Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer

NEW QUESTION 1
You have just deployed your infrastructure on Google Cloud. You now need to configure the DNS to meet the following requirements:
Your on-premises resources should resolve your Google Cloud zones. Your Google Cloud resources should resolve your on-premises zones.
You need the ability to resolve ".internal" zones provisioned by Google Cloud. What should you do?

A. Configure an outbound server policy, and set your alternative name server to be your on-premises DNS resolve
B. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.
C. Configure both an inbound server policy and outbound DNS forwarding zones with the target as the on-premises DNS resolve
D. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
E. Configure an outbound DNS server policy, and set your alternative name server to be your on-premises DNS resolve
F. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
G. Configure Cloud DNS to DNS peer with your on-premises DNS resolve
H. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

Answer: A

NEW QUESTION 2
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.

Answer: DE

Explanation:
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrieving-loas

NEW QUESTION 3
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 1Configure a hierarchical firewall policy to the organization node to deny TCP port 22 for all IP addresses with priority 0.

Answer: A

NEW QUESTION 4
You have enabled HTTP(S) load balancing for your application, and your application developers have reported that HTTP(S) requests are not being distributed correctly to your Compute Engine Virtual Machine instances. You want to find data about how the request are being distributed.
Which two methods can accomplish this? (Choose two.)

A. On the Load Balancer details page of the GCP Console, click on the Monitoring tab, select your backend service, and look at the graphs.
B. In Stackdriver Error Reporting, look for any unacknowledged errors for the Cloud Load Balancers service.
C. In Stackdriver Monitoring, select Resources > Metrics Explorer and search for https/request_bytes_count metric.
D. In Stackdriver Monitoring, select Resources > Google Cloud Load Balancers and review the Key Metrics graphs in the dashboard.
E. In Stackdriver Monitoring, create a new dashboard and track the https/backend_request_count metric for the load balancer.

Answer: AE

NEW QUESTION 5
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 6
You need to centralize the Identity and Access Management permissions and email distribution for the WebServices Team as efficiently as possible.
What should you do?

A. Create a Google Group for the WebServices Team.
B. Create a G Suite Domain for the WebServices Team.
C. Create a new Cloud Identity Domain for the WebServices Team.
D. Create a new Custom Role for all members of the WebServices Team.

Answer: A

NEW QUESTION 7
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 addresse
B. Configure a (25 secondary IP range for the Pod
C. Configure a /22 secondary IP range for the Services.
D. Configure a /28 primary IP address range for the node IP addresse
E. Configure a /25 secondary IP range for the Pod
F. Configure a /21 secondary IP range for the Services.
G. Configure a /28 primary IP address range for the node IP addresse
H. Configure a /28 secondary IP range for the Pod
I. Configure a /21 secondary IP range for the Services.
J. Configure a /28 primary IP address range for the node IP addresse
K. Configure a /24 secondary IP range for the Pad
L. Configure a /22 secondary IP range for the Services.

Answer: A

NEW QUESTION 8
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 9
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 perimete
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 10
You recently noticed a recurring daily spike in network usage in your Google Cloud project. You need to identify the virtual machine (VM) instances and type of traffic causing the spike in traffic utilization while minimizing the cost and management overhead required. What should you do?

A. Enable VPC Flow Logs and send the output to BigQuery for analysis.
B. Enable Firewall Rules Logging for all allowed traffic and send the output to BigQuery for analysis.
C. Configure Packet Mirroring to send all traffic to a V
D. Use Wireshark on the VM to identity traffic utilization for each VM in the VPC.
E. Deploy a third-party network appliance and configure it as the default gatewa
F. Use the third-party network appliance to identify users with high network traffic.

Answer: C

NEW QUESTION 10
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.

Answer: C

NEW QUESTION 11
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.

Answer: AB

NEW QUESTION 15
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 onpremises data center.

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 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 24
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 bucke
E. Wait 90 second
F. Re-enable Cloud CDN on the storage bucket.

Answer: B

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

NEW QUESTION 25
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 29
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.

Answer: B

NEW QUESTION 31
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 mod
B. Peer both the VPCs via network peering.
C. Create the Distribution VPC in custom mod
D. Use the CIDR range 10.0.0.0/9. Create the necessary subnets, and then peer them via network peering.
E. Create the Distribution VPC in custom mod
F. Use the CIDR range 10.128.0.0/9. Create the necessary subnets, and then peer them via network peering.
G. Rename the default VPC as "Distribution" and peer it via network peering.

Answer: B

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

NEW QUESTION 35
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.

Answer: D

NEW QUESTION 38
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 De
D. Create a firewall rule to deny traffic to port 8080 in the Shared VPC for Pro
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 43
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 thisinstance 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 44
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 48
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 hour
E. The attack will stop when the application is offline.
F. SSH into the backend compute engine instances, and view the auth logs and syslogs to further understand the nature of the attack.

Answer: BE

NEW QUESTION 53
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.

Answer: B

NEW QUESTION 58
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

Answer: D

NEW QUESTION 59
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

Answer: AC

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.

NEW QUESTION 60
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.

Answer: D

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.

NEW QUESTION 64
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.

Answer: D

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

NEW QUESTION 69
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 * googieapis.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 * googieapis.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.

Answer: C

NEW QUESTION 71
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 76
You have the following private Google Kubernetes Engine (GKE) cluster deployment:

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

...

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

...

masterAuthorizedNetworksConfig:

...

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

...

servicesIpv4Cidr: 192.162.37.0/24

...

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

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

A. Add the network 192.168.40.0/24 to the masterAuthorizedNetworksConfi
B. Configure kubectl tocommunicate with the endpoint 192.168.38.2.
C. Add the network 192.168.38.0/28 to the masterAuthorizedNetworksConfi
D. Configure kubectl to communicate with the endpoint 192.168.38.2
E. Add the network 192.168.36.0/24 to the masterAuthorizedNetworksConfi
F. Configure kubectl to communicate with the endpoint 192.168.38.2
G. Add an external IP address to the VM, and add this IP address in the masterAuthorizedNetworksConfig.Configure kubectl to communicate with the endpoint 35.224.37.17.

Answer: A

NEW QUESTION 78
You need to create the network infrastructure to deploy a highly available web application in the us-east1 and us-west1 regions. The application runs on Compute Engine instances, and it does not require the use of a database. You want to follow Google-recommended practices. What should you do?

A. Create one VPC with one subnet in each region.Create a regional network load balancer in each region with a static IP addres
B. Enable Cloud CDN on the load balancers.Create an A record in Cloud DNS with both IP addresses for the load balancers.
C. Create one VPC with one subnet in each region.Create a global load balancer with a static IP address.Enable Cloud CDN and Google Cloud Armor on the load balancer.Create an A record using the IP address of the load balancer in Cloud DNS.
D. Create one VPC in each region, and peer both VPCs.Create a global load balancer.Enable Cloud CDN on the load balancer.Create a CNAME for the load balancer in Cloud DNS.
E. Create one VPC with one subnet in each region.Create an HTTP(S) load balancer with a static IP address.Choose the standard tier for the networ
F. Enable Cloud CDN on the load balancer.Create a CNAME record using the load balancer's IP address in Cloud DNS.

Answer: C

NEW QUESTION 80
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.

Answer: B

Explanation:
https://cloud.google.com/vpc/docs/shared-vpc

NEW QUESTION 82
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 83
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 are configuring load balancing for a standard three-tier (web, application, and database) application. You have configured an external HTTP(S) load balancer for the web servers. You need to configure load balancing for the application tier of servers. What should you do?

A. Configure a forwarding rule on the existing load balancer for the application tier.
B. Configure equal cost multi-path routing on the application servers.
C. Configure a new internal HTTP(S) load balancer for the application tier.
D. Configure a URL map on the existing load balancer to route traffic to the application tier.

Answer: A

NEW QUESTION 86
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 88
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?

A. /21
B. /22
C. /23
D. /25

Answer: B

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

NEW QUESTION 92
You are configuring an HA VPN connection between your Virtual Private Cloud (VPC) and on-premises 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.

Answer: B

NEW QUESTION 96
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 identity 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 99
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 104
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.

Answer: C

NEW QUESTION 105
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.

Answer: CE

Explanation:
https://cloud.google.com/sql/docs/mysql/configure-private-services-access#console_1
C: If you are using private IP for any of your Cloud SQL instances, you only need to configure private services access one time for every Google Cloud project that has or needs to connect to a Cloud SQL instance. If your Google Cloud project has a Cloud SQL instance, you can either configure it yourself or let Cloud SQL do it for you to use private IP. Cloud SQL configures private services access for you when all the conditions below are true:
https://cloud.google.com/sql/docs/postgres/configure-private-services-access#before_you_begin

E: You can enable Private Google access on a subnet level and any VMs on that subnet can access Google APIs by using their internal IP address.
https://cloud.google.com/vpc/docs/configure-private-google-access

NEW QUESTION 108
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.

Answer: DE

NEW QUESTION 112
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 114
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 accoun
B. Apply the new firewall rule with a priority of 50.
C. Assign the compute.securityAdmin and logging.bucketWriter role to the new user accoun
D. Apply the new firewall rule with a priority of 150.
E. Assign the compute.orgSecurityPolicyAdmin and logging.viewer role to the new user accoun
F. Apply the new firewall rule with a priority of 50.
G. Assign the compute.orgSecurityPolicyAdmin and logging.bucketWriter role to the new user account.Apply the new firewall rule with a priority of 150.

Answer: A

NEW QUESTION 118
You have deployed an HTTP(s) load balancer, but health checks to port 80 on the Compute Engine virtual machine instance are failing, and no traffic is sent to your instances. You want to resolve the problem. Which commands should you run?

A. gcloud compute instances add-access-config instance-1
B. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destination-ranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS
C. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS
D. gcloud compute health-checks update http health-check --unhealthy-threshold 10

Answer: A

NEW QUESTION 122
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 125
You are developing an HTTP API hosted on a Compute Engine virtual machine instance that must be invoked only by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?

A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rul
B. Clients should use this IP address to connect to the service.
C. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/.
D. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwardingrul
E. Then, define an A record in Cloud DN
F. Clients should use the name of the A record to connect to the service.
G. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Answer: C

NEW QUESTION 129
You need to create a new VPC network that allows instances to have IP addresses in both the 10.1.1.0/24 network and the 172.16.45.0/24 network. What should you do?

A. Configure global load balancing to point 172.16.45.0/24 to the correct instance.
B. Create unique DNS records for each service that sends traffic to the desired IP address.
C. Configure an alias-IP range of 172.16.45.0/24 on the virtual instances within the VPC subnet of 10.1.1.0/24.
D. Use VPC peering to allow traffic to route between the 10.1.0.0/24 network and the 172.16.45.0/24network.

Answer: C

NEW QUESTION 132
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 \/[a-z]{2}\/video and \/[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 137
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

Answer: B

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.

NEW QUESTION 140
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.

Answer: C

NEW QUESTION 142
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 144
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 regio
B. Configure the on-premises router to advertise routes with the same multi-exit discriminator (MED).
C. Configure an additional VLAN attachment of 10 Gbps in the same regio
D. Configure the on-premises router to advertise routes with the same multi-exit discriminator (MED).
E. From the Google Cloud Console, modify the bandwidth of the VLAN attachment to 20 Gbps.
F. From the Google Cloud Console, request a new Dedicated Interconnect connection of 20 Gbps, and configure a VLAN attachment of 10 Gbps.
G. Configure Link Aggregation Control Protocol (LACP) on the on-premises router to use the 20-Gbps Dedicated Interconnect connection.

Answer: CE

NEW QUESTION 147
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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