

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 built a web application with several containerized microservices. You want to run those microservices on Cloud Run. You must also ensure that the services are highly available to your customers with low latency. What should you do?

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

Answer: B

NEW QUESTION 3

You have deployed a new internal application that provides HTTP and TFTP services to on-premises hosts. You want to be able to distribute traffic across multiple Compute Engine instances, but need to ensure that clients are sticky to a particular instance across both services. Which session affinity should you choose?

- A. None
- B. Client IP
- C. Client IP and protocol
- D. Client IP, port and protocol

Answer: B

NEW QUESTION 4

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

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

Answer: D

NEW QUESTION 5

You have an application that is running in a managed instance group. Your development team has released an updated instance template which contains a new feature which was not heavily tested. You want to minimize impact to users if there is a bug in the new template. How should you update your instances?

- A. Manually patch some of the instances, and then perform a rolling restart on the instance group.
- B. Using the new instance template, perform a rolling update across all instances in the instance group. Verify the new feature once the rollout completes.
- C. Deploy a new instance group and canary the updated template in that group
- D. Verify the new feature in the new canary instance group, and then update the original instance group.
- E. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template
- F. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Answer: D

Explanation:

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

NEW QUESTION 6

You have ordered Dedicated Interconnect in the GCP Console and need to give the Letter of Authorization/Connecting Facility Assignment (LOA-CFA) to your cross-connect provider to complete the physical connection.
Which two actions can accomplish this? (Choose two.)

- A. Open a Cloud Support ticket under the Cloud Interconnect category.
- B. Download the LOA-CFA from the Hybrid Connectivity section of the GCP Console.
- C. Run `gcloud compute interconnects describe <interconnect>`.
- D. Check the email for the account of the NOC contact that you specified during the ordering process.
- E. Contact your cross-connect provider and inform them that Google automatically sent the LOA/CFA to them via email, and to complete the connection.

Answer: DE

Explanation:

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

NEW QUESTION 7

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

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

Answer: D

Explanation:

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

NEW QUESTION 8

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 9

You just finished your company's migration to Google Cloud and configured an architecture with 3 Virtual Private Cloud (VPC) networks: one for Sales, one for Finance, and one for Engineering. Every VPC contains over 100 Compute Engine instances, and now developers using instances in the Sales VPC and the Finance VPC require private connectivity between each other. You need to allow communication between Sales and Finance without compromising performance or security. What should you do?

- A. Configure an HA VPN gateway between the Finance VPC and the Sales VPC.
- B. Configure the instances that require communication between each other with an external IP address.
- C. Create a VPC Network Peering connection between the Finance VPC and the Sales VPC.
- D. Configure Cloud NAT and a Cloud Router in the Sales and Finance VPCs.

Answer: C

NEW QUESTION 10

You 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

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 12

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 14

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 address
- 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 address
- 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 address
- 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 address
- K. Configure a /24 secondary IP range for the Pod
- L. Configure a /22 secondary IP range for the Services.

Answer: A

NEW QUESTION 17

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

What should you do on your on-premises servers?

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

Answer: A

Explanation:

<https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid> <https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid>

<https://cloud.google.com/blog/products/gcp/5-steps-to-better-gcp-network-performance?hl=ml>

NEW QUESTION 19

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

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

Answer: B

NEW QUESTION 23

You want to establish a dedicated connection to Google that can access Cloud SQL via a public IP address and that does not require a third-party service provider. Which connection type should you choose?

- A. Carrier Peering
- B. Direct Peering
- C. Dedicated Interconnect
- D. Partner Interconnect

Answer: B

Explanation:

When established, Direct Peering provides a direct path from your on-premises network to Google services, including Google Cloud products that can be exposed through one or more public IP addresses. Traffic from Google's network to your on-premises network also takes that direct path, including traffic from VPC networks in your projects. Google Cloud customers must request that direct egress pricing be enabled for each of their projects after they have established Direct Peering with Google. For more information, see Pricing.

NEW QUESTION 25

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 26

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 28

Your organization uses a Shared VPC architecture with a host project and three service projects. You have Compute Engine instances that reside in the service projects. You have critical workloads in your on-premises data center. You need to ensure that the Google Cloud instances can resolve on-premises hostnames via the Dedicated Interconnect you deployed to establish hybrid connectivity. What should you do?

- A. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 35.199.192.0/19 to the on-premises environment.
- B. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the Private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 169.254 169.254 to the on-premises environment.
- C. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. In your Cloud Router, add a custom route advertisement for the IP 169.254 169 254 to the on-premises environment.
- D. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. Configure a DNS policy in the Shared VPC to allow inbound query forwarding with your on-premises DNS server as the alternative DNS server.

Answer: D

NEW QUESTION 31

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

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

Answer: C

Explanation:

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

NEW QUESTION 35

You have created an HTTP(S) load balanced service. You need to verify that your backend instances are responding properly. How should you configure the health check?

- A. Set request-path to a specific URL used for health checking, and set proxy-header to PROXY_V1.
- B. Set request-path to a specific URL used for health checking, and set host to include a custom host header that identifies the health check.
- C. Set request-path to a specific URL used for health checking, and set response to a string that the backend service will always return in the response body.
- D. Set proxy-header to the default value, and set host to include a custom host header that identifies the health check.

Answer: C

Explanation:

https://cloud.google.com/load-balancing/docs/health-check-concepts#content-based_health_checks

NEW QUESTION 36

You have created a firewall with rules that only allow traffic over HTTP, HTTPS, and SSH ports. While testing, you specifically try to reach the server over multiple ports and protocols; however, you do not see any denied connections in the firewall logs. You want to resolve the issue. What should you do?

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

Answer: D

Explanation:

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

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

NEW QUESTION 40

Your company has a security team that manages firewalls and SSL certificates. It also has a networking team that manages the networking resources. The networking team needs to be able to read firewall rules, but should not be able to create, modify, or delete them. How should you set up permissions for the networking team?

- A. Assign members of the networking team the compute.networkUser role.
- B. Assign members of the networking team the compute.networkAdmin role.
- C. Assign members of the networking team a custom role with only the compute.networks.* and the compute.firewalls.list permissions.
- D. Assign members of the networking team the compute.networkViewer role, and add the compute.networks.use permission.

Answer: B

NEW QUESTION 44

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 45

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 47

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

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

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

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

Answer: A

NEW QUESTION 51

You created a new VPC for your development team. You want to allow access to the resources in this VPC via SSH only.

How should you configure your firewall rules?

- A. Create two firewall rules: one to block all traffic with priority 0, and another to allow port 22 with priority 1000.
- B. Create two firewall rules: one to block all traffic with priority 65536, and another to allow port 3389 with priority 1000.
- C. Create a single firewall rule to allow port 22 with priority 1000.
- D. Create a single firewall rule to allow port 3389 with priority 1000.

Answer: C

NEW QUESTION 55

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 57

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 59

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 64

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 69

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 70

You are responsible for designing a new connectivity solution for your organization's enterprise network to access and use Google Workspace. You have an existing Shared VPC with Compute Engine instances in us-west1. Currently, you access Google Workspace via your service provider's internet access. You want to set up a direct connection between your network and Google. What should you do?

- A. Order a Dedicated Interconnect connection in the same metropolitan area
- B. Create a VLAN attachment, a Cloud Router in us-west1, and a Border Gateway Protocol (BGP) session between your Cloud Router and your router.
- C. Order a Direct Peering connection in the same metropolitan area
- D. Configure a Border Gateway Protocol (BGP) session between Google and your router.
- E. Configure HA VPN in us-west1. Configure a Border Gateway Protocol (BGP) session between your Cloud Router and your on-premises data center.
- F. Order a Carrier Peering connection in the same metropolitan area
- G. Configure a Border Gateway Protocol (BGP) session between Google and your router.

Answer: B

NEW QUESTION 72

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

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

Answer: C

NEW QUESTION 77

You are designing a hub-and-spoke network architecture for your company's cloud-based environment. You need to make sure that all spokes are peered with the hub. The spokes must use the hub's virtual appliance for internet access.

The virtual appliance is configured in high-availability mode with two instances using an internal load balancer with IP address 10.0.0.5. What should you do?

- A. Create a default route in the hub VPC that points to IP address 10.0.0.5. Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway. Export the custom routes in the hub.
- B. Import the custom routes in the spokes.
- C. Create a default route in the hub VPC that points to IP address 10.0.0.5. Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway. Export the custom routes in the hub.
- D. Import the custom routes in the spoke.
- E. Delete the default internet gateway route of the spokes.
- F. Create two default routes in the hub VPC that point to the next hop instances of the virtual appliances. Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway. Export the custom routes in the hub.
- G. Import the custom routes in the spokes.
- H. Create a default route in the hub VPC that points to IP address 10.0.0.5. Delete the default internet gateway route in the hub VPC, and create a new higher-priority route that is tagged only to the appliances with a next hop of the default internet gateway. Create a new route in the spoke VPC that points to IP address 10.0.0.5.

Answer: B

NEW QUESTION 79

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 81

You need to give each member of your network operations team least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments. What should you do?

- A. Assign each user the editor role.
- B. Assign each user the compute.networkAdmin role.
- C. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get.
- D. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get, compute.routers.create, compute.routers.get, compute.routers.update.

Answer: D

Explanation:

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

NEW QUESTION 82

Your company offers a popular gaming service. Your instances are deployed with private IP addresses, and external access is granted through a global load balancer. You have recently engaged a traffic-scrubbing service and want to restrict your origin to allow connections only from the traffic-scrubbing service. What should you do?

- A. Create a Cloud Armor Security Policy that blocks all traffic except for the traffic-scrubbing service.
- B. Create a VPC Firewall rule that blocks all traffic except for the traffic-scrubbing service.
- C. Create a VPC Service Control Perimeter that blocks all traffic except for the traffic-scrubbing service.
- D. Create IPTables firewall rules that block all traffic except for the traffic-scrubbing service.

Answer: A

Explanation:

Global load balancer will proxy the connection . thus no trace of session origin IP. you should use Cloud Armor to geofence your service.

<https://cloud.google.com/load-balancing/docs/https>

NEW QUESTION 84

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 88

You work for a multinational enterprise that is moving to GCP. These are the cloud requirements:

- An on-premises data center located in the United States in Oregon and New York with Dedicated Interconnects connected to Cloud regions us-west1 (primary HQ) and us-east4 (backup)
- Multiple regional offices in Europe and APAC
- Regional data processing is required in europe-west1 and australia-southeast1
- Centralized Network Administration Team

Your security and compliance team requires a virtual inline security appliance to perform L7 inspection for URL filtering. You want to deploy the appliance in us-west1.

What should you do?

- A. • Create 2 VPCs in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Host Project. • Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. • Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- B. • Create 2 VPCs in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Service Project. • Attach NIC0 in VPC #1 us-west1 subnet of the Host Project. • Attach NIC1 in VPC #2 us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- C. • Create 1 VPC in a Shared VPC Host Project. • Configure a 2-NIC instance in zone us-west1-a in the Host Project. • Attach NIC0 in us-west1 subnet of the Host Project. • Attach NIC1 in us-west1 subnet of the Host Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.
- D. • Create 1 VPC in a Shared VPC Service Project. • Configure a 2-NIC instance in zone us-west1-a in the Service Project. • Attach NIC0 in us-west1 subnet of the Service Project. • Attach NIC1 in us-west1 subnet of the Service Project. • Deploy the instance. • Configure the necessary routes and firewall rules to pass traffic through the instance.

Answer: B

Explanation:

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

NEW QUESTION 89

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 91

You recently deployed Cloud VPN to connect your on-premises data center to Google Cloud. You need to monitor the usage of this VPN and set up alerts in case traffic exceeds the maximum allowed. You need to be able to quickly decide whether to add extra links or move to a Dedicated Interconnect. What should you do?

- A. In the Network Intelligence Center, check for the number of packet drops on the VPN.
- B. In the Google Cloud Console, use Monitoring Query Language to create a custom alert for bandwidth utilization.
- C. In the Monitoring section of the Google Cloud Console, use the Dashboard section to select a default dashboard for VPN usage.
- D. In the VPN section of the Google Cloud Console, select the VPN under hybrid connectivity, and then select monitoring to display utilization on the dashboard.

Answer: A

NEW QUESTION 93

You want to apply a new Cloud Armor policy to an application that is deployed in Google Kubernetes Engine (GKE). You want to find out which target to use for your Cloud Armor policy.

Which GKE resource should you use?

- A. GKE Node
- B. GKE Pod
- C. GKE Cluster
- D. GKE Ingress

Answer: D

Explanation:

Cloud Armour is applied at load balancers Configuring Google Cloud Armor through Ingress. <https://cloud.google.com/kubernetes-engine/docs/how-to/ingress-features> Security policy features Google Cloud Armor security policies have the following core features: You can optionally use the QUIC protocol with load balancers that use Google Cloud Armor. You can use Google Cloud Armor with external HTTP(S) load balancers that are in either Premium Tier or Standard Tier. You can use security policies with GKE and the default Ingress controller.

NEW QUESTION 98

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

- A. Configure the remote autonomous system number (ASN) to 4096.
- B. Configure a second Cloud Router to scale bandwidth in and out of the VPC.
- C. Configure the maximum transmission unit (MTU) to its highest supported value.
- D. Configure a second set of active/passive VPN tunnels.

Answer: D

NEW QUESTION 101

One instance in your VPC is configured to run with a private IP address only. You want to ensure that even if this instance is deleted, its current private IP address will not be automatically assigned to a different instance.

In the GCP Console, what should you do?

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

Answer: C

Explanation:

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

NEW QUESTION 102

Your organization is implementing a new security policy to control how firewall rules are applied to control flows between virtual machines (VMs). Using Google-recommended practices, you need to set up a firewall rule to enforce strict control of traffic between VM A and VM B. You must ensure that communications flow only from VM A to VM B within the VPC, and no other communication paths are allowed. No other firewall rules exist in the VPC. Which firewall rule should you configure to allow only this communication path?

- A. Firewall rule direction: ingress Action: allowTarget: VM B service accountSource ranges: VM A service account Priority: 1000
- B. Firewall rule direction: ingress Action: allowTarget: specific VM B tagSource ranges: VM A tag and VM A source IP address Priority: 1000
- C. Firewall rule direction: ingress Action: allowTarget: VM A service accountSource ranges: VM B service account and VM B source IP address Priority: 100
- D. Firewall rule direction: ingress Action: allowTarget: specific VM A tagSource ranges: VM B tag and VM B source IP address Priority: 100

Answer: D

NEW QUESTION 105

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 108

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 interconnec
- 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 Interconnec
- 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 112

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 116

You are using the gcloud command line tool to create a new custom role in a project by coping 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 121

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 125

Your organization uses a hub-and-spoke architecture with critical Compute Engine instances in your Virtual Private Clouds (VPCs). You are responsible for the design of Cloud DNS in Google Cloud. You need to be able to resolve Cloud DNS private zones from your on-premises data center and enable on-premises name resolution from your hub-and-spoke VPC design. What should you do?

- A. Configure a private DNS zone in the hub VPC, and configure DNS forwarding to the on-premises server. Configure DNS peering from the spoke VPCs to the hub VPC.
- B. Configure a DNS policy in the hub VPC to allow inbound query forwarding from the spoke VPCs. Configure the spoke VPCs with a private zone, and set up DNS peering to the hub VPC.
- C. Configure a DNS policy in the spoke VPCs, and configure your on-premises DNS as an alternate DNS server. Configure the hub VPC with a private zone, and set up DNS peering to each of the spoke VPCs.
- D. Configure a DNS policy in the hub VPC, and configure the on-premises DNS as an alternate DNS server. Configure the spoke VPCs with a private zone, and set up DNS peering to the hub VPC.

Answer: C

NEW QUESTION 127

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

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

Answer: A

NEW QUESTION 131

You are designing a Partner Interconnect hybrid cloud connectivity solution with geo-redundancy across two metropolitan areas. You want to follow Google-recommended practices to set up the following region/metro pairs:
(region 1/metro 1)
(region 2/metro 2) What should you do?

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

Answer: B

NEW QUESTION 132

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 134

You are disabling DNSSEC for one of your Cloud DNS-managed zones. You removed the DS records from your zone file, waited for them to expire from the cache, and disabled DNSSEC for the zone. You receive reports that DNSSEC validating resolves are unable to resolve names in your zone. What should you do?

- A. Update the TTL for the zone.
- B. Set the zone to the TRANSFER state.
- C. Disable DNSSEC at your domain registrar.
- D. Transfer ownership of the domain to a new registrar.

Answer: C

Explanation:

Before disabling DNSSEC for a managed zone you want to use, you must deactivate DNSSEC at your domain registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.

NEW QUESTION 135

You are creating a new application and require access to Cloud SQL from VPC instances without public IP addresses. Which two actions should you take? (Choose two.)

- A. Activate the Service Networking API in your project.
- B. Activate the Cloud Datastore API in your project.
- C. Create a private connection to a service producer.
- D. Create a custom static route to allow the traffic to reach the Cloud SQL API.
- E. Enable Private Google Access.

Answer: CE

Explanation:

https://cloud.google.com/sql/docs/mysql/configure-private-services-access#console_1

C: If you are using private IP for any of your Cloud SQL instances, you only need to configure private services access one time for every Google Cloud project that has or needs to connect to a Cloud SQL instance. If your Google Cloud project has a Cloud SQL instance, you can either configure it yourself or let Cloud SQL do it for you to use private IP. Cloud SQL configures private services access for you when all the conditions below are true:

https://cloud.google.com/sql/docs/postgres/configure-private-services-access#before_you_begin

E: You can enable Private Google access on a subnet level and any VMs on that subnet can access Google APIs by using their internal IP address.

<https://cloud.google.com/vpc/docs/configure-private-google-access>

NEW QUESTION 139

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 140

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

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

Answer: C

NEW QUESTION 143

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 144

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 145

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 148

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 correct answer 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 150

Your company has recently expanded their EMEA-based operations into APAC. Globally distributed users report that their SMTP and IMAP services are slow. Your company requires end-to-end encryption, but you do not have access to the SSL certificates. Which Google Cloud load balancer should you use?

- A. SSL proxy load balancer
- B. Network load balancer

- C. HTTPS load balancer
- D. TCP proxy load balancer

Answer: D

Explanation:

<https://cloud.google.com/security/encryption-in-transit/> Automatic encryption between GFEs and backends For the following load balancer types, Google automatically encrypts traffic between Google Front Ends (GFEs) and your backends that reside within Google Cloud VPC networks: HTTP(S) Load Balancing TCP Proxy Load Balancing SSL Proxy Load Balancing

NEW QUESTION 151

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

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

Answer: A

NEW QUESTION 153

In your project my-project, you have two subnets in a Virtual Private Cloud (VPC): subnet-a with IP range 10.128.0.0/20 and subnet-b with IP range 172.16.0.0/24. You need to deploy database servers in subnet-a. You will also deploy the application servers and web servers in subnet-b. You want to configure firewall rules that only allow database traffic from the application servers to the database servers. What should you do?

- A. Create network tag app-server and service account sa-db@my-project.iam.gserviceaccount.com
- B. Add the tag to the application servers, and associate the service account with the database server
- C. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules top:3306 --source-tags app-server --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- D. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com
- E. Associate service account sa-app with the application servers, and associate the service account sa-db with the database server
- F. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-service-accounts sa-app@demo-cloud-idp-demo.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- G. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com
- H. Associate the service account sa-app with the application servers, and associate the service account sa-db with the database server
- I. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-ranges 10.128.0.0/20 --source-service-accounts sa-app@my-project.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- J. Create network tags app-server and db-server
- K. Add the app-server tag to the application servers, and add the db-server tag to the database server
- L. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules tcp:3306 --source-ranges 10.128.0.0/20 --source-tags app-server --target-tags db-server`

Answer: D

NEW QUESTION 156

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 158

Your company's on-premises network is connected to a VPC using a Cloud VPN tunnel. You have a static route of 0.0.0.0/0 with the VPN tunnel as its next hop defined in the VPC. All internet bound traffic currently passes through the on-premises network. You configured Cloud NAT to translate the primary IP addresses of Compute Engine instances in one region. Traffic from those instances will now reach the internet directly from their VPC and not from the on-premises network. Traffic from the virtual machines (VMs) is not translating addresses as expected. What should you do?

- A. Lower the TCP Established Connection Idle Timeout for the NAT gateway.
- B. Add firewall rules that allow ingress and egress of the external NAT IP address, have a target tag that is on the Compute Engine instances, and have a priority value higher than the priority value of the default route to the VPN gateway.
- C. Add a default static route to the VPC with the default internet gateway as the next hop, the network tag associated with the Compute Engine instances, and a higher priority than the priority of the default route to the VPN tunnel.

D. Increase the default min-ports-per-vm setting for the Cloud NAT gateway.

Answer: A

NEW QUESTION 161

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 164

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 165

Your company's Google Cloud-deployed, streaming application supports multiple languages. The application development team has asked you how they should support splitting audio and video traffic to different backend Google Cloud storage buckets. They want to use URL maps and minimize operational overhead. They are currently using the following directory structure:

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

Which solution should you recommend?

- A. Rearrange the directory structure, create a URL map and leverage a path rule such as `/video/*` and `/audio/*`.
- B. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as `/video/*` and `/audio/*`.
- C. Leave the directory structure as-is, create a URL map and leverage a path rule such as `V[a-z]{2}Vvideo` and `V[a-z]{2}Vaudio`.
- 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 169

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 173

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 174

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 175

You are using a third-party next-generation firewall to inspect traffic. You created a custom route of 0.0.0.0/0 to route egress traffic to the firewall. You want to allow your VPC instances without public IP addresses to access the BigQuery and Cloud Pub/Sub APIs, without sending the traffic through the firewall.

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

- A. Turn on Private Google Access at the subnet level.
- B. Turn on Private Google Access at the VPC level.
- C. Turn on Private Services Access at the VPC level.
- D. Create a set of custom static routes to send traffic to the external IP addresses of Google APIs and services via the default internet gateway.
- E. Create a set of custom static routes to send traffic to the internal IP addresses of Google APIs and services via the default internet gateway.

Answer: AD

Explanation:

<https://cloud.google.com/vpc/docs/private-access-options#pga> Private Google Access VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the `_external IP addresses_` of Google APIs and services.

NEW QUESTION 177

.....

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>