



Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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 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 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

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 5

You are responsible for configuring firewall policies for your company in Google Cloud. Your security team has a strict set of requirements that must be met to configure firewall rules. Always allow Secure Shell (SSH) from your corporate IP address. Restrict SSH access from all other IP addresses. There are multiple projects and VPCs in your Google Cloud organization. You need to ensure that other VPC firewall rules cannot bypass the security team's requirements. What should you do?

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

Answer: A

NEW QUESTION 6

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 7

You 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 8

You recently deployed two network virtual appliances in us-central1. Your network appliances provide connectivity to your on-premises network, 10.0.0.0/8. You need to configure the routing for your Virtual Private Cloud (VPC). Your design must meet the following requirements:
All access to your on-premises network must go through the network virtual appliances. Allow on-premises access in the event of a single network virtual appliance failure.
Both network virtual appliances must be used simultaneously. Which method should you use to accomplish this?

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

Answer: B

NEW QUESTION 9

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 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 10

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 12

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 13

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 16

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

- A. Use the Network Intelligence Center Connectivity Tests to test the connectivity between the VPC and the on-premises network.
- B. Use Network Intelligence Center Network Topology to check the traffic flow, and replay the traffic from the time period when the connectivity issue occurred.
- C. Configure VPC Flow Log
- D. Review the logs by filtering on the source and destination.
- E. Configure a Compute Engine instance on the same VPC as the service running on Google Cloud to run a traceroute targeted at the on-premises service.

Answer: B

NEW QUESTION 18

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 20

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 21

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 22

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

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

Answer: B

Explanation:

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

NEW QUESTION 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 28

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 33

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

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

Answer: B

NEW QUESTION 34

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 38

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

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

Answer: BE

NEW QUESTION 39

Your on-premises data center has 2 routers connected to your GCP through a VPN on each router. All applications are working correctly; however, all of the traffic is passing across a single VPN instead of being load-balanced across the 2 connections as desired. During troubleshooting you find:

- Each on-premises router is configured with the same ASN.
- Each on-premises router is configured with the same routes and priorities.
- Both on-premises routers are configured with a VPN connected to a single Cloud Router.
- The VPN logs have no-proposal-chosen lines when the VPNs are connecting.
- BGP session is not established between one on-premises router and the Cloud Router. What is the most likely cause of this problem?

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

Answer: A

Explanation:

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

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

NEW QUESTION 42

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

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

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

Answer: AB

Explanation:

A: Using VPC Flow Logs VPC Flow Logs records a sample of network flows sent from and received by VM instances, including instances used as GKE nodes.

These logs can be used for network monitoring, forensics, real-time security analysis, and expense optimization. <https://cloud.google.com/vpc/docs/using-flow-logs>

(B): Firewall Rules Logging overview Firewall Rules Logging allows you to audit, verify, and analyze the effects of your firewall rules. For example, you can determine if a firewall rule designed to deny traffic is functioning

as intended. Firewall Rules Logging is also useful if you need to determine how many connections are affected by a given firewall rule. You enable Firewall Rules Logging individually for each firewall rule whose connections you need to log. Firewall Rules Logging is an option for any firewall rule, regardless of the action (allow or deny) or direction (ingress or egress) of the rule.

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

NEW QUESTION 44

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 45

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 47

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 49

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 50

You have the following private Google Kubernetes Engine (GKE) cluster deployment:

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

...

```
clusterIpv4Cidr: 192.168.36.0/24
endpoint: 192.168.38.2
ipAllocationPolicy:
  clusterIpv4Cidr: 192.168.36.0/24
  clusterIpv4CidrBlock: 192.168.36.0/24
  clusterSecondaryRangeName: customer-1-pods
  servicesIpv4Cidr: 192.168.37.0/24
  servicesIpv4CidrBlock: 192.168.37.0/24
  servicesSecondaryRangeName: customer-1-svc
  useIpAliases: true
```

...

```
masterAuthorizedNetworksConfig:
```

...

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

...

```
servicesIpv4Cidr: 192.162.37.0/24
```

...

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

You have a virtual machine (VM) deployed in the same VPC in the subnetwork 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 masterAuthorizedNetworksConfig
- B. Configure kubectl to communicate with the endpoint 192.168.38.2.
- C. Add the network 192.168.38.0/28 to the masterAuthorizedNetworksConfig
- D. Configure kubectl to communicate with the endpoint 192.168.38.2
- E. Add the network 192.168.36.0/24 to the masterAuthorizedNetworksConfig
- 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 53

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 54

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 56

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 60

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

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

Answer: D

NEW QUESTION 63

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 64

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 65

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 VPC
- N. Export custom routes from the hub VPC and import on all projects' VPC networks.

Answer: D

NEW QUESTION 66

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 68

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 70

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 71

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 76

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 78

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 83

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 86

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 90

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 91

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 94

.....

Relate Links

100% Pass Your Professional-Cloud-Network-Engineer Exam with ExamBible Prep Materials

<https://www.exambible.com/Professional-Cloud-Network-Engineer-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>