

## Professional-Cloud-Network-Engineer Dumps

### Google Cloud Certified - Professional Cloud Network Engineer

<https://www.certleader.com/Professional-Cloud-Network-Engineer-dumps.html>



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

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

**Answer:** D

**NEW QUESTION 3**

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

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

**Answer:** A

**NEW QUESTION 4**

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

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

You want to integrate networking and DNS infrastructure of both organizations as quickly as possible and with minimal downtime. Which two steps should you take? (Choose two.)

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

**Answer:** BC

**Explanation:**

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

**NEW QUESTION 5**

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 6**

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

**NEW QUESTION 7**

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

**NEW QUESTION 8**

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 9**

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 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 10

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

- A. Configure a /28 primary IP address range for the node IP addresses
- 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 addresses
- 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 addresses
- 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 addresses
- K. Configure a /24 secondary IP range for the Pod
- L. Configure a /22 secondary IP range for the Services.

**Answer: A**

#### NEW QUESTION 14

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 17

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 19

You want to use Partner Interconnect to connect your on-premises network with your VPC. You already have an Interconnect partner. What should you first?

- A. Log in to your partner's portal and request the VLAN attachment there.
- B. Ask your Interconnect partner to provision a physical connection to Google.
- C. Create a Partner Interconnect type VLAN attachment in the GCP Console and retrieve the pairing key.
- D. Run `gcloud compute interconnect attachments partner update <attachment> / -- region <region>--admin-enabled`.

**Answer: B**

**Explanation:**

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

**NEW QUESTION 20**

You are designing a new global application using Compute Engine instances that will be exposed by a global HTTP(S) load balancer. You need to secure your application from distributed denial-of-service and application layer (layer 7) attacks. What should you do?

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

**Answer:** C

**NEW QUESTION 22**

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 23**

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 24**

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](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 25**

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

- A. Create the Distribution VPC in auto mode.
- B. Peer both the VPCs via network peering.
- C. Create the Distribution VPC in custom mode.
- 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 mode.
- 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 27**

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 32**

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

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

**Answer:** B

**NEW QUESTION 35**

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 40**

You need to configure the Border Gateway Protocol (BGP) session for a VPN tunnel you just created between two Google Cloud VPCs, 10.1.0.0/16 and 172.16.0.0/16. You have a Cloud Router (router-1) in the 10.1.0.0/16 network and a second Cloud Router (router-2) in the 172.16.0.0/16 network. Which configuration should you use for the BGP session?

A. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

| Router   | BGP Interface Name    | BGP IP        | BGP Peer IP   | Peer ASN |
|----------|-----------------------|---------------|---------------|----------|
| router-1 | if-tunnel-a-to-b-if-0 | 169.254.0.254 | 169.254.0.254 | 65502    |
| router-2 | if-tunnel-b-to-a-if-0 | 169.254.0.254 | 169.254.0.254 | 65501    |

B. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

| Router   | BGP Interface Name    | BGP IP     | BGP Peer IP | Peer ASN |
|----------|-----------------------|------------|-------------|----------|
| router-1 | if-tunnel-a-to-b-if-0 | 10.1.0.1   | 172.16.0.1  | 15052    |
| router-2 | if-tunnel-b-to-a-if-0 | 172.16.0.1 | 10.1.0.1    | 15501    |

C. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

| Router   | BGP Interface Name    | BGP IP       | BGP Peer IP  | Peer ASN |
|----------|-----------------------|--------------|--------------|----------|
| router-1 | if-tunnel-a-to-b-if-0 | 169.254.20.1 | 169.254.20.2 | 65002    |
| router-2 | if-tunnel-b-to-a-if-0 | 169.254.20.2 | 169.254.20.1 | 65001    |

D. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

| Router   | BGP Interface Name    | BGP IP       | BGP Peer IP  | Peer ASN |
|----------|-----------------------|--------------|--------------|----------|
| router-1 | if-tunnel-a-to-b-if-0 | 172.16.0.254 | 10.1.0.254   | 16552    |
| router-2 | if-tunnel-b-to-a-if-0 | 10.1.0.254   | 172.16.0.254 | 16551    |

**Answer:** C

**NEW QUESTION 44**

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 46**

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 caused autoscaling to reach the maximum number of instances, and users of your application cannot complete transactions. After an investigation, you think it is a DDOS attack. You want to quickly restore user access to your application and allow successful transactions while minimizing cost.

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

- A. Use Cloud Armor to blacklist the attacker's IP addresses.
- B. Increase the maximum autoscaling backend to accommodate the severe bursty traffic.
- C. Create a global HTTP(s) load balancer and move your application backend to this load balancer.
- D. Shut down the entire application in GCP for a few hours.
- 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 49**

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 54**

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 58**

Your end users are located in close proximity to `us-east1` and `eu-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 subnets
- F. Create a global load balancer to establish connectivity between the regions.
- G. Create 1 VPC with 2 regional subnets
- 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 `eu-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 59**

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 63**

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 67**

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

What should you do?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 70**

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 to communicate 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 73

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 75

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 76**

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

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

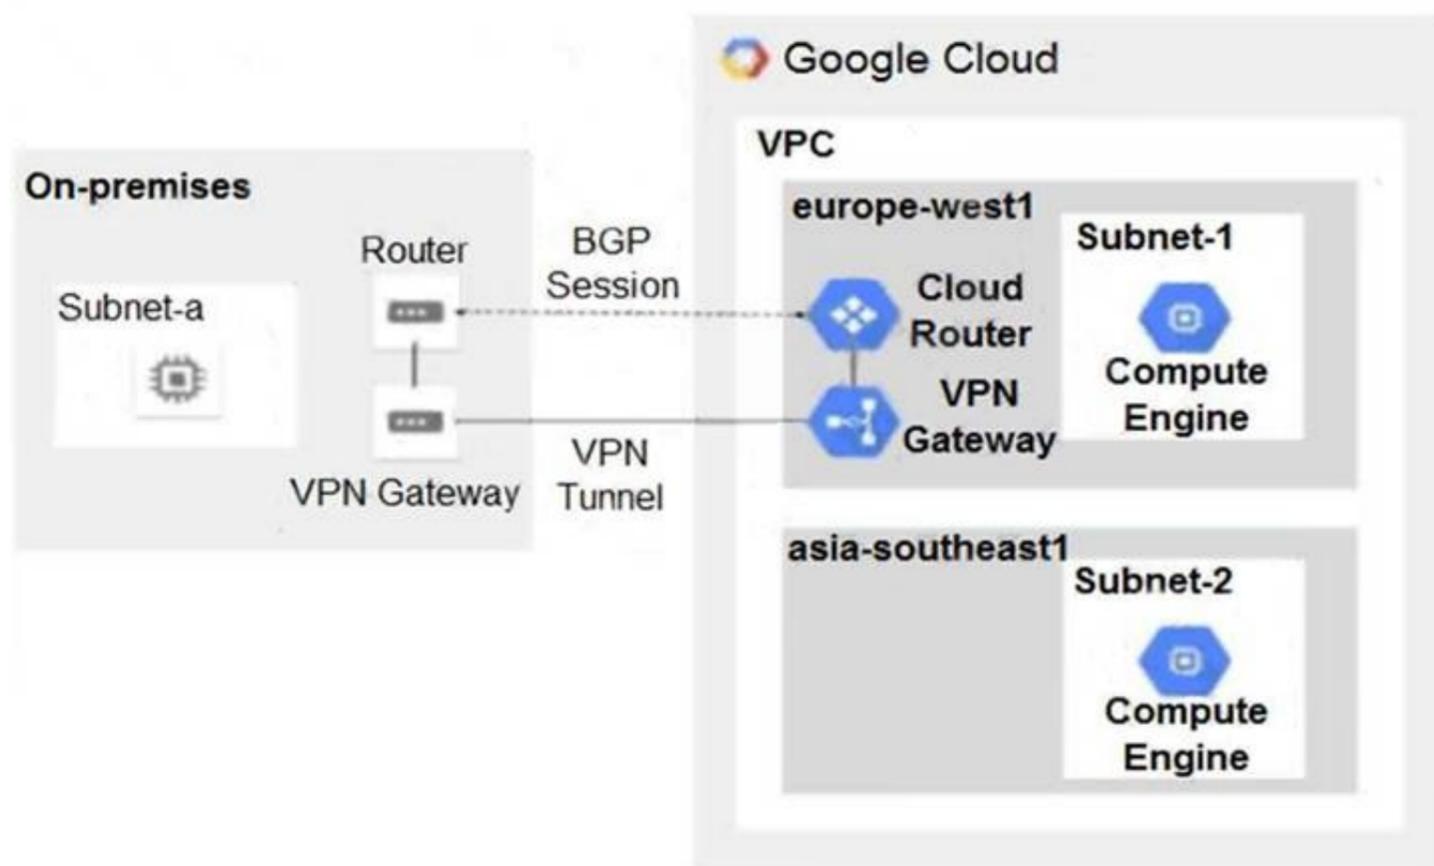
**Answer: D**

**Explanation:**

"TFTP is a UDP-based protocol. Servers listen on port 69 for the initial client-to-server packet to establish the TFTP session, then use a port above 1023 for all further packets during that session. Clients use ports above 1023" [https://docstore.mik.ua/oreilly/networking\\_2ndEd/fire/ch17\\_02.htm](https://docstore.mik.ua/oreilly/networking_2ndEd/fire/ch17_02.htm) Besides, Google Cloud external TCP/UDP Network Load Balancing (after this referred to as Network Load Balancing) is a regional, non-proxied load balancer. Network Load Balancing distributes traffic among virtual machine (VM) instances in the same region in a Virtual Private Cloud (VPC) netw

**NEW QUESTION 77**

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 81**

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 85**

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

- A. Keep the existing Dedicated interconnect
- B. Deploy a VLAN attachment to a Cloud Router in us-west2, and use VPC global routing to access workloads in us-east4 and us-central1.
- C. Keep the existing Dedicated Interconnect
- D. Deploy a VLAN attachment to a Cloud Router in us-east4, and deploy another VLAN attachment to a Cloud Router in us-central1.
- E. Order a new Dedicated Interconnect for a colocation facility closest to us-east4, and use VPC globalrouting to access workloads in us-central1.
- F. Order a new Dedicated Interconnect for a colocation facility closest to us-central1, and use VPC global routing to access workloads in us-east4.

**Answer: C**

#### NEW QUESTION 86

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 90

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 91

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

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

**Answer: A**

#### NEW QUESTION 94

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 98

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 102

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](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](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 105**

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 109**

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

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

**Answer:** B

**NEW QUESTION 110**

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 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 116**

You have several microservices running in a private subnet in an existing Virtual Private Cloud (VPC). You need to create additional serverless services that use Cloud Run and Cloud Functions to access the

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

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

I. Configure the serverless service to use the connector for communication to the microservices.

**Answer:** D

**NEW QUESTION 117**

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 role 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 118**

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 Loggin
- B. Use Firewall Insights to display the number of hits.
- C. Configure Firewall Rules Loggin
- 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 Marketplac
- F. Route all traffic through this appliance, and apply the firewall rules at this laye
- G. Use the firewall appliance to display the number of hits.
- H. Configure Packet Mirroring on the VP
- 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 121**

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

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

**Answer:** A

**Explanation:**

The service range setting is permanent and cannot be changed. Please see <https://stackoverflow.com/questions/60957040/how-to-increase-the-service-address-range-of-a-gke-cluster> I think the correc tanswer is A since: Grow is expected to up to 100 nodes (that would be /25), then up to 200 pods per node (100 times 200 = 20000 so /17 is 32768), then 1500 services in a /21 (up to 2048) <https://docs.netgate.com/pfsense/en/latest/book/network/understanding-cidr-subnet-mask-notation.html>

**NEW QUESTION 122**

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 i
- 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 125**

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 129**

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

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

**Answer:** B

**NEW QUESTION 130**

You are designing the network architecture for your organization. Your organization has three developer teams: Web, App, and Database. All of the developer teams require access to Compute Engine instances to perform their critical tasks. You are part of a small network and security team that needs to provide network access to the developers. You need to maintain centralized control over network resources, including subnets, routes, and firewalls. You want to minimize operational overhead. How should you design this topology?

- A. Configure a host project with a Shared VPC
- B. Create service projects for Web, App, and Database.
- C. Configure one VPC for Web, one VPC for App, and one VPC for Database
- D. Configure HA VPN between each VPC.
- E. Configure three Shared VPC host projects, each with a service project: one for Web, one for App, and one for Database.
- F. Configure one VPC for Web, one VPC for App, and one VPC for Database
- G. Use VPC Network Peering to connect all VPCs in a full mesh.

**Answer:** C

**NEW QUESTION 132**

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 rule
- 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 forwarding rule
- E. Then, define an A record in Cloud DNS
- 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 135**

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

What should you do?

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

**Answer: C**

**Explanation:**

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

**NEW QUESTION 137**

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 172.16.45.0/24 network.

**Answer: C**

**NEW QUESTION 142**

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 147**

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 148**

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 149**

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

What should you do?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 153**

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 158**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your Professional-Cloud-Network-Engineer Exam with Our Prep Materials Via below:**

<https://www.certleader.com/Professional-Cloud-Network-Engineer-dumps.html>