurations differ significantly across these backend servers. The migration to GCP will be lift-and-shift, and all requests to the servers will be served by a single network load balancer frontend. You want to use a GCP-native solution when possible. How should you deploy this service in GCP?

- A. Create a managed instance group from one of the images of the on-premises servers, and link this instance group to a target pool behind your load balancer.
- B. Create a target pool, add all backend instances to this target pool, and deploy the target pool behind your load balancer.
- C. Deploy a third-party virtual appliance as frontend to these servers that will accommodate the significant differences between these backend servers.
- D. Use GCP's ECMP capability to load-balance traffic to the backend servers by installing multiple equal-priority static routes to the backend servers.

Correct Answer: B

Section:

QUESTION 33

You decide to set up Cloud NAT. After completing the configuration, you find that one of your instances is not using the Cloud NAT for outbound NAT. What is the most likely cause of this problem?

- A. The instance has been configured with multiple interfaces.
- B. An external IP address has been configured on the instance.
- C. You have created static routes that use RFC1918 ranges.
- D. The instance is accessible by a load balancer external IP address.

Correct Answer: B

Section:

QUESTION 34

You want to set up two Cloud Routers so that one has an active Border Gateway Protocol (BGP) session, and the other one acts as a standby. Which BGP attribute should you use on your on-premises router?

- A. AS-Path
- B. Community
- C. Local Preference
- D. Multi-exit Discriminator

Correct Answer: D

Section:

QUESTION 35

You are increasing your usage of Cloud VPN between on-premises and GCP, and you want to support more traffic than a single tunnel can handle. You want to increase the available bandwidth using Cloud VPN. What should you do?

- A. Double the MTU on your on-premises VPN gateway from 1460 bytes to 2920 bytes.
- B. Create two VPN tunnels on the same Cloud VPN gateway that point to the same destination VPN gateway IP address.
- C. Add a second on-premises VPN gateway with a different public IP address. Create a second tunnel on the existing Cloud VPN gateway that forwards the same IP range, but points at the new on-premises gateway IP.
- D. Add a second Cloud VPN gateway in a different region than the existing VPN gateway. Create a new tunnel on the second Cloud VPN gateway that forwards the same IP range, but points to the existing on-premises VPN gateway IP address.

Correct Answer: C

Section:

Explanation:

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#redundancyoptions>

QUESTION 36

You are disabling DNSSEC for one of your Cloud DNS-managed zones. You removed the DS records from your zone file, waited for them to expire from the cache, and disabled DNSSEC for the zone. You receive reports that DNSSEC validating resolves are unable to resolve names in your zone. What should you do?

- A. Update the TTL for the zone.
- B. Set the zone to the TRANSFER state.
- C. Disable DNSSEC at your domain registrar.
- D. Transfer ownership of the domain to a new registrar.

Correct Answer: C

Section:

Explanation:

Before disabling DNSSEC for a managed zone you want to use, you must deactivate DNSSEC at your domain registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.

QUESTION 37

You have an application hosted on a Compute Engine virtual machine instance that cannot communicate with a resource outside of its subnet. When you review the flow and firewall logs, you do not see any denied traffic listed.

During troubleshooting you find:

- Flow logs are enabled for the VPC subnet, and all firewall rules are set to log.
- The subnetwork logs are not excluded from Stackdriver.
- The instance that is hosting the application can communicate outside the subnet.
- Other instances within the subnet can communicate outside the subnet.
- The external resource initiates communication.

What is the most likely cause of the missing log lines?

- A. The traffic is matching the expected ingress rule.
- B. The traffic is matching the expected egress rule.
- C. The traffic is not matching the expected ingress rule.
- D. The traffic is not matching the expected egress rule.

Correct Answer: C

Section:

QUESTION 38

You have configured Cloud CDN using HTTP(S) load balancing as the origin for cacheable content. Compression is configured on the web servers, but responses served by Cloud CDN are not compressed.

What is the most likely cause of the problem?

- A. You have not configured compression in Cloud CDN.
- B. You have configured the web servers and Cloud CDN with different compression types.
- C. The web servers behind the load balancer are configured with different compression types.
- D. You have to configure the web servers to compress responses even if the request has a Via header.

Correct Answer: D

Section:

Explanation:

If responses served by Cloud CDN are not compressed but should be, check that the web server software running on your instances is configured to compress responses. By default, some web server software will automatically disable compression for requests that include a Via header. The presence of a Via header indicates the request was forwarded by a proxy. HTTP proxies such as HTTP(S) load balancing add a Via header to each request as required by the HTTP specification. To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a Via header.

QUESTION 39

You have a web application that is currently hosted in the us-central1 region. Users experience high latency when traveling in Asia. You've configured a network load balancer, but users have not experienced a performance improvement.

You want to decrease the latency.

What should you do?

- A. Configure a policy-based route rule to prioritize the traffic.



- B. Configure an HTTP load balancer, and direct the traffic to it.
- C. Configure Dynamic Routing for the subnet hosting the application.
- D. Configure the TTL for the DNS zone to decrease the time between updates.

Correct Answer: B

Section:

QUESTION 40

You have an application running on Compute Engine that uses BigQuery to generate some results that are stored in Cloud Storage. You want to ensure that none of the application instances have external IP addresses. Which two methods can you use to accomplish this? (Choose two.)

- A. Enable Private Google Access on all the subnets.
- B. Enable Private Google Access on the VPC.
- C. Enable Private Services Access on the VPC.
- D. Create network peering between your VPC and BigQuery.
- E. Create a Cloud NAT, and route the application traffic via NAT gateway.

Correct Answer: A, E

Section:

Explanation:

<https://cloud.google.com/nat/docs/overview#interaction-pga> Specifications
<https://cloud.google.com/vpc/docs/configure-private-google-access#specifications>

QUESTION 41

You are designing a shared VPC architecture. Your network and security team has strict controls over which routes are exposed between departments. Your Production and Staging departments can communicate with each other, but only via specific networks. You want to follow Google recommended practices. How should you design this topology?

- A. Create 2 shared VPCs within the shared VPC Host Project, and enable VPC peering between them. Use firewall rules to filter access between the specific networks.
- B. Create 2 shared VPCs within the shared VPC Host Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.
- C. Create 2 shared VPCs within the shared VPC Service Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.
- D. Create 1 VPC within the shared VPC Host Project, and share individual subnets with the Service Projects to filter access between the specific networks.

Correct Answer: D

Section:

QUESTION 42

You are adding steps to a working automation that uses a service account to authenticate. You need to drive the automation the ability to retrieve files from a Cloud Storage bucket. Your organization requires using the least privilege possible. What should you do?

- A. Grant the compute.instanceAdmin to your user account.
- B. Grant the iam.serviceAccountUser to your user account.
- C. Grant the read-only privilege to the service account for the Cloud Storage bucket.
- D. Grant the cloud-platform privilege to the service account for the Cloud Storage bucket.

Correct Answer: C

Section:

QUESTION 43

You converted an auto mode VPC network to custom mode. Since the conversion, some of your Cloud Deployment Manager templates are no longer working. You want to resolve the problem. What should you do?

- A. Apply an additional IAM role to the Google API's service account to allow custom mode networks.
- B. Update the VPC firewall to allow the Cloud Deployment Manager to access the custom mode networks.
- C. Explicitly reference the custom mode networks in the Cloud Armor whitelist.
- D. Explicitly reference the custom mode networks in the Deployment Manager templates.

Correct Answer: D

Section:

QUESTION 44

You have recently been put in charge of managing identity and access management for your organization. You have several projects and want to use scripting and automation wherever possible. You want to grant the editor role to a project member.

Which two methods can you use to accomplish this? (Choose two.)

- A. GetIamPolicy() via REST API
- B. setIamPolicy() via REST API
- C. `gcloud pubsub add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor`
- D. `gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor`
- E. Enter an email address in the Add members field, and select the desired role from the drop-down menu in the GCP Console.

Correct Answer: D, E

Section:

QUESTION 45

You are using a 10-Gbps direct peering connection to Google together with the gsutil tool to upload files to Cloud Storage buckets from on-premises servers. The on-premises servers are 100 milliseconds away from the Google peering point. You notice that your uploads are not using the full 10-Gbps bandwidth available to you. You want to optimize the bandwidth utilization of the connection.

What should you do on your on-premises servers?

- A. Tune TCP parameters on the on-premises servers.
- B. Compress files using utilities like tar to reduce the size of data being sent.
- C. Remove the -m flag from the gsutil command to enable single-threaded transfers.
- D. Use the perfdiag parameter in your gsutil command to enable faster performance: `gsutil perfdiag gs://[BUCKET NAME]`.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid>

<https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid>

<https://cloud.google.com/blog/products/gcp/5-steps-to-better-gcp-network-performance?hl=ml>

QUESTION 46

You work for a multinational enterprise that is moving to GCP.

These are the cloud requirements:

- An on-premises data center located in the United States in Oregon and New York with Dedicated Interconnects connected to Cloud regions us-west1 (primary HQ) and us-east4 (backup)
- Multiple regional offices in Europe and APAC
- Regional data processing is required in europe-west1 and australia-southeast1
- Centralized Network Administration Team

Your security and compliance team requires a virtual inline security appliance to perform L7 inspection for URL filtering. You want to deploy the appliance in us-west1.

What should you do?

- A. • Create 2 VPCs in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Host Project. • Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. • Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- B. • Create 2 VPCs in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Service Project. • Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. • Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- C. • Create 1 VPC in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Host Project. • Attach NIC0 in us-west1 subnet of the Host Project. • Attach NIC1 in us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- D. • Create 1 VPC in a Shared VPC Service Project. • Configure a 2-NIC instance in zone us-west1-a in the Service Project. • Attach NIC0 in us-west1 subnet of the Service Project. • Attach NIC1 in us-west1 subnet of the Service Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/vpc/docs/shared-vpc>

QUESTION 47

You are designing a Google Kubernetes Engine (GKE) cluster for your organization. The current cluster size is expected to host 10 nodes, with 20 Pods per node and 150 services. Because of the migration of new services over the next 2 years, there is a planned growth for 100 nodes, 200 Pods per node, and 1500 services. You want to use VPC-native clusters with alias IP ranges, while minimizing address consumption. How should you design this topology?

- A. Create a subnet of size /25 with 2 secondary ranges of: /17 for Pods and /21 for Services. Create a VPC-native cluster and specify those ranges.
- B. Create a subnet of size /28 with 2 secondary ranges of: /24 for Pods and /24 for Services. Create a VPC-native cluster and specify those ranges. When the services are ready to be deployed, resize the subnets.
- C. Use gcloud container clusters create [CLUSTER NAME]--enable-ip-alias to create a VPC-native cluster.
- D. Use gcloud container clusters create [CLUSTER NAME] to create a VPC-native cluster.

Correct Answer: A

Section:

Explanation:

The service range setting is permanent and cannot be changed. Please see

<https://stackoverflow.com/questions/60957040/how-to-increase-the-service-address-range-of-agke-cluster>

I think the correct answer is A since: Grow is expected to up to 100 nodes (that would be /25), then up to 200 pods per node (100 times 200 = 20000 so /17 is 32768), then 1500 services in a /21 (up to 2048)

<https://docs.netgate.com/pfsense/en/latest/book/network/understanding-cidr-subnet-masknotation.html>

QUESTION 48

Your company has recently expanded their EMEA-based operations into APAC. Globally distributed users report that their SMTP and IMAP services are slow. Your company requires end-to-end encryption, but you do not have access to the SSL certificates.

Which Google Cloud load balancer should you use?

- A. SSL proxy load balancer
- B. Network load balancer

- C. HTTPS load balancer
- D. TCP proxy load balancer

Correct Answer: D

Section:

Explanation:

<https://cloud.google.com/security/encryption-in-transit/> Automatic encryption between GFEs and backends For the following load balancer types, Google automatically encrypts traffic between Google Front Ends (GFEs) and your backends that reside within Google Cloud VPC networks: HTTP(S) Load Balancing TCP Proxy Load Balancing SSL Proxy Load Balancing

QUESTION 49

Your company is working with a partner to provide a solution for a customer. Both your company and the partner organization are using GCP. There are applications in the partner's network that need access to some resources in your company's VPC. There is no CIDR overlap between the VPCs.

Which two solutions can you implement to achieve the desired results without compromising the security? (Choose two.)

- A. VPC peering
- B. Shared VPC
- C. Cloud VPN
- D. Dedicated Interconnect
- E. Cloud NAT

Correct Answer: A, C

Section:

Explanation:

Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.

QUESTION 50

You have a storage bucket that contains the following objects:

- folder-a/image-a-1.jpg
- folder-a/image-a-2.jpg
- folder-b/image-b-1.jpg
- folder-b/image-b-2.jpg

Cloud CDN is enabled on the storage bucket, and all four objects have been successfully cached. You want to remove the cached copies of all the objects with the prefix folder-a, using the minimum number of commands. What should you do?

- A. Add an appropriate lifecycle rule on the storage bucket.
- B. Issue a cache invalidation command with pattern /folder-a/*.
- C. Make sure that all the objects with prefix folder-a are not shared publicly.
- D. Disable Cloud CDN on the storage bucket. Wait 90 seconds. Re-enable Cloud CDN on the storage bucket.

Correct Answer: B

Section:

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html>

QUESTION 51

Your company is running out of network capacity to run a critical application in the on-premises data center. You want to migrate the application to GCP. You also want to ensure that the Security team does not lose their ability to monitor traffic to and from Compute Engine instances.

Which two products should you incorporate into the solution? (Choose two.)

- A. VPC flow logs
- B. Firewall logs
- C. Cloud Audit logs
- D. Stackdriver Trace
- E. Compute Engine instance system logs

Correct Answer: A, B

Section:

Explanation:

A: Using VPC Flow Logs VPC Flow Logs records a sample of network flows sent from and received by VM instances, including instances used as GKE nodes. These logs can be used for network monitoring, forensics, real-time security analysis, and expense optimization.

<https://cloud.google.com/vpc/docs/using-flow-logs> (B): Firewall Rules Logging overview Firewall Rules Logging allows you to audit, verify, and analyze the effects of your firewall rules. For example, you can determine if a firewall rule designed to deny traffic is functioning as intended. Firewall Rules Logging is also useful if you need to determine how many connections are affected by a given firewall rule. You enable Firewall Rules Logging individually for each firewall rule whose connections you need to log. Firewall Rules Logging is an option for any firewall rule, regardless of the action (allow or deny) or direction (ingress or egress) of the rule.

<https://cloud.google.com/vpc/docs/firewall-rules-logging>

QUESTION 52

You want to apply a new Cloud Armor policy to an application that is deployed in Google Kubernetes Engine (GKE). You want to find out which target to use for your Cloud Armor policy. Which GKE resource should you use?

- A. GKE Node
- B. GKE Pod
- C. GKE Cluster
- D. GKE Ingress



Correct Answer: D

Section:

Explanation:

Cloud Armour is applied at load balancers Configuring Google Cloud Armor through Ingress.

<https://cloud.google.com/kubernetes-engine/docs/how-to/ingress-features> Security policy features Google Cloud Armor security policies have the following core features: You can optionally use the QUIC protocol with load balancers that use Google Cloud Armor. You can use Google Cloud Armor with external HTTP(S) load balancers that are in either Premium Tier or Standard Tier. You can use security policies with GKE and the default Ingress controller.

QUESTION 53

You need to establish network connectivity between three Virtual Private Cloud networks, Sales, Marketing, and Finance, so that users can access resources in all three VPCs. You configure VPC peering between the Sales VPC and the Finance VPC. You also configure VPC peering between the Marketing VPC and the Finance VPC. After you complete the configuration, some users cannot connect to resources in the Sales VPC and the Marketing VPC. You want to resolve the problem.

You want to resolve the problem.

What should you do?

- A. Configure VPC peering in a full mesh.
- B. Alter the routing table to resolve the asymmetric route.
- C. Create network tags to allow connectivity between all three VPCs.
- D. Delete the legacy network and recreate it to allow transitive peering.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/vpc/docs/using-vpc-peering>

QUESTION 54

You create multiple Compute Engine virtual machine instances to be used as TFTP servers. Which type of load balancer should you use?

- A. HTTP(S) load balancer
- B. SSL proxy load balancer
- C. TCP proxy load balancer
- D. Network load balancer

Correct Answer: D

Section:

Explanation:

"TFTP is a UDP-based protocol. Servers listen on port 69 for the initial client-to-server packet to establish the TFTP session, then use a port above 1023 for all further packets during that session.

Clients use ports above 1023" https://docstore.mik.ua/oreilly/networking_2ndEd/fire/ch17_02.htm Besides, Google Cloud external TCP/UDP Network Load Balancing (after this referred to as Network Load Balancing) is a regional, non-proxied load balancer. Network Load Balancing distributes traffic among virtual machine (VM) instances in the same region in a Virtual Private Cloud (VPC) netw

QUESTION 55

You want to configure load balancing for an internet-facing, standard voice-over-IP (VOIP) application. Which type of load balancer should you use?

- A. HTTP(S) load balancer
- B. Network load balancer
- C. Internal TCP/UDP load balancer
- D. TCP/SSL proxy load balancer

Correct Answer: B

Section:

QUESTION 56

You want to configure a NAT to perform address translation between your on-premises network blocks and GCP. Which NAT solution should you use?

- A. Cloud NAT
- B. An instance with IP forwarding enabled
- C. An instance configured with iptables DNAT rules
- D. An instance configured with iptables SNAT rules

Correct Answer: A

Section:

QUESTION 57

You need to ensure your personal SSH key works on every instance in your project. You want to accomplish this as efficiently as possible. What should you do?

- A. Upload your public ssh key to the project Metadata.
- B. Upload your public ssh key to each instance Metadata.
- C. Create a custom Google Compute Engine image with your public ssh key embedded.



D. Use `gcloud compute ssh` to automatically copy your public ssh key to the instance.

Correct Answer: A

Section:

Explanation:

Overview By creating and managing SSH keys, you can let users access a Linux instance through thirdparty tools. An SSH key consists of the following files: A public SSH key file that is applied to instancelevel metadata or project-wide metadata. A private SSH key file that the user stores on their local devices. If a user presents their private SSH key, they can use a third-party tool to connect to any instance that is configured with the matching public SSH key file, even if they aren't a member of your Google Cloud project. Therefore, you can control which instances a user can access by changing the public SSH key metadata for one or more instances.

<https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys#addkey>

QUESTION 58

In order to provide subnet level isolation, you want to force instance-A in one subnet to route through a security appliance, called instance-B, in another subnet. What should you do?

- A. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with no tag.
- B. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with a tag applied to instance-A.
- C. Delete the system-generated subnet route and create a specific route to instance-B with a tag applied to instance-A.
- D. Move instance-B to another VPC and, using multi-NIC, connect instance-B's interface to instance- A's network. Configure the appropriate routes to force traffic through to instance-A.

Correct Answer: B

Section:

QUESTION 59

Your company offers a popular gaming service. Your instances are deployed with private IP addresses, and external access is granted through a global load balancer. You have recently engaged a traffic-scrubbing service and want to restrict your origin to allow connections only from the trafficscrubbing service.

What should you do?

- A. Create a Cloud Armor Security Policy that blocks all traffic except for the traffic-scrubbing service.
- B. Create a VPC Firewall rule that blocks all traffic except for the traffic-scrubbing service.
- C. Create a VPC Service Control Perimeter that blocks all traffic except for the traffic-scrubbing service.
- D. Create IPTables firewall rules that block all traffic except for the traffic-scrubbing service.

Correct Answer: A

Section:

Explanation:

Global load balancer will proxy the connection . thus no trace of session origin IP. you should use Cloud Armor to geofence your service.

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

QUESTION 60

Your software team is developing an on-premises web application that requires direct connectivity to Compute Engine Instances in GCP using the RFC 1918 address space. You want to choose a connectivity solution from your on-premises environment to GCP, given these specifications:

Your ISP is a Google Partner Interconnect provider.

Your on-premises VPN device's internet uplink and downlink speeds are 10 Gbps.

A test VPN connection between your on-premises gateway and GCP is performing at a maximum speed of 500 Mbps due to packet losses.

Most of the data transfer will be from GCP to the on-premises environment.

The application can burst up to 1.5 Gbps during peak transfers over the Interconnect.

Cost and the complexity of the solution should be minimal.

How should you provision the connectivity solution?

- A. Provision a Partner Interconnect through your ISP.
- B. Provision a Dedicated Interconnect instead of a VPN.
- C. Create multiple VPN tunnels to account for the packet losses, and increase bandwidth using ECMP.
- D. Use network compression over your VPN to increase the amount of data you can send over your VPN.

Correct Answer: A

Section:

Explanation:

Direct Interconnect will be too expensive and also an overkill for this requirement. Managing multiple tunnels that too with packet loss consideration is complex also. Whereas partner interconnect fits the bill with providing required bandwidth but not super expensive also once setup not too complex too manage.

QUESTION 61

Your company has just launched a new critical revenue-generating web application. You deployed the application for scalability using managed instance groups, autoscaling, and a network load balancer as frontend. One day, you notice severe bursty traffic that the caused autoscaling to reach the maximum number of instances, and users of your application cannot complete transactions. After an investigation, you think it as a DDOS attack. You want to quickly restore user access to your application and allow successful transactions while minimizing cost.

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

- A. Use Cloud Armor to blacklist the attacker's IP addresses.
- B. Increase the maximum autoscaling backend to accommodate the severe bursty traffic.
- C. Create a global HTTP(s) load balancer and move your application backend to this load balancer.
- D. Shut down the entire application in GCP for a few hours. The attack will stop when the application is offline.
- E. SSH into the backend compute engine instances, and view the auth logs and syslogs to further understand the nature of the attack.

Correct Answer: B, E

Section:

QUESTION 62

You are creating a new application and require access to Cloud SQL from VPC instances without public IP addresses.

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

- A. Activate the Service Networking API in your project.
- B. Activate the Cloud Datastore API in your project.
- C. Create a private connection to a service producer.
- D. Create a custom static route to allow the traffic to reach the Cloud SQL API.
- E. Enable Private Google Access.

Correct Answer: A, C

Section:

QUESTION 63

You want to use Cloud Interconnect to connect your on-premises network to a GCP VPC. You cannot meet Google at one of its point-of-presence (POP) locations, and your on-premises router cannot run a Border Gateway Protocol (BGP) configuration.

Which connectivity model should you use?

- A. Direct Peering
- B. Dedicated Interconnect

- C. Partner Interconnect with a layer 2 partner
- D. Partner Interconnect with a layer 3 partner

Correct Answer: D

Section:

Explanation:

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview> For Layer 3 connections, your service provider establishes a BGP session between your Cloud Routers and their edge routers for each VLAN attachment. You don't need to configure BGP on your on-premises router. Google and your service provider automatically set the correct configurations.

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#connectivity-type>

QUESTION 64

You have configured a Compute Engine virtual machine instance as a NAT gateway. You execute the following command:

```
gcloud compute routes create no-ip-internet-route \  
--network custom-network1 \  
--destination-range 0.0.0.0/0 \  
--next-hop instance nat-gateway \  
--next-hop instance-zone us-central1-a \  
--tags no-ip --priority 800
```

You want existing instances to use the new NAT gateway. Which command should you execute?

- A. `sudo sysctl -w net.ipv4.ip_forward=1`
- B. `gcloud compute instances add-tags [existing-instance] --tags no-ip`
- C. `gcloud builds submit --config=cloudbuild.waml --substitutions=TAG_NAME=no-ip`
- D. `gcloud compute instances create example-instance --network custom-network1 \
--subnet subnet-us-central \
--no-address \
--zone us-central1-a \
--image-family debian-9 \
--image-project debian-cloud \
--tags no-ip`

Correct Answer: B

Section:

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/compute/routes/create>

In order to apply a route to an existing instance we should use a tag to bind the route to it.

Reference: <https://cloud.google.com/vpc/docs/special-configurations>

QUESTION 65

You need to configure a static route to an on-premises resource behind a Cloud VPN gateway that is configured for policy-based routing using the gcloud command.

Which next hop should you choose?

- A. The default internet gateway
- B. The IP address of the Cloud VPN gateway
- C. The name and region of the Cloud VPN tunnel
- D. The IP address of the instance on the remote side of the VPN tunnel

Correct Answer: C

Section:



Explanation:

When you create a route based tunnel using the Cloud Console, Classic VPN performs both of the following tasks: Sets the tunnel's local and remote traffic selectors to any IP address (0.0.0.0/0) For each range in Remote network IP ranges, Google Cloud creates a custom static route whose destination (prefix) is the range's CIDR, and whose next hop is the tunnel.

<https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-static-vpns>

Reference: <https://cloud.google.com/vpn/docs/how-to/creating-static-vpns>

QUESTION 66

You need to enable Cloud CDN for all the objects inside a storage bucket. You want to ensure that all the object in the storage bucket can be served by the CDN. What should you do in the GCP Console?

- A. Create a new cloud storage bucket, and then enable Cloud CDN on it.
- B. Create a new TCP load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- C. Create a new SSL proxy load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- D. Create a new HTTP load balancer, select the storage bucket as a backend, enable Cloud CDN on the backend, and make sure each object inside the storage bucket is shared publicly.

Correct Answer: D

Section:**Explanation:**

https://cloud.google.com/load-balancing/docs/https/adding-backend-buckets-to-loadbalancers#using_cloud_cdn_with_cloud_storage_buckets

Cloud CDN needs HTTP(S) Load Balancers and Cloud Storage bucket has to be shared publicly.

<https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket>

QUESTION 67

Your company's Google Cloud-deployed, streaming application supports multiple languages. The application development team has asked you how they should support splitting audio and video traffic to different backend Google Cloud storage buckets. They want to use URL maps and minimize operational overhead. They are currently using the following directory structure:

/fr/video

/en/video

/es/video

- A. ./video
/fr/audio
/en/audio
/es/audio
- B. ./audio
Which solution should you recommend?
- C. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/* and /audio/*.
- D. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as /video/* and /audio/*.
- E. Leave the directory structure as-is, create a URL map and leverage a path rule such as $\sqrt{[az]\{2\}}$ video and $\sqrt{[a-z]\{2\}}$ audio.
- F. Leave the directory structure as-is, create a URL map and leverage a path rule such as /*/video and /*/ audio.

Correct Answer: A

Section:**Explanation:**

https://cloud.google.com/load-balancing/docs/url-map#configuring_url_maps Path matcher constraints Path matchers and path rules have the following constraints: A path rule can only include a wildcard character (*) after a forward slash character (/). For example, /videos/* and /videos/hd/* are valid for path rules, but /videos* and /videos/hd* are not. Path rules do not use regular expression or substring matching. For example, path rules for either /videos/hd or /videos/hd/* do not apply to a URL with the path /video/hd-abcd. However, a path rule for /video/* does apply to that path. <https://cloud.google.com/load-balancing/docs/url-map-concepts#pmconstraints>

QUESTION 68

You want to establish a dedicated connection to Google that can access Cloud SQL via a public IP address and that does not require a third-party service provider.

Which connection type should you choose?

- A. Carrier Peering
- B. Direct Peering
- C. Dedicated Interconnect
- D. Partner Interconnect

Correct Answer: B

Section:

Explanation:

When established, Direct Peering provides a direct path from your on-premises network to Google services, including Google Cloud products that can be exposed through one or more public IP addresses. Traffic from Google's network to your on-premises network also takes that direct path, including traffic from VPC networks in your projects. Google Cloud customers must request that direct egress pricing be enabled for each of their projects after they have established Direct Peering with Google. For more information, see Pricing.

Reference: <https://cloud.google.com/interconnect/docs/how-to/direct-peering>

QUESTION 69

You are configuring a new instance of Cloud Router in your Organization's Google Cloud environment to allow connection across a new Dedicated Interconnect to your data center Sales, Marketing, and IT each have a service project attached to the Organization's host project.

Where should you create the Cloud Router instance?

- A. VPC network in all projects
- B. VPC network in the IT Project
- C. VPC network in the Host Project
- D. VPC network in the Sales, Marketing, and IT Projects

Correct Answer: C

Section:

Explanation:

Reference: <https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnectsother-projects>

QUESTION 70

You created a new VPC for your development team. You want to allow access to the resources in this VPC via SSH only.

How should you configure your firewall rules?

- A. Create two firewall rules: one to block all traffic with priority 0, and another to allow port 22 with priority 1000.
- B. Create two firewall rules: one to block all traffic with priority 65536, and another to allow port 3389 with priority 1000.
- C. Create a single firewall rule to allow port 22 with priority 1000.
- D. Create a single firewall rule to allow port 3389 with priority 1000.

Correct Answer: C

Section:

Explanation:

Reference: <https://geekflare.com/gcp-firewall-configuration/>

QUESTION 71

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?

- One of the VPN sessions is configured incorrectly.
- A firewall is blocking the traffic across the second VPN connection.
- You do not have a load balancer to load-balance the network traffic.
- BGP sessions are not established between both on-premises routers and the Cloud Router.

Correct Answer: A

Section:

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/networkconnectivity/docs/vpn/support/troubleshooting#:~:text=If%20the%20VPN%20logs%20show,of%20ciphers%20must%20match%20exactly.&text=Make%20sure%20that%20you%20use,confi%20your%20peer%20VPN%20gateway.>

QUESTION 72

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?

- /21
- /22
- /23
- /25

Correct Answer: B

Section:

Explanation:

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

Reference: <https://cloud.google.com/kubernetes-engine/docs/how-to/alias-ips>

<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 73

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?

- Enable logging on the default Deny Any Firewall Rule.
- Enable logging on the VM Instances that receive traffic.
- Create a logging sink forwarding all firewall logs with no filters.
- Create an explicit Deny Any rule and enable logging on the new rule.



Correct Answer: D

Section:

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-ruleslogging#specifications>

QUESTION 74

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.)

- 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.

Correct Answer: B, D

Section:

QUESTION 75

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?

- 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.

Correct Answer: D

Section:



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/creatingkubernetes-engine-private-clusters-with-net-proxies>

QUESTION 76

You are creating an instance group and need to create a new health check for HTTP(s) load balancing. Which two methods can you use to accomplish this? (Choose two.)

- A. Create a new health check using the gcloud command line tool.
- B. Create a new health check using the VPC Network section in the GCP Console.
- C. Create a new health check, or select an existing one, when you complete the load balancer's backend configuration in the GCP Console.
- D. Create a new legacy health check using the gcloud command line tool.
- E. Create a new legacy health check using the Health checks section in the GCP Console.

Correct Answer: A, C

Section:**Explanation:**

https://cloud.google.com/load-balancing/docs/healthchecks#creating_and_modifying_health_checks

QUESTION 77

You are maintaining a Shared VPC in a host project. Several departments within your company have infrastructure in different service projects attached to the Shared VPC and use Identity and Access Management (IAM) permissions to manage the cloud resources in those projects. VPC Network Peering is also set up between the Shared VPC and a common services VPC that is not in a service project. Several users are experiencing failed connectivity between certain instances in different Shared VPC service projects and between certain instances and the internet. You need to validate the network configuration to identify whether a misconfiguration is the root cause of the problem. What should you do?

- A. Review the VPC audit logs in Cloud Logging for the affected instances.
- B. Use Secure Shell (SSH) to connect to the affected Compute Engine instances, and run a series of PING tests to the other affected endpoints and the 8.8.8.8 IPv4 address.
- C. Run Connectivity Tests from Network Intelligence Center to check connectivity between the affected endpoints in your network and the internet.
- D. Enable VPC Flow Logs for all VPCs, and review the logs in Cloud Logging for the affected instances.

Correct Answer: C

Section:**QUESTION 78**

Your organization has Compute Engine instances in us-east1, us-west2, and us-central1. Your organization also has an existing Cloud Interconnect physical connection in the East Coast of the United States with a single VLAN attachment and Cloud Router in us-east1. You need to provide a design with high availability and ensure that if a region goes down, you still have access to all your other Virtual Private Cloud (VPC) subnets. You need to accomplish this in the most cost-effective manner possible. What should you do?

- A. Configure your VPC routing in regional mode.
Add an additional Cloud Interconnect VLAN attachment in the us-east1 region, and configure a Cloud Router in us-east1.
- B. Configure your VPC routing in global mode.
Add an additional Cloud Interconnect VLAN attachment in the us-east1 region, and configure a Cloud Router in us-east1.
- C. Configure your VPC routing in global mode.
Add an additional Cloud Interconnect VLAN attachment in the us-west2 region, and configure a Cloud Router in us-west2.
- D. Configure your VPC routing in regional mode.
Add additional Cloud Interconnect VLAN attachments in the us-west2 and us-central1 regions, and configure Cloud Routers in us-west2 and us-central1.

Correct Answer: B

Section:

QUESTION 79

You are migrating a three-tier application architecture from on-premises to Google Cloud. As a first step in the migration, you want to create a new Virtual Private Cloud (VPC) with an external HTTP(S) load balancer. This load balancer will forward traffic back to the on-premises compute resources that run the presentation tier. You need to stop malicious traffic from entering your VPC and consuming resources at the edge, so you must configure this policy to filter IP addresses and stop cross-site scripting (XSS) attacks. What should you do?

- A. Create a Google Cloud Armor policy, and apply it to a backend service that uses an unmanaged instance group backend.
- B. Create a hierarchical firewall ruleset, and apply it to the VPC's parent organization resource node.
- C. Create a Google Cloud Armor policy, and apply it to a backend service that uses an internet network endpoint group (NEG) backend.
- D. Create a VPC firewall ruleset, and apply it to all instances in unmanaged instance groups.

Correct Answer: C

Section:

QUESTION 80

You just finished your company's migration to Google Cloud and configured an architecture with 3 Virtual Private Cloud (VPC) networks: one for Sales, one for Finance, and one for Engineering. Every VPC contains over 100 Compute Engine instances, and now developers using instances in the Sales VPC and the Finance VPC require private connectivity between each other. You need to allow communication between Sales and Finance without compromising performance or security. What should you do?

- A. Configure an HA VPN gateway between the Finance VPC and the Sales VPC.
- B. Configure the instances that require communication between each other with an external IP address.
- C. Create a VPC Network Peering connection between the Finance VPC and the Sales VPC.
- D. Configure Cloud NAT and a Cloud Router in the Sales and Finance VPCs.

Correct Answer: C

Section:

QUESTION 81

You have provisioned a Partner Interconnect connection to extend connectivity from your onpremises data center to Google Cloud. You need to configure a Cloud Router and create a VLAN attachment to connect to resources inside your VPC. You need to configure an Autonomous System number (ASN) to use with the associated Cloud Router and create the VLAN attachment. What should you do?

- A. Use a 4-byte private ASN 4200000000-4294967294.
- B. Use a 2-byte private ASN 64512-65535.
- C. Use a public Google ASN 15169.
- D. Use a public Google ASN 16550.

Correct Answer: B

Section:

QUESTION 82

You are configuring a new application that will be exposed behind an external load balancer with both IPv4 and IPv6 addresses and support TCP pass-through on port 443. You will have backends in two regions: us-west1 and us-east1. You want to serve the content with the lowest possible latency while ensuring high availability and autoscaling. Which configuration should you use?

- A. Use global SSL Proxy Load Balancing with backends in both regions.

- B. Use global TCP Proxy Load Balancing with backends in both regions.
- C. Use global external HTTP(S) Load Balancing with backends in both regions.
- D. Use Network Load Balancing in both regions, and use DNS-based load balancing to direct traffic to the closest region.

Correct Answer: D

Section:

QUESTION 83

In your project my-project, you have two subnets in a Virtual Private Cloud (VPC): subnet-a with IP range 10.128.0.0/20 and subnet-b with IP range 172.16.0.0/24. You need to deploy database servers in subnet-a. You will also deploy the application servers and web servers in subnet-b. You want to configure firewall rules that only allow database traffic from the application servers to the database servers. What should you do?

- A. Create network tag app-server and service account sa-db@my-project.iam.gserviceaccount.com. Add the tag to the application servers, and associate the service account with the database servers. Run the following command:

```
gcloud compute firewall-rules create app-db-firewall-rule \  
--action allow \  
--direction ingress \  
--rules top:3306 \  
--source-tags app-server \  
--target-service-accounts sa-db@myproject.iam.gserviceaccount.com
```
- B. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@myproject.iam.gserviceaccount.com. Associate service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command:

```
gcloud compute firewall-rules create app-db-firewall-ru \  
--allow TCP:3306 \  
--source-service-accounts sa-app@democloud-idpdemo.iam.gserviceaccount.com \  
--target-service-accounts sa-db@my-
```
- C. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@myproject.iam.gserviceaccount.com. Associate the service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command:

```
gcloud compute firewall-rules create app-db-firewall-ru \  
--allow TCP:3306 \  
--source-ranges 10.128.0.0/20 \  
--source-service-accounts sa-app@myproject.iam.gserviceaccount.com \  
--target-service-accounts sa-db@myproject.iam.gserviceaccount.com
```
- D. Create network tags app-server and db-server. Add the app-server tag to the application servers, and add the db-server tag to the database servers. Run the following command:

```
gcloud compute firewall-rules create app-db-firewall-rule \  
--action allow \  
--direction ingress \  
--rules tcp:3306 \  
--source-ranges 10.128.0.0/20 \  
--source-tags app-server \  
--target-tags db-server
```

Correct Answer: D

Section:

QUESTION 84

You are planning a large application deployment in Google Cloud that includes on-premises connectivity. The application requires direct connectivity between workloads in all regions and on-premises locations without address translation, but all RFC 1918 ranges are already in use in the on-premises locations. What should you do?

- A. Use multiple VPC networks with a transit network using VPC Network Peering.
- B. Use overlapping RFC 1918 ranges with multiple isolated VPC networks.
- C. Use overlapping RFC 1918 ranges with multiple isolated VPC networks and Cloud NAT.
- D. Use non-RFC 1918 ranges with a single global VPC.

Correct Answer: D

Section:

QUESTION 85

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?

- 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 80 and 443.

Correct Answer: C

Section:

Explanation:

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

**QUESTION 86**

You successfully provisioned a single Dedicated Interconnect. The physical connection is at a colocation facility closest to us-west2. Seventy-five percent of your workloads are in us-east4, and the remaining twenty-five percent of your workloads are in us-central1. All workloads have the same network traffic profile. You need to minimize data transfer costs when deploying VLAN attachments.

What should you do?

- A. Keep the existing Dedicated interconnect. Deploy a VLAN attachment to a Cloud Router in uswest2, and use VPC global routing to access workloads in us-east4 and us-central1.
- B. Keep the existing Dedicated Interconnect. Deploy a VLAN attachment to a Cloud Router in useast4, and deploy another VLAN attachment to a Cloud Router in us-central1.
- C. Order a new Dedicated Interconnect for a colocation facility closest to us-east4, and use VPC global routing to access workloads in us-central1.
- D. Order a new Dedicated Interconnect for a colocation facility closest to us-central1, and use VPC global routing to access workloads in us-east4.

Correct Answer: C

Section:

QUESTION 87

You are designing a hybrid cloud environment. Your Google Cloud environment is interconnected with your on-premises network using HA VPN and Cloud Router in a central transit hub VPC. The Cloud Router is configured with the default settings. Your on-premises DNS server is located at 192.168.20.88. You need to ensure that your Compute Engine resources in multiple spoke VPCs can resolve on-premises private hostnames using the domain corp.altostrat.com while also resolving Google Cloud hostnames. You want to follow Google-recommended practices. What should you do?

- A. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Configure VPC peering in the spoke VPCs to peer with the hub VPC.

- B. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke PCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.
- C. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. Create a hub-and-spoke VPN deployment in each spoke VPC to connect back to the on-premises network directly.
- D. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. Create a hub and spoke VPN deployment in each spoke VPC to connect back to the hub VPC.

Correct Answer: A

Section:

QUESTION 88

You have the following firewall ruleset applied to all instances in your Virtual Private Cloud (VPC):

Direction	Action	Address range	Port	Priority
egress	deny	192.0.2.0/24	80	100
egress	deny	198.51.100.0/24	80	200
ingress	allow	203.0.113.0/24	80	300

You need to update the firewall rule to add the following rule to the ruleset:

Direction	Action	Address range	Port	Logging
egress	deny	192.0.2.42/32	80	true



You are using a new user account. You must assign the appropriate identity and Access Management (IAM) user roles to this new user account before updating the firewall rule. The new user account must be able to apply the update and view firewall logs. What should you do?

- A. Assign the compute.securityAdmin and logging.viewer rule to the new user account. Apply the new firewall rule with a priority of 50.
- B. Assign the compute.securityAdmin and logging.bucketWriter role to the new user account. Apply the new firewall rule with a priority of 150.
- C. Assign the compute.orgSecurityPolicyAdmin and logging.viewer role to the new user account. Apply the new firewall rule with a priority of 50.
- D. Assign the compute.orgSecurityPolicyAdmin and logging.bucketWriter role to the new user account. Apply the new firewall rule with a priority of 150.

Correct Answer: A

Section:

QUESTION 89

Your organization has a single project that contains multiple Virtual Private Clouds (VPCs). You need to secure API access to your Cloud Storage buckets and BigQuery datasets by allowing API access only from resources in your corporate public networks. What should you do?

- A. Create an access context policy that allows your VPC and corporate public network IP ranges, and then attach the policy to Cloud Storage and BigQuery.
- B. Create a VPC Service Controls perimeter for your project with an access context policy that allows your corporate public network IP ranges.
- C. Create a firewall rule to block API access to Cloud Storage and BigQuery from unauthorized networks.
- D. Create a VPC Service Controls perimeter for each VPC with an access context policy that allows your corporate public network IP ranges.

Correct Answer: B

Section:

QUESTION 90

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?

- 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.

Correct Answer: B

Section:



QUESTION 91

You configured Cloud VPN with dynamic routing via Border Gateway Protocol (BGP). You added a custom route to advertise a network that is reachable over the VPN tunnel. However, the onpremises clients still cannot reach the network over the VPN tunnel. You need to examine the logs in Cloud Logging to confirm that the appropriate routers are being advertised over the VPN tunnel.

Which filter should you use in Cloud Logging to examine the logs?

- A. resource.type= "gce_router"
- B. resource.type= "gce_network_region"
- C. resource.type= "vpn_tunnel"
- D. resource.type= "vpn_gateway"

Correct Answer: C

Section:

QUESTION 92

Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with access from on-premises locations using Cloud Interconnect connections. Your company must be able to send traffic to Cloud Storage only through the Interconnect links while accessing other Google APIs and services over the public internet. What should you do?

- A. Use the default public domains for all Google APIs and services.
- B. Use Private Service Connect to access Cloud Storage, and use the default public domains for all other Google APIs and services.
- C. Use Private Google Access, with restricted.googleapis.com virtual IP addresses for Cloud Storage and private.googleapis.com for all other Google APIs and services.
- D. Use Private Google Access, with private.googleapis.com virtual IP addresses for Cloud Storage and restricted.googleapis.com virtual IP addresses for all other Google APIs and services.

Correct Answer: B

Section:

QUESTION 93

Your organization has a Google Cloud Virtual Private Cloud (VPC) with subnets in us-east1, us-west4, and europe-west4 that use the default VPC configuration. Employees in a branch office in Europe need to access the resources in the VPC using HA VPN. You configured the HA VPN associated with the Google Cloud VPC for your organization with a Cloud Router deployed in europe-west4. You need to ensure that the users in the branch office can quickly and easily access all resources in the VPC.

What should you do?

- A. Create custom advertised routes for each subnet.
- B. Configure each subnet's VPN connections to use Cloud VPN to connect to the branch office.
- C. Configure the VPC dynamic routing mode to Global.
- D. Set the advertised routes to Global for the Cloud Router.

Correct Answer: C

Section:

QUESTION 94

Your organization uses a Shared VPC architecture with a host project and three service projects. You have Compute Engine instances that reside in the service projects. You have critical workloads in your on-premises data center. You need to ensure that the Google Cloud instances can resolve on-premises hostnames via the Dedicated Interconnect you deployed to establish hybrid connectivity.

What should you do?

- A. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 35.199.192.0/19 to the on-premises environment.
- B. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the Private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 169.254.169.254 to the on-premises environment.
- C. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. In your Cloud Router, add a custom route advertisement for the IP 169.254.169.254 to the on-premises environment.
- D. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. Configure a DNS policy in the Shared VPC to allow inbound query forwarding with your on-premises DNS server as the alternative DNS server.

Correct Answer: D

Section:

QUESTION 95

Your organization is implementing a new security policy to control how firewall rules are applied to control flows between virtual machines (VMs). Using Google-recommended practices, you need to set up a firewall rule to enforce strict control of traffic between VM A and VM B. You must ensure that communications flow only from VM A to VM B within the VPC, and no other communication paths are allowed. No other firewall rules exist in the VPC. Which firewall rule should you configure to allow only this communication path?

- A. Firewall rule direction: ingress
Action: allow
Target: VM B service account
Source ranges: VM A service account
Priority: 1000
- B. Firewall rule direction: ingress
Action: allow

Target: specific VM B tag

Source ranges: VM A tag and VM A source IP address

Priority: 1000

C. Firewall rule direction: ingress

Action: allow

Target: VM A service account

Source ranges: VM B service account and VM B source IP address

Priority: 100

D. Firewall rule direction: ingress

Action: allow

Target: specific VM A tag

Source ranges: VM B tag and VM B source IP address

Priority: 100

Correct Answer: D

Section:

QUESTION 96

You have configured a service on Google Cloud that connects to an on-premises service via a Dedicated Interconnect. Users are reporting recent connectivity issues. You need to determine whether the traffic is being dropped because of firewall rules or a routing decision. What should you do?

A. Use the Network Intelligence Center Connectivity Tests to test the connectivity between the VPC and the on-premises network.

B. Use Network Intelligence Center Network Topology to check the traffic flow, and replay the traffic from the time period when the connectivity issue occurred.

C. Configure VPC Flow Logs. Review the logs by filtering on the source and destination.

D. Configure a Compute Engine instance on the same VPC as the service running on Google Cloud to run a traceroute targeted at the on-premises service.

Correct Answer: B

Section:

QUESTION 97

You are configuring a new HTTP application that will be exposed externally behind both IPv4 and IPv6 virtual IP addresses, using ports 80, 8080, and 443. You will have backends in two regions: us-west1 and us-east1. You want to serve the content with the lowest-possible latency while ensuring high availability and autoscaling, and create native content-based rules using the HTTP hostname and request path. The IP addresses of the clients that connect to the load balancer need to be visible to the backends. Which configuration should you use?

A. Use Network Load Balancing

B. Use TCP Proxy Load Balancing with PROXY protocol enabled

C. Use External HTTP(S) Load Balancing with URL Maps and custom headers

D. Use External HTTP(S) Load Balancing with URL Maps and an X-Forwarded-For header

Correct Answer: D

Section:

QUESTION 98

You need to define an address plan for a future new Google Kubernetes Engine (GKE) cluster in your Virtual Private Cloud (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?

A. /21

- B. /22
- C. /23
- D. /25

Correct Answer: A

Section:

QUESTION 99

You are responsible for designing a new connectivity solution for your organization's enterprise network to access and use Google Workspace. You have an existing Shared VPC with Compute Engine instances in us-west1. Currently, you access Google Workspace via your service provider's internet access. You want to set up a direct connection between your network and Google. What should you do?

- A. Order a Dedicated Interconnect connection in the same metropolitan area. Create a VLAN attachment, a Cloud Router in us-west1, and a Border Gateway Protocol (BGP) session between your Cloud Router and your router.
- B. Order a Direct Peering connection in the same metropolitan area. Configure a Border Gateway Protocol (BGP) session between Google and your router.
- C. Configure HA VPN in us-west1. Configure a Border Gateway Protocol (BGP) session between your Cloud Router and your on-premises data center.
- D. Order a Carrier Peering connection in the same metropolitan area. Configure a Border Gateway Protocol (BGP) session between Google and your router.

Correct Answer: B

Section:

QUESTION 100

You suspect that one of the virtual machines (VMs) in your default Virtual Private Cloud (VPC) is under a denial-of-service attack. You need to analyze the incoming traffic for the VM to understand where the traffic is coming from. What should you do?

- A. Enable Data Access audit logs of the VPC. Analyze the logs and get the source IP addresses from the subnetworks.get field.
- B. Enable VPC Flow Logs for the subnet. Analyze the logs and get the source IP addresses from the connection field.
- C. Enable VPC Flow Logs for the VPC. Analyze the logs and get the source IP addresses from the src_location field.
- D. Enable Data Access audit logs of the subnet. Analyze the logs and get the source IP addresses from the networks.get field.

Correct Answer: B

Section:

QUESTION 101

You are responsible for configuring firewall policies for your company in Google Cloud. Your security team has a strict set of requirements that must be met to configure firewall rules.

Always allow Secure Shell (SSH) from your corporate IP address.

Restrict SSH access from all other IP addresses.

There are multiple projects and VPCs in your Google Cloud organization. You need to ensure that other VPC firewall rules cannot bypass the security team's requirements. What should you do?

- A. Configure a hierarchical firewall policy to the organization node to allow TCP port 22 for your corporate IP address with priority 0.
Configure a hierarchical firewall policy to the organization node to deny TCP port 22 for all IP addresses with priority 1.
- B. Configure a VPC firewall rule to allow TCP port 22 for your corporate IP address with priority 0.
Configure a VPC firewall rule to deny TCP port 22 for all IP addresses with priority 1.
- C. Configure a VPC firewall rule to allow TCP port 22 for your corporate IP address with priority 1.
Configure a VPC firewall rule to deny TCP port 22 for all IP addresses with priority 0.
- D. Configure a hierarchical firewall policy to the organization node to allow TCP port 22 for your corporate IP address with priority 1
Configure a hierarchical firewall policy to the organization node to deny TCP port 22 for all IP addresses with priority 0.

Correct Answer: A

Section:

QUESTION 102

You are designing a new application that has backends internally exposed on port 800. The application will be exposed externally using both IPv4 and IPv6 via TCP on port 700. You want to ensure high availability for this application. What should you do?

- A. Create a network load balancer that used backend services containing one instance group with two instances.
- B. Create a network load balancer that uses a target pool backend with two instances.
- C. Create a TCP proxy that uses a zonal network endpoint group containing one instance.
- D. Create a TCP proxy that uses backend services containing an instance group with two instances.

Correct Answer: D

Section:

QUESTION 103

You work for a university that is migrating to Google Cloud.

These are the cloud requirements:

On-premises connectivity with 10 Gbps

Lowest latency access to the cloud

Centralized Networking Administration Team

New departments are asking for on-premises connectivity to their projects. You want to deploy the most cost-efficient interconnect solution for connecting the campus to Google Cloud.

What should you do?

- A. Use Shared VPC, and deploy the VLAN attachments and Dedicated Interconnect in the host project.
- B. Use Shared VPC, and deploy the VLAN attachments in the service projects. Connect the VLAN attachment to the Shared VPC's host project.
- C. Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects' Dedicated Interconnects.
- D. Use standalone projects and deploy the VLAN attachments and Dedicated Interconnects in each of the individual projects.

Correct Answer: A

Section:

QUESTION 104

You have several microservices running in a private subnet in an existing Virtual Private Cloud (VPC).

You need to create additional serverless services that use Cloud Run and Cloud Functions to access the microservices. The network traffic volume between your serverless services and private microservices is low. However, each serverless service must be able to communicate with any of your microservices. You want to implement a solution that minimizes cost. What should you do?

- A. Deploy your serverless services to the serverless VPC. Peer the serverless service VPC to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.
- B. Create a serverless VPC access connector for each serverless service. Configure the connectors to allow traffic between the serverless services and your existing microservices.
- C. Deploy your serverless services to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.
- D. Create a serverless VPC access connector. Configure the serverless service to use the connector for communication to the microservices.

Correct Answer: D

Section:

QUESTION 105

You have provisioned a Dedicated Interconnect connection of 20 Gbps with a VLAN attachment of 10 Gbps. You recently noticed a steady increase in ingress traffic on the Interconnect connection from the on-premises data center. You need to ensure that your end users can achieve the full 20 Gbps throughput as quickly as possible. Which two methods can you use to accomplish this? (Choose two.)

- A. Configure an additional VLAN attachment of 10 Gbps in another region. Configure the on-premises router to advertise routes with the same multi-exit discriminator (MED).
- B. Configure an additional VLAN attachment of 10 Gbps in the same region. Configure the on-premises router to advertise routes with the same multi-exit discriminator (MED).
- C. From the Google Cloud Console, modify the bandwidth of the VLAN attachment to 20 Gbps.
- D. From the Google Cloud Console, request a new Dedicated Interconnect connection of 20 Gbps, and configure a VLAN attachment of 10 Gbps.
- E. Configure Link Aggregation Control Protocol (LACP) on the on-premises router to use the 20-Gbps Dedicated Interconnect connection.

Correct Answer: C, E

Section:

QUESTION 106

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?

- 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 uswest1 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

Correct Answer: A

Section:

QUESTION 107

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
  servicesIpv4CidrBlock: 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
...
subnetwork: customer-1-nodes
zone: us-central1-c
```



You have a virtual machine (VM) deployed in the same VPC in the subnetwork kubernetesmanagement 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?

- 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.

Correct Answer: A

Section:

Explanation:

QUESTION 108

You have the networking configuration shown in the diagram. Two VLAN attachments associated with two Dedicated Interconnect connections terminate on the same Cloud Router (myclouddrouter). The Interconnect connections terminate on two separate on-premises routers. You advertise the same prefixes from the Border Gateway Protocol (BGP) sessions associated with each of the VLAN attachments. You notice an asymmetric traffic flow between the two Interconnect connections. Which of the following actions should you take to troubleshoot the asymmetric traffic flow?



- A. From the Google Cloud console, navigate to the Hybrid Connectivity select the Cloud Router, and view BGP sessions.
- B. From the Cloud CLI, run `gcloud compute --protect_id router get---status mycloudrouter ----region REGION` and review the results.
- C. From the Google Cloud console, navigate to Cloud Logging to view VPC Flow Logs and review the results
- D. From the Cloud CLI. run `gcloud compute routers describe mycloudrouter --region REGION` and review the results

Correct Answer: A

Section:

Explanation:

QUESTION 109

You are in the process of deploying an internal HTTP(S) load balancer for your web server virtual machine (VM) Instances What two prerequisite tasks must be completed before creating the load balancer?
Choose 2 answers

- A. Choose a region.
- B. Create firewall rules for health checks
- C. Reserve a static IP address for the load balancer
- D. Determine the subnet mask for a proxy-only subnet.
- E. Determine the subnet mask for Serverless VPC Access.

Correct Answer: B, C

Section:

Explanation:

The correct answer is B and C. You must create firewall rules for health checks and reserve a static IP address for the load balancer before creating the internal HTTP(S) load balancer.

The other options are not correct because:

Option A is not a prerequisite task. You can choose a region when you create the load balancer, but you do not need to do it beforehand.

Option D is not a prerequisite task. You can determine the subnet mask for a proxy-only subnet when you create the subnet, but you do not need to do it beforehand.

Option E is not related to the internal HTTP(S) load balancer. Serverless VPC Access is a feature that allows you to connect your serverless applications to your VPC network, but it is not required for the load balancer.

QUESTION 110

You want Cloud CDN to serve the <https://www.example.com/images/spacetime.png> static image file that is hosted in a private Cloud Storage bucket, You are using the VSE ORIG.-X_NZADERS cache mode You receive an HTTP 403 error when opening the file In your browser and you see that the HTTP response has a Cache-control: private, max-age=0 header How should you correct this Issue?

- A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role
- B. Change the cache mode to cache all content.
- C. Increase the default time-to-live (TTL) for the backend service.
- D. Enable negative caching for the backend bucket

Correct Answer: A

Section:

Explanation:

The correct answer is A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role.

This answer is based on the following facts:

Cloud CDN can serve private content from Cloud Storage buckets, but you need to grant the appropriate permissions to the Google-managed service account that represents your load balancer¹.

The Storage Legacy Object Reader role grants read access to objects in a bucket².

The Cache-control: private header indicates that the object is not publicly readable and requires authentication³.

The USE_ORIGIN_HEADERS cache mode instructs Cloud CDN to cache responses based on the Cache-Control and Expires headers from the origin server⁴. Changing the cache mode, increasing the TTL, or enabling negative caching will not affect the 403 error.

QUESTION 111

You are designing an IP address scheme for new private Google Kubernetes Engine (GKE) clusters. Due to IP address exhaustion of the RFC 1918 address space In your enterprise, you plan to use privately used public IP space for the new clusters. You want to follow Google-recommended practices. What should you do after designing your IP scheme?

- A. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters. Re-use the secondary address range for the pods across multiple private GKE clusters
- B. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters Re-use the secondary address range for the services across multiple private GKE clusters
- C. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected and
- D. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster With the following options selected --disable-default-snat, ---enable-ip-alias, and---enable-private-nodes

Correct Answer: D

Section:

Explanation:

This answer follows the Google-recommended practices for using privately used public IP (PUPI) addresses for GKE Pod address blocks¹. The benefits of this approach are:

It allows you to use any public IP addresses that are not owned by Google or your organization for your Pods, which can help mitigate address exhaustion in your enterprise.

It prevents any external traffic from reaching your Pods, as Google Cloud does not route PUPI addresses to the internet or to other VPC networks by default.

It enables you to use VPC Network Peering to connect your GKE cluster to other VPC networks that use different PUPI addresses, as long as you enable the export and import of custom routes for the peering connection.

It preserves the fully integrated network model of GKE, where Pods can communicate with nodes and other resources in the same VPC network without NAT.

The options that you need to select when creating a private GKE cluster with PUPI addresses are:

--disable-default-snat: This option disables source NAT for outbound traffic from Pods to destinations outside the cluster's VPC network.This is necessary to prevent Pods from using RFC 1918 addresses as their source IP addresses, which could cause conflicts with other networks that use the same address space².

--enable-ip-alias: This option enables alias IP ranges for Pods and Services, which allows you to use separate subnet ranges for them.This is required to use PUPI addresses for Pods¹.

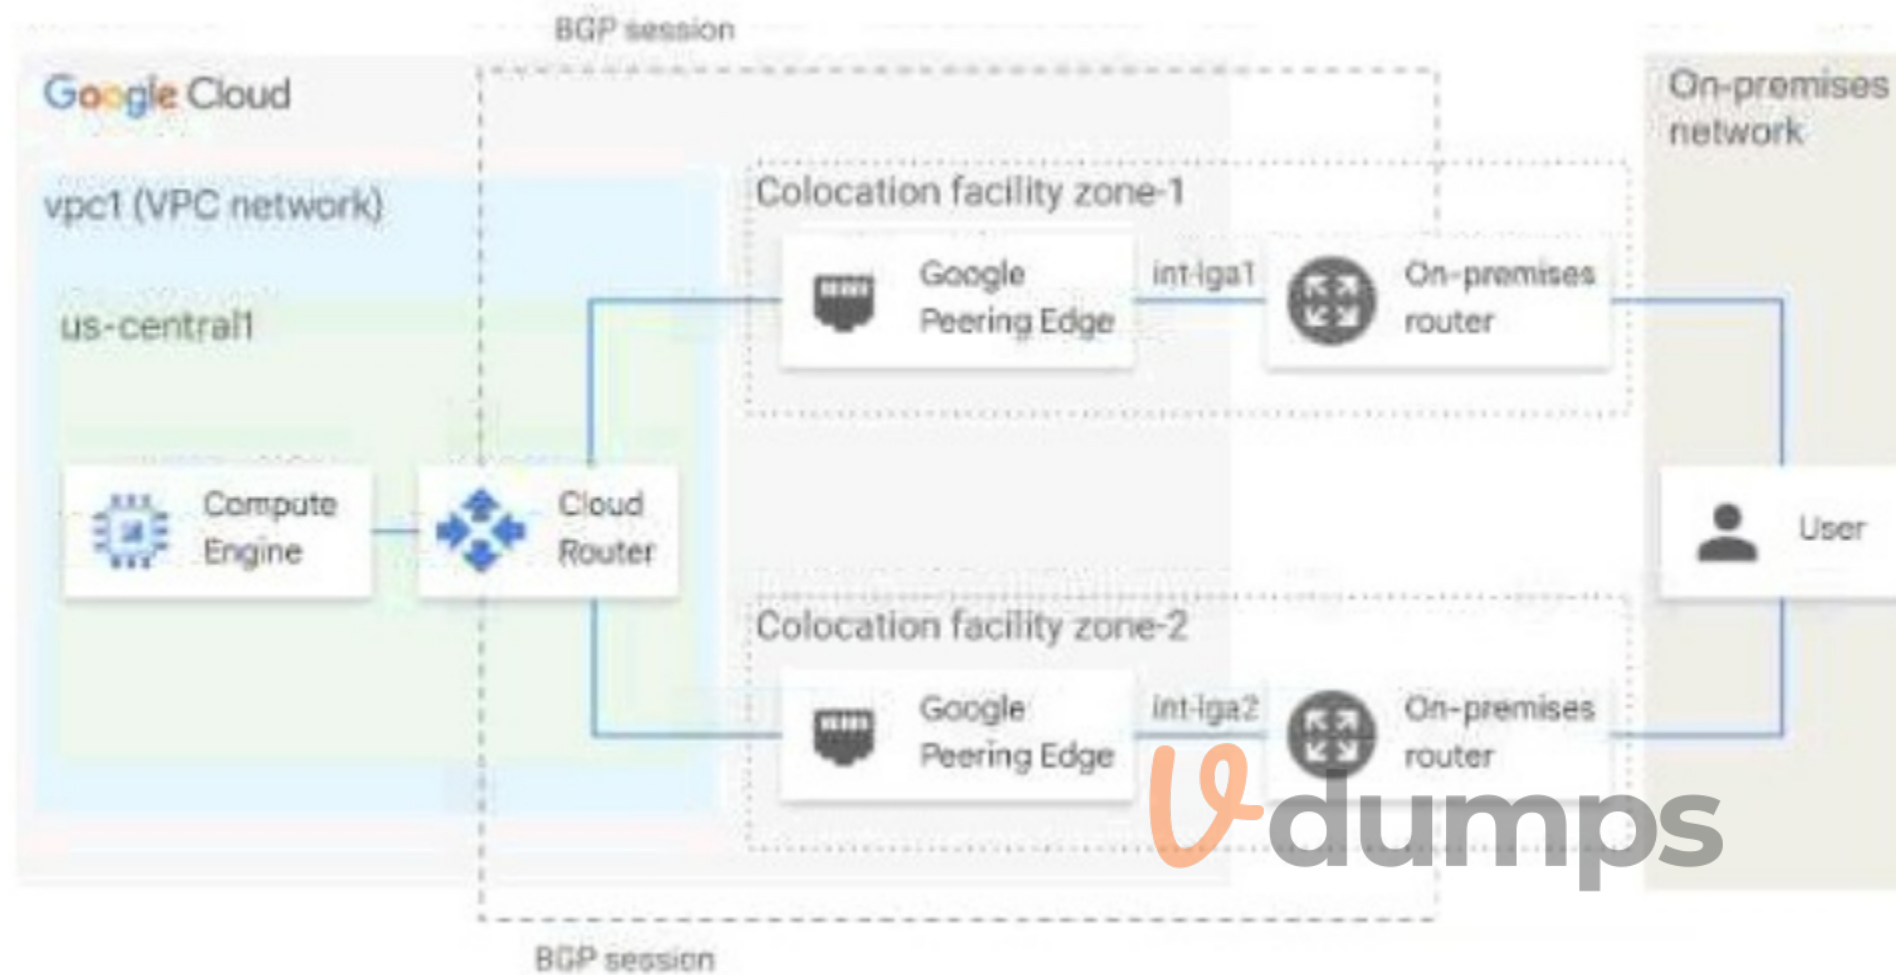
--enable-private-nodes: This option creates a private cluster, where nodes do not have external IP addresses and can only communicate with the control plane through a private endpoint.This enhances the security and privacy of your cluster³.

Option A is incorrect because it does not use PUPI addresses for Pods, but rather RFC 1918 addresses. This does not solve the problem of address exhaustion in your enterprise. Option B is incorrect because it reuses the secondary address range for Services across multiple private GKE clusters, which could cause IP conflicts and routing issues. Option C is incorrect because it does not specify the options that are needed to create a private GKE cluster with PUPI addresses.

1:Configuring privately used public IPs for GKE | Kubernetes Engine | Google Cloud2:Using Cloud NAT with GKE | Kubernetes Engine | Google Cloud3:Private clusters | Kubernetes Engine | Google Cloud

QUESTION 112

You have the networking configuration shown in the diagram. A pair of redundant Dedicated Interconnect connections (int-lga1 and int-lga2) terminate on the same Cloud Router. The Interconnect connections terminate on two separate on-premises routers. You are advertising the same prefixes from the Border Gateway Protocol (BGP) sessions associated with the Dedicated Interconnect connections. You need to configure one connection as Active for both ingress and egress traffic. If the active Interconnect connection fails, you want the passive Interconnect connection to automatically begin routing all traffic. Which two actions should you take to meet this requirement? (Choose Two)



- A. Configure the advertised route priority $> 10,200$ on the active Interconnect connection.
- B. Advertise a lower MED on the passive Interconnect connection from the on-premises router
- C. Configure the advertised route priority as 200 for the BGP session associated with the active Interconnect connection.
- D. Configure the advertised route priority as 200 for the BGP session associated with the passive Interconnect connection.
- E. Advertise a lower MED on the active Interconnect connection from the on-premises router

Correct Answer: C, E

Section:

Explanation:

This answer meets the requirement of configuring one connection as Active for both ingress and egress traffic, and enabling automatic failover to the passive connection in case of failure. The reason is:

The advertised route priority is a value that Cloud Router uses to set the route priority when advertising routes to your on-premises router. The lower the value, the higher the priority¹. By setting the advertised route priority as 200 for the active connection, you ensure that it has a higher priority than the passive connection, which has the default value of 1001. This way, your on-premises router will prefer the routes from the active connection over the passive one for ingress traffic.

The MED (Multi-Exit Discriminator) is a value that your on-premises router uses to indicate its preference for receiving traffic from Cloud Router. The lower the value, the higher the preference². By advertising a lower MED on the active connection from your on-premises router, you ensure that Cloud Router will prefer sending traffic to the active connection over the passive one for egress traffic.

If the active connection fails, Cloud Router will stop receiving routes from it and will start using the routes from the passive connection for egress traffic. Similarly, your on-premises router will stop receiving routes with priority 200 from the active connection and will start using the routes with priority 100 from the passive connection for ingress traffic. This achieves automatic failover without any manual intervention.

Option A is incorrect because setting the advertised route priority $> 10,200$ on the active connection would deprioritize it globally in your VPC network, which is not what you want¹. Option B is incorrect because advertising a

lower MED on the passive connection would make Cloud Router prefer sending traffic to it over the active one, which is not what you want. Option D is incorrect because setting the advertised route priority as 200 for both connections would make them equally preferred by your on-premises router, which is not what you want.

Update the base route priority | Cloud Router | Google Cloud

Configuring BGP sessions | Cloud Router | Google Cloud

QUESTION 113

You create a Google Kubernetes Engine private cluster and want to use kubectl to get the status of the pods. In one of your instances you notice the master is not responding, even though the cluster is up and running. What should you do to solve the problem?

- A. Assign a public IP address to the instance.
- B. Create a route to reach the Master, pointing to the default internet gateway.
- C. Create the appropriate firewall policy in the VPC to allow traffic from Master node IP address to the instance.
- D. Create the appropriate master authorized network entries to allow the instance to communicate to the master.

Correct Answer: D

Section:

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/private-clusters#cant_reach_cluster

<https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks>

QUESTION 114

Your company has a security team that manages firewalls and SSL certificates. It also has a networking team that manages the networking resources. The networking team needs to be able to read firewall rules, but should not be able to create, modify, or delete them.

How should you set up permissions for the networking team?

- A. Assign members of the networking team the compute.networkUser role.
- B. Assign members of the networking team the compute.networkAdmin role.
- C. Assign members of the networking team a custom role with only the compute.networks.* and the compute.firewalls.list permissions.
- D. Assign members of the networking team the compute.networkViewer role, and add the compute.networks.use permission.



Correct Answer: B

Section:

QUESTION 115

You have created an HTTP(S) load balanced service. You need to verify that your backend instances are responding properly.

How should you configure the health check?

- A. Set request-path to a specific URL used for health checking, and set proxy-header to PROXY_V1.
- B. Set request-path to a specific URL used for health checking, and set host to include a custom host header that identifies the health check.
- C. Set request-path to a specific URL used for health checking, and set response to a string that the backend service will always return in the response body.
- D. Set proxy-header to the default value, and set host to include a custom host header that identifies the health check.

Correct Answer: C

Section:

Explanation:

https://cloud.google.com/load-balancing/docs/health-check-concepts#contentbased_health_checks

QUESTION 116

You need to give each member of your network operations team least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments. What should you do?

- A. Assign each user the editor role.
- B. Assign each user the compute.networkAdmin role.
- C. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get.
- D. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get, compute.routers.create, compute.routers.get, compute.routers.update.

Correct Answer: D

Section:

Explanation:

<https://cloud.google.com/interconnect/docs/how-to/dedicated/creating-vlan-attachments>

QUESTION 117

You have an application that is running in a managed instance group. Your development team has released an updated instance template which contains a new feature which was not heavily tested. You want to minimize impact to users if there is a bug in the new template. How should you update your instances?

- A. Manually patch some of the instances, and then perform a rolling restart on the instance group.
- B. Using the new instance template, perform a rolling update across all instances in the instance group. Verify the new feature once the rollout completes.
- C. Deploy a new instance group and canary the updated template in that group. Verify the new feature in the new canary instance group, and then update the original instance group.
- D. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Correct Answer: D

Section:

Explanation:

https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instancegroups#starting_a_canary_update

<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instancegroups>

QUESTION 118

You have deployed a proof-of-concept application by manually placing instances in a single Compute Engine zone. You are now moving the application to production, so you need to increase your application availability and ensure it can autoscale. How should you provision your instances?

- A. Create a single managed instance group, specify the desired region, and select Multiple zones for the location.
- B. Create a managed instance group for each region, select Single zone for the location, and manually distribute instances across the zones in that region.
- C. Create an unmanaged instance group in a single zone, and then create an HTTP load balancer for the instance group.
- D. Create an unmanaged instance group for each zone, and manually distribute the instances across the desired zones.

Correct Answer: A

Section:

Explanation:

<https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances>

QUESTION 119

You have a storage bucket that contains two objects. Cloud CDN is enabled on the bucket, and both objects have been successfully cached. Now you want to make sure that one of the two objects will not be cached anymore, and will always be served to the internet directly from the origin.

What should you do?

- A. Ensure that the object you don't want to be cached anymore is not shared publicly.
- B. Create a new storage bucket, and move the object you don't want to be checked anymore inside it. Then edit the bucket setting and enable the private attribute.
- C. Add an appropriate lifecycle rule on the storage bucket containing the two objects.
- D. Add a Cache-Control entry with value private to the metadata of the object you don't want to be cached anymore. Invalidate all the previously cached copies.

Correct Answer: D

Section:

Explanation:

<https://cloud.google.com/cdn/docs/invalidating-cached-content>

QUESTION 120

Your company's logo is published as an image file across multiple websites that are hosted by your company. You have implemented Cloud CDN, however, you want to improve the performance of the cache hit ratio associated with this image file. What should you do?

- A. Configure custom cache keys for the backend service that holds the image file, and clear the Host and Protocol checkboxes.
- B. Configure Cloud Storage as a custom origin backend to host the image file, and select multi-region as the location type.
- C. Configure versioned IURLs for each domain to serve users the *mage file before the cache entry expires.
- D. Configure the default time to live (TTL) as 0 for the image file.

Correct Answer: A

Section:

Explanation:

This answer meets the requirement of improving the performance of the cache hit ratio associated with the image file. The reason is:

Custom cache keys allow you to control which parts of the request URL are used to build the cache key. The cache key is a unique identifier that Cloud CDN uses to store and retrieve cached content¹.

By default, Cloud CDN uses the complete request URL, including the protocol (http or https) and the host (the domain name), to build the cache key. This means that if the same image file is requested from different domains or protocols, Cloud CDN will cache multiple copies of it, which reduces the cache hit ratio¹.

By clearing the Host and Protocol checkboxes, you can tell Cloud CDN to ignore these parts of the request URL when building the cache key. This way, Cloud CDN will cache only one copy of the image file, regardless of which domain or protocol it is requested from, which improves the cache hit ratio¹.

Option B is incorrect because configuring Cloud Storage as a custom origin backend does not affect the cache hit ratio. It only affects how Cloud CDN retrieves the content from the origin if it is not cached. Option C is incorrect because configuring versioned URLs for each domain does not improve the cache hit ratio. It actually worsens it, because it creates more variations of the request URL that Cloud CDN has to cache separately. Option D is incorrect because configuring the default TTL as 0 for the image file means that Cloud CDN will not cache it at all, which defeats the purpose of using Cloud CDN.

Custom cache keys | Cloud CDN | Google Cloud

QUESTION 121

You are responsible for designing a new connectivity solution between your organization's on-premises data center and your Google Cloud Virtual Private Cloud (VPC) network. Currently, there is no end-to-end connectivity. You must ensure a service level agreement (SLA) of 99.99% availability. What should you do?

- A. Use one Dedicated Interconnect connection in a single metropolitan area. Configure one Cloud Router and enable global routing in the VPC.
- B. Use a Direct Peering connection between your on-premises data center and Google Cloud. Configure Classic VPN with two tunnels and one Cloud Router.
- C. Use two Dedicated Interconnect connections in a single metropolitan area. Configure one Cloud Router and enable global routing in the VPC.
- D. Use HA VPN. Configure one tunnel from each Interface of the VPN gateway to connect to the corresponding interfaces on the peer gateway on-premises. Configure one Cloud Router and enable global routing in the VPC.

Correct Answer: D

Section:

Explanation:

For Dedicated Interconnects: At least four Dedicated Interconnect connections, two connections in one metropolitan area (metro) and two connections in another metro. Connections that are in the same metro must be

placed in different edge availability domains (metro availability zones) to achieve 99.99% availability.

For HA VPN:

HA VPN to peer VPN gateways Connect an HA VPN gateway to one or two separate peer VPN devices 99.99%

HA VPN between two Google Cloud networks Connect two Google Cloud VPC networks in a single region by using an HA VPN gateway in each network 99.99%

QUESTION 122

You need to create the technical architecture for hybrid connectivity from your data center to Google Cloud This will be managed by a partner. You want to follow Google-recommended practices for production-level applications. What should you do?

- A. Ask the partner to install two security appliances in the data center. Configure one VPN connection from each of these devices to Google Cloud, and ensure that the VPN devices on-premises are in separate racks on separate power and cooling systems.
- B. Configure two Partner Interconnect connections in one metropolitan area (metro). Make sure the Interconnect connections are placed in different metro edge availability domains. Configure two VLAN attachments in a single region, and configure regional dynamic routing on the VPC
- C. Configure two Partner Interconnect connections in one metro and two connections in another metro Make sure the Interconnect connections are placed in different metro edge availability domains. Configure two VLAN attachments in one region and two VLAN attachments in another region, and configure global dynamic routing on the VPC
- D. Configure two Partner Interconnect connections in one metro and two connections in another metro. Make sure the Interconnect connections are placed in different metro edge availability domains. Configure two VLAN attachments in one region and two VLAN attachments in another region, and configure regional dynamic routing on the VPC.

Correct Answer: D

Section:

Explanation:

'Google's recommended practices for production-level applications' and then see overview of these 2 pages- <https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/production-level-overview> and <https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/non-critical-overview> .