

Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer



NEW QUESTION 1

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 zone
- B. Create a global TCP load balance
- C. Add the Cloud Run endpoints to its backend service.
- D. Deploy the Cloud Run services to multiple region
- E. Create serverless network endpoint groups (NEGs) that point to the service
- F. Create a global HTTPS load balancer, and attach the serverless NEGs as backend services of the load balancer.
- G. Deploy the Cloud Run services to multiple availability zone
- H. Create Cloud Endpoints that point to the service
- I. Create a global HTTPS load balancer, and attach the Cloud Endpoints to its backend
- J. Deploy the Cloud Run services to multiple region
- K. Configure a round-robin A record in Cloud DNS.

Answer: B

NEW QUESTION 2

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.

Answer: BC

Explanation:

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

NEW QUESTION 3

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
- D. Verify the new feature in the new canary instance group, and then update the original instance group.
- E. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template
- F. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Answer: D

Explanation:

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

NEW QUESTION 4

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.

Answer: A

NEW QUESTION 5

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.

Answer: D

Explanation:

https://cloud.google.com/load-balancing/docs/https/adding-backend-buckets-to-load-balancers#using_cloud_cdn 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>

NEW QUESTION 6

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.

Answer: D

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

NEW QUESTION 7

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.

Answer: C

NEW QUESTION 8

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 appliance
- D. Configure a route for 10.0.0.0/8 with the network load balancer as the next hop.
- E. 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.

Answer: B

NEW QUESTION 9

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.

Answer: A

Explanation:

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

NEW QUESTION 10

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

Answer: B

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.

NEW QUESTION 10

You are deploying a global external TCP load balancing solution and want to preserve the source IP address of the original layer 3 payload. Which type of load balancer should you use?

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

Answer: D

Explanation:

By default TCP/SSL proxy load balancer original client IP address and port information is not preserved, but it can be preserved using the PROXY protocol:

<https://cloud.google.com/load-balancing/docs/tcp#target-proxies>

<https://medium.com/google-cloud/preserving-client-ips-through-google-clouds-global-tcp-and-ssl-proxy-load-ba>

NEW QUESTION 14

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

Answer: C

NEW QUESTION 18

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.

Answer: A

NEW QUESTION 19

You need to restrict access to your Google Cloud load-balanced application so that only specific IP addresses can connect. What should you do?

- A. Create a secure perimeter using the Access Context Manager feature of VPC Service Controls and restrict access to the source IP range of the allowed clients and Google health check IP ranges.
- B. Create a secure perimeter using VPC Service Controls, and mark the load balancer as a service restricted to the source IP range of the allowed clients and Google health check IP ranges.
- C. Tag the backend instances "application," and create a firewall rule with target tag "application" and the source IP range of the allowed clients and Google health check IP ranges.
- D. Label the backend instances "application," and create a firewall rule with the target label "application" and the source IP range of the allowed clients and Google health check IP ranges.

Answer: C

Explanation:

<https://cloud.google.com/load-balancing/docs/https/setting-up-https#sendtraffic>

NEW QUESTION 23

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

Answer: D

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>

NEW QUESTION 27

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?

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

Answer: D

Explanation:

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

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

NEW QUESTION 31

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.

Answer: B

Explanation:

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

NEW QUESTION 32

You have applications running in the us-west1 and us-east1 regions. You want to build a highly available VPN that provides 99.99% availability to connect your applications from your project to the cloud services provided by your partner's project while minimizing the amount of infrastructure required. Your partner's services are also in the us-west1 and us-east1 regions. You want to implement the simplest solution. What should you do?

- A. Create one Cloud Router and one HA VPN gateway in each region of your VPC and your partner's VP
- B. Connect your VPN gateways to the partner's gateway
- C. Enable global dynamic routing in each VPC.
- D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VP
- E. Create one OpenVPN Access Server in each region of your partner's VP
- F. Connect your VPN gateway to your partner's servers.
- G. Create one OpenVPN Access Server in each region of your VPC and your partner's VP
- H. Connect your servers to the partner's servers.
- I. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VP
- J. Connect your VPN gateways to the partner's gateways with a pair of tunnel
- K. Enable global dynamic routing in each VPC.

Answer: A

NEW QUESTION 35

In your Google Cloud organization, you have two folders: Dev and Prod. You want a scalable and consistent way to enforce the following firewall rules for all virtual

machines (VMs) with minimal cost:

Port 8080 should always be open for VMs in the projects in the Dev folder.

Any traffic to port 8080 should be denied for all VMs in your projects in the Prod folder. What should you do?

- A. Create and associate a firewall policy with the Dev folder with a rule to open port 8080. Create and associate a firewall policy with the Prod folder with a rule to deny traffic to port 8080.
- B. Create a Shared VPC for the Dev projects and a Shared VPC for the Prod project
- C. Create a VPC firewall rule to open port 8080 in the Shared VPC for Dev
- D. Create a firewall rule to deny traffic to port 8080 in the Shared VPC for Prod
- E. Deploy VMs to those Shared VPCs.
- F. In all VPCs for the Dev projects, create a VPC firewall rule to open port 8080. In all VPCs for the Prod projects, create a VPC firewall rule to deny traffic to port 8080.
- G. Use Anthos Config Connector to enforce a security policy to open port 8080 on the Dev VMs and deny traffic to port 8080 on the Prod VMs.

Answer: A

NEW QUESTION 38

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.

Answer: B

NEW QUESTION 39

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.

Answer: AC

Explanation:

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

NEW QUESTION 42

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?

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

Answer: A

Explanation:

If the VPN logs show a no-proposal-chosen error, this error indicates that Cloud VPN and your peer VPN gateway were unable to agree on a set of ciphers. For IKEv1, the set of ciphers must match exactly. For IKEv2, there must be at least one common cipher proposed by each gateway. Make sure that you use supported ciphers to configure your peer VPN gateway.

<https://cloud.google.com/network-connectivity/docs/vpn/support/troubleshooting#:~:text=If%20the%20VPN%2>

NEW QUESTION 46

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

Answer: AB

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>

NEW QUESTION 48

You are configuring your Google Cloud environment to connect to your on-premises network. Your configuration must be able to reach Cloud Storage APIs and your Google Kubernetes Engine nodes across your private Cloud Interconnect network. You have already configured a Cloud Router with your Interconnect VLAN attachments. You now need to set up the appropriate router advertisement configuration on the Cloud Router. What should you do?

- A. Configure the route advertisement to the default setting.
- B. On the on-premises router, configure a static route for the storage API virtual IP address which points to the Cloud Router's link-local IP address.
- C. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisement
- D. Leave all other options as their default settings.
- E. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisement
- F. Advertise all visible subnets to the Cloud Router.

Answer: C

NEW QUESTION 52

Your end users are located in close proximity to us-east1 and europe-west1. Their workloads need to communicate with each other. You want to minimize cost and increase network efficiency. How should you design this topology?

- A. Create 2 VPCs, each with their own regions and individual subnet
- B. Create 2 VPN gateways to establish connectivity between these regions.
- C. Create 2 VPCs, each with their own region and individual subnet
- D. Use external IP addresses on the instances to establish connectivity between these regions.
- E. Create 1 VPC with 2 regional subnet
- F. Create a global load balancer to establish connectivity between the regions.
- G. Create 1 VPC with 2 regional subnet
- H. Deploy workloads in these subnets and have them communicate using private RFC1918 IP addresses.

Answer: D

Explanation:

<https://cloud.google.com/vpc/docs/using-vpc#create-auto-network>

We create one VPC network in auto mode that creates one subnet in each Google Cloud region automatically. So, region us-east1 and europe-west1 are in the same network and they can communicate using their internal IP address even though they are in different Regions. They take advantage of Google's global fiber network.

NEW QUESTION 57

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
- D. Verify the format of the private key and add it to the instance
- E. SSH into the instance using a third-party tool like putty or ssh.
- F. Generate a new SSH key pair
- G. Verify the format of the public key and add it to the project
- H. SSH into the instance using a third-party tool like putty or ssh.

Answer: A

NEW QUESTION 59

Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with access from your on-premises network using Cloud Interconnect. You must configure access only to Google APIs and services that are supported by VPC Service Controls through hybrid connectivity with a service level agreement (SLA) in place. What should you do?

- A. Configure the existing Cloud Routers to advertise the Google API's public virtual IP addresses.
- B. Use Private Google Access for on-premises hosts with restricted.googleapis.com virtual IP addresses.
- C. Configure the existing Cloud Routers to advertise a default route, and use Cloud NAT to translate traffic from your on-premises network.
- D. Add Direct Peering links, and use them for connectivity to Google APIs that use public virtual IP addresses.

Answer: B

NEW QUESTION 61

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 traffic-scrubbing 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.

Answer: A

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>

NEW QUESTION 66

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 plan
- B. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Router
- C. Configure authorized networks to specify the desired on-premises subnets.
- D. Create a GKE private cluster with a public endpoint for the control plan
- E. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Routers.
- F. Create a GKE private cluster with a private endpoint for the control plan
- G. Configure authorized networks to specify the desired on-premises subnets.
- H. Create a GKE public cluste
- I. Configure authorized networks to specify the desired on-premises subnets.

Answer: C

NEW QUESTION 70

After a network change window one of your company's applications stops working. The application uses an on-premises database server that no longer receives any traffic from the application. The database server IP address is 10.2.1.25. You examine the change request, and the only change is that 3 additional VPC subnets were created. The new VPC subnets created are 10.1.0.0/16, 10.2.0.0/16, and 10.3.1.0/24/ The on-premises router is advertising 10.0.0.0/8. What is the most likely cause of this problem?

- A. The less specific VPC subnet route is taking priority.
- B. The more specific VPC subnet route is taking priority.
- C. The on-premises router is not advertising a route for the database server.
- D. A cloud firewall rule that blocks traffic to the on-premises database server was created during the change.

Answer: B

NEW QUESTION 75

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

Answer: A

NEW QUESTION 77

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

Answer: D

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.

NEW QUESTION 78

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.

Answer: C

Explanation:

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

NEW QUESTION 79

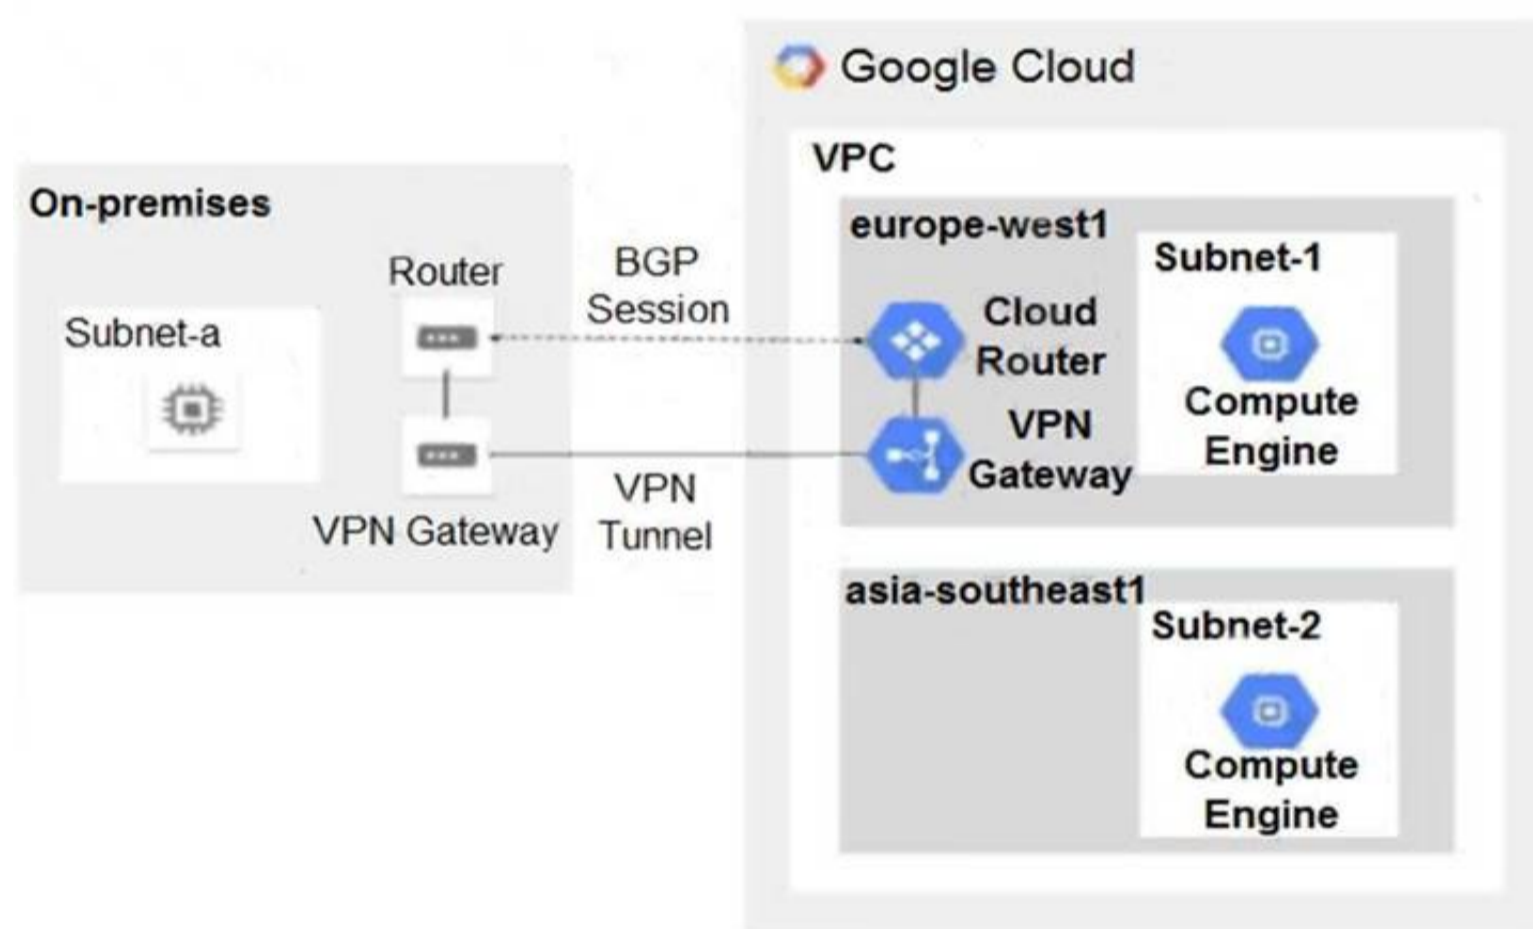
You have an HA VPN connection with two tunnels running in active/passive mode between your Virtual Private Cloud (VPC) and on-premises network. Traffic over the connection has recently increased from 1 gigabit per second (Gbps) to 4 Gbps, and you notice that packets are being dropped. You need to configure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

Answer: D

NEW QUESTION 84

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.

Answer: C

NEW QUESTION 86

One instance in your VPC is configured to run with a private IP address only. You want to ensure that even if this instance is deleted, its current private IP address will not be automatically assigned to a different instance. In the GCP Console, what should you do?

- A. Assign a public IP address to the instance.
- B. Assign a new reserved internal IP address to the instance.
- C. Change the instance's current internal IP address to static.
- D. Add custom metadata to the instance with key internal-address and value reserved.

Answer: C

Explanation:

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

NEW QUESTION 89

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.

Answer: D

NEW QUESTION 90

You are in the early stages of planning a migration to GCP. You want to test the functionality of your hybrid cloud design before you start to implement it in production. The design includes services running on a Compute Engine Virtual Machine instance that need to communicate to on-premises servers using private IP addresses. The on-premises servers have connectivity to the internet, but you have not yet established any Cloud Interconnect connections. You want to choose the lowest cost method of enabling connectivity between your instance and on-premises servers and complete the test in 24 hours.

Which connectivity method should you choose?

- A. Cloud VPN
- B. 50-Mbps Partner VLAN attachment
- C. Dedicated Interconnect with a single VLAN attachment
- D. Dedicated Interconnect, but don't provision any VLAN attachments

Answer: A

NEW QUESTION 92

You are designing a Partner Interconnect hybrid cloud connectivity solution with geo-redundancy across two metropolitan areas. You want to follow Google-recommended practices to set up the following region/metro pairs:

(region 1/metro 1)

(region 2/metro 2) What should you do?

- A. Create a Cloud Router in region 1 with two VLAN attachments connected to metro1-zone1-x. Create a Cloud Router in region 2 with two VLAN attachments connected to metro1-zone2-x.
- B. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone1-x. Create a Cloud Router in region 2 with two VLAN attachments connected to metro2-zone2-x.
- C. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone2-x. Create a Cloud Router in region 2 with one VLAN attachment connected to metro2-zone2-x.
- D. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone1-x and one VLAN attachment connected to metro1-zone2-x. Create a Cloud Router in region 2 with one VLAN attachment connected to metro2-zone1-x and one VLAN attachment to metro2-zone2-x.

Answer: B

NEW QUESTION 97

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 Cloud Logging.
- 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 CloudConsole.

Answer: A

NEW QUESTION 102

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.

Answer: C

NEW QUESTION 106

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.

Answer: B

NEW QUESTION 107

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

Answer: D

NEW QUESTION 112

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 VP
- B. Peer the serverless service VPC to the existing VP
- C. Configure firewall rules to allow traffic between the serverless services and your existing microservices.
- D. Create a serverless VPC access connector for each serverless servic
- E. Configure the connectors to allow traffic between the serverless services and your existing microservices.
- F. Deploy your serverless services to the existing VP
- G. Configure firewall rules to allow traffic between the serverless services and your existing microservices.
- H. Create a serverless VPC access connecto
- I. Configure the serverless service to use the connector for communication to the microservices.

Answer: D

NEW QUESTION 113

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

Answer: A

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-multiple-networks-access-sa>

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

NEW QUESTION 116

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

Answer: A

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-a-gke-cluster> I think the correc tanswer 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-mask-notation.html>

NEW QUESTION 119

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

Answer: D

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

NEW QUESTION 120

Your company has separate Virtual Private Cloud (VPC) networks in a single region for two departments: Sales and Finance. The Sales department's VPC network already has connectivity to on-premises locations using HA VPN, and you have confirmed that the subnet ranges do not overlap. You plan to peer both VPC networks to use the same HA tunnels for on-premises connectivity, while providing internet connectivity for the Google Cloud workloads through Cloud NAT. Internet access from the on-premises locations should not flow through Google Cloud. You need to propagate all routes between the Finance department and on-premises locations. What should you do?

- A. Peer the two VPCs, and use the default configuration for the Cloud Routers.
- B. Peer the two VPCs, and use Cloud Router's custom route advertisements to announce the peered VPC network ranges to the on-premises locations.
- C. Peer the two VPC
- D. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network
- E. Use Cloud Router's custom route advertisements to announce a default route to the on-premises locations.
- F. Peer the two VPC
- G. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network
- H. Use Cloud Router's custom route advertisements to announce the peered VPC network ranges to the on-premises locations.

Answer: A

NEW QUESTION 121

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

Answer: C

NEW QUESTION 124

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.

Answer: B

NEW QUESTION 128

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.

Answer: C

NEW QUESTION 132

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
- D. 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.
- E. Add a second Cloud VPN gateway in a different region than the existing VPN gateway
- F. 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.

Answer: C

Explanation:

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

NEW QUESTION 134

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
/./video
/fr/audio
/en/audio
/es/audio
/./audio
```

Which solution should you recommend?

- A. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/* and /audio/*.
- B. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as /video/* and /audio/*.
- C. Leave the directory structure as-is, create a URL map and leverage a path rule such as V[a-z]{2}\video and V[a-z]{2}\audio.
- D. Leave the directory structure as-is, create a URL map and leverage a path rule such as /*/video and /*/ audio.

Answer: A

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#pm-constraints>

NEW QUESTION 139

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 (
- D. g., code-dev), and make the second project (
- E. g., data-dev) a service project.
- F. 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.
- G. Create a route in the code-dev project to the destination prefixes in project data-dev and use nexthop as the default gateway, and vice versa.

Answer: BD

NEW QUESTION 144

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.

Answer: A

Explanation:

Overview By creating and managing SSH keys, you can let users access a Linux instance through third-party tools. An SSH key consists of the following files: A public SSH key file that is applied to instance-level 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>

NEW QUESTION 146

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.

Answer: AD

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.

NEW QUESTION 148

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-Network-Engineer Practice Exam Features:

- * Professional-Cloud-Network-Engineer Questions and Answers Updated Frequently
- * Professional-Cloud-Network-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Network-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * Professional-Cloud-Network-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Network-Engineer Practice Test Here](#)