



Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer

NEW QUESTION 1

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

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

Answer: A

Explanation:

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

NEW QUESTION 2

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 3

You have a Cloud Storage bucket in Google Cloud project XYZ. The bucket contains sensitive data. You need to design a solution to ensure that only instances belonging to VPCs under project XYZ can access the data stored in this Cloud Storage bucket. What should you do?

- A. Configure Private Google Access to privately access the Cloud Storage service using private IP addresses.
- B. Configure a VPC Service Controls perimeter around project XYZ, and include storage.googleapis.com as a restricted service in the service perimeter.
- C. Configure Cloud Storage with projectPrivate Access Control List (ACL) that gives permission to the project team based on their roles.
- D. Configure Private Service Connect to privately access Cloud Storage from all VPCs under project XYZ.

Answer: C

NEW QUESTION 4

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

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

Answer: D

NEW QUESTION 5

You need to define an address plan for a future new Google Kubernetes Engine (GKE) cluster in your Virtual Private Cloud (VPC). This will be a VPC-native cluster, and the default Pod IP range allocation will be used. You must pre-provision all the needed VPC subnets and their respective IP address ranges before cluster creation. The cluster will initially have a single node, but it will be scaled to a maximum of three nodes if necessary. You want to allocate the minimum number of Pod IP addresses. Which subnet mask should you use for the Pod IP address range?

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

Answer: A

NEW QUESTION 6

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 7

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

What should you do to solve the problem?

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

Answer: D

Explanation:

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

NEW QUESTION 8

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 9

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 10

You need to enable Private Google Access for use by some subnets within your Virtual Private Cloud (VPC). Your security team set up the VPC to send all internet-bound traffic back to the on-premises data center for inspection before egressing to the internet, and is also implementing VPC Service Controls in the environment for API-level security control. You have already enabled the subnets for Private Google Access. What configuration changes should you make to enable Private Google Access while adhering to your security team's requirements?

- A. Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range. Create a custom route that points Google's restricted API address range to the default internet gateway as the next hop.
- B. Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range. Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.
- C. Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record pointing to Google's private API address range. Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.
- D. Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record pointing to Google's private API address range. Create a custom route that points Google's private API address range to the default internet gateway as the next hop.

Answer: C

NEW QUESTION 10

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

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

Answer: B

NEW QUESTION 13

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 17

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

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

Answer: A

Explanation:

<https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid> <https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid>
<https://cloud.google.com/blog/products/gcp/5-steps-to-better-gcp-network-performance?hl=ml>

NEW QUESTION 22

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

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

Answer: B

NEW QUESTION 23

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 28

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

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

Answer: AE

Explanation:

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

NEW QUESTION 29

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 identify traffic utilization for each VM in the VPC.
- E. Deploy a third-party network appliance and configure it as the default gateway
- F. Use the third-party network appliance to identify users with high network traffic.

Answer: C

NEW QUESTION 31

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 35

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 37

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 42

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 45

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 46

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 51

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

- A. Use regional routin
- B. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 CloudRouter to a base priority of 1
- C. Use global routin
- D. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1
- E. Use regional routin
- F. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1
- G. Use global routin
- H. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1

Answer: A

NEW QUESTION 54

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 59

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 62

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

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

Answer: B

NEW QUESTION 65

You are migrating to Cloud DNS and want to import your BIND zone file. Which command should you use?

- A. `gcloud dns record-sets import ZONE_FILE --zone MANAGED_ZONE`
- B. `gcloud dns record-sets import ZONE_FILE --replace-origin-ns --zone MANAGED_ZONE`
- C. `gcloud dns record-sets import ZONE_FILE --zone-file-format --zone MANAGED_ZONE`
- D. `gcloud dns record-sets import ZONE_FILE --delete-all-existing --zone MANAGED_ZONE`

Answer: C

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/dns/record-sets/import>

NEW QUESTION 69

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 73

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

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

Answer: B

Explanation:

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

NEW QUESTION 75

Your organization is deploying a single project for 3 separate departments. Two of these departments require network connectivity between each other, but the

third department should remain in isolation. Your design should create separate network administrative domains between these departments. You want to minimize operational overhead.
How should you design the topology?

- A. Create a Shared VPC Host Project and the respective Service Projects for each of the 3 separate departments.
- B. Create 3 separate VPCs, and use Cloud VPN to establish connectivity between the two appropriate VPCs.
- C. Create 3 separate VPCs, and use VPC peering to establish connectivity between the two appropriate VPCs.
- D. Create a single project, and deploy specific firewall rule
- E. Use network tags to isolate access between the departments.

Answer: C

Explanation:

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

NEW QUESTION 76

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 79

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 81

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 83

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 88

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 93

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 address.
- 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 network.
- 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 98

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

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

Answer: B

NEW QUESTION 100

You deployed a hub-and-spoke architecture in your Google Cloud environment that uses VPC Network Peering to connect the spokes to the hub. For security reasons, you deployed a private Google Kubernetes Engine (GKE) cluster in one of the spoke projects with a private endpoint for the control plane. You configured authorized networks to be the subnet range where the GKE nodes are deployed. When you attempt to reach the GKE control plane from a different spoke project, you cannot access it. You need to allow access to the GKE control plane from the other spoke projects. What should you do?

- A. Add a firewall rule that allows port 443 from the other spoke projects.
- B. Enable Private Google Access on the subnet where the GKE nodes are deployed.
- C. Configure the authorized networks to be the subnet ranges of the other spoke projects.
- D. Deploy a proxy in the spoke project where the GKE nodes are deployed and connect to the control plane through the proxy.

Answer: C

NEW QUESTION 103

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 108

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.

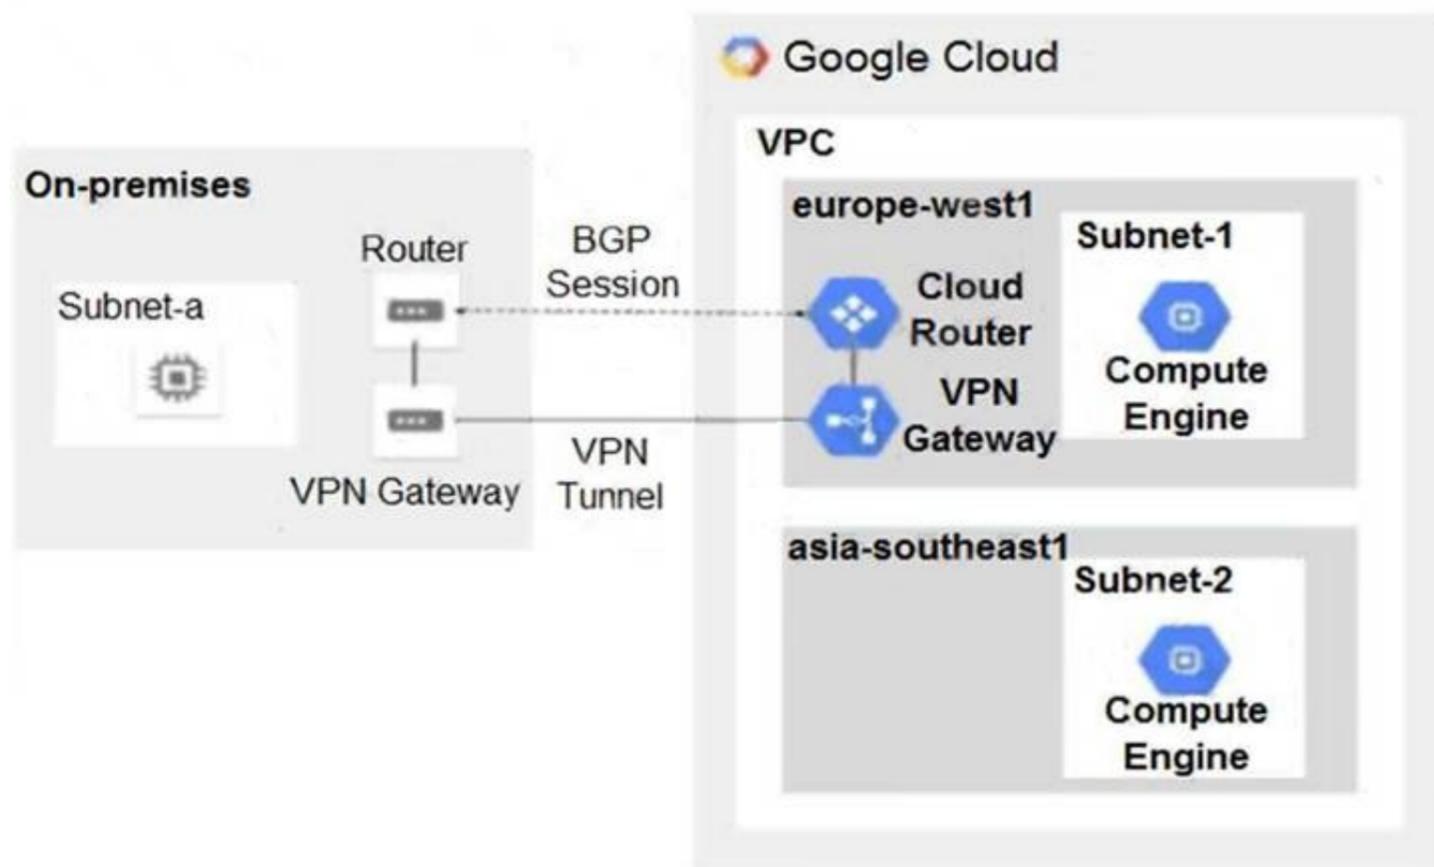
Answer: C

Explanation:

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

NEW QUESTION 111

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 116

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 118

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 122

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

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

Answer: C

NEW QUESTION 123

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

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

Answer: B

NEW QUESTION 127

Your organization has a new security policy that requires you to monitor all egress traffic payloads from your virtual machines in region us-west2. You deployed an intrusion detection system (IDS) virtual appliance in the same region to meet the new policy. You now need to integrate the IDS into the environment to monitor all egress traffic payloads from us-west2. What should you do?

- A. Enable firewall logging, and forward all filtered egress firewall logs to the IDS.
- B. Enable VPC Flow Log
- C. Create a sink in Cloud Logging to send filtered egress VPC Flow Logs to the IDS.
- D. Create an internal TCP/UDP load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.
- E. Create an internal HTTP(S) load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

Answer: B

NEW QUESTION 131

You are using the gcloud command line tool to create a new custom role in a project by copying a predefined role. You receive this error message: INVALID_ARGUMENT: Permission resourcemanager.projects.list is not valid What should you do?

- A. Add the resourcemanager.projects.get permission, and try again.
- B. Try again with a different role with a new name but the same permissions.
- C. Remove the resourcemanager.projects.list permission, and try again.
- D. Add the resourcemanager.projects.setIamPolicy permission, and try again.

Answer: C

NEW QUESTION 134

You want to implement an IPSec tunnel between your on-premises network and a VPC via Cloud VPN. You need to restrict reachability over the tunnel to specific local subnets, and you do not have a device capable of speaking Border Gateway Protocol (BGP). Which routing option should you choose?

- A. Dynamic routing using Cloud Router
- B. Route-based routing using default traffic selectors
- C. Policy-based routing using a custom local traffic selector
- D. Policy-based routing using the default local traffic selector

Answer: C

NEW QUESTION 136

Your company has 10 separate Virtual Private Cloud (VPC) networks, with one VPC per project in a single region in Google Cloud. Your security team requires each VPC network to have private connectivity to the main on-premises location via a Partner Interconnect connection in the same region. To optimize cost and operations, the same connectivity must be shared with all projects. You must ensure that all traffic between different projects, on-premises locations, and the internet can be inspected using the same third-party appliances. What should you do?

- A. Configure the third-party appliances with multiple interfaces and specific Partner Interconnect VLAN attachments per project
- B. Create the relevant routes on the third-party appliances and VPC networks.
- C. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network
- D. Create separate VPC networks for on-premises and internet connectivity
- E. Create the relevant routes on the third-party appliances and VPC networks.
- F. Consolidate all existing projects' subnetworks into a single VPC
- G. Create separate VPC networks for on-premises and internet connectivity
- H. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network
- I. Create the relevant routes on the third-party appliances and VPC networks.
- J. Configure the third-party appliances with multiple interfaces
- K. Create a hub VPC network for all projects, and create separate VPC networks for on-premises and internet connectivity
- L. Create the relevant routes on the third-party appliances and VPC network

- M. Use VPC Network Peering to connect all projects' VPC networks to the hub VP
- N. Export custom routes from the hub VPC and import on all projects' VPC networks.

Answer: D

NEW QUESTION 139

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 142

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 147

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 149

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 151

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 154

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 155

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 159

You recently deployed your application in Google Cloud. You need to verify your Google Cloud network configuration before deploying your on-premises workloads. You want to confirm that your Google Cloud network configuration allows traffic to flow from your cloud resources to your on-premises network. This validation should also analyze and diagnose potential failure points in your Google Cloud network configurations without sending any data plane test traffic. What should you do?

- A. Use Network Intelligence Center's Connectivity Tests.
- B. Enable Packet Mirroring on your application and send test traffic.
- C. Use Network Intelligence Center's Network Topology visualizations.
- D. Enable VPC Flow Logs and send test traffic.

Answer: C

NEW QUESTION 160

You are the network administrator responsible for hybrid connectivity at your organization. Your developer team wants to use Cloud SQL in the us-west1 region in your Shared VPC. You configured a Dedicated Interconnect connection and a Cloud Router in us-west1, and the connectivity between your Shared VPC and on-premises data center is working as expected. You just created the private services access connection required for Cloud SQL using the reserved IP address range and default settings. However, your developers cannot access the Cloud SQL instance from on-premises. You want to resolve the issue. What should you do?

- A. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes. Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.
- B. Change the VPC routing mode to global. Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.
- C. Create an additional Cloud Router in us-west2. Create a new Border Gateway Protocol (BGP) peering connection to your on-premises data center.
- D. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.
- E. Change the VPC routing mode to global. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.

Answer: A

NEW QUESTION 162

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

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

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

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

You are using a new user account. You must assign the appropriate identity and Access Management (IAM) user roles to this new user account before updating

the firewall rule. The new user account must be able to apply the update and view firewall logs. What should you do?

- A. Assign the compute.securityAdmin and logging.viewer rule to the new user account
- B. Apply the new firewall rule with a priority of 50.
- C. Assign the compute.securityAdmin and logging.bucketWriter role to the new user account
- D. Apply the new firewall rule with a priority of 150.
- E. Assign the compute.orgSecurityPolicyAdmin and logging.viewer role to the new user account
- 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 167

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 170

Your company's security team tends to use managed services when possible. You need to build a dashboard to show the number of deny hits that occur against configured firewall rules without increasing operational overhead. What should you do?

- A. Configure Firewall Rules Logging
- B. Use Firewall Insights to display the number of hits.
- C. Configure Firewall Rules Logging
- D. View the logs in Cloud Logging, and create a custom dashboard in Cloud Monitoring to display the number of hits.
- E. Configure a firewall appliance from the Google Cloud Marketplace
- F. Route all traffic through this appliance, and apply the firewall rules at this layer
- G. Use the firewall appliance to display the number of hits.
- H. Configure Packet Mirroring on the VPC
- I. Apply a filter with an IP address list of the Denied Firewall rule
- J. Configure an intrusion detection system (IDS) appliance as the receiver to display the number of hits.

Answer: A

NEW QUESTION 175

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 176

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 178

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 179

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

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

Answer: D

Explanation:

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

NEW QUESTION 184

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

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

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

All DNS resolution must be done on-premises.

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

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

Answer: C

NEW QUESTION 189

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

- > Your ISP is a Google Partner Interconnect provider.
- > Your on-premises VPN device's internet uplink and downlink speeds are 10 Gbps.
- > A test VPN connection between your on-premises gateway and GCP is performing at a maximum speed of 500 Mbps due to packet losses.
- > Most of the data transfer will be from GCP to the on-premises environment.
- > The application can burst up to 1.5 Gbps during peak transfers over the Interconnect.
- > Cost and the complexity of the solution should be minimal.

How should you provision the connectivity solution?

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

Answer: A

Explanation:

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

NEW QUESTION 190

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 195

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

- > IP ranges for pods and services must be as small as possible.
- > The nodes and the master must not be reachable from the internet.
- > You must be able to use kubectl commands from on-premises subnets to manage the cluster.

How should you create the GKE cluster?

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

Answer: D

Explanation:

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

NEW QUESTION 197

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

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

Answer: C

NEW QUESTION 199

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 201

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 204

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 205

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 209

.....

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](#)