



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

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

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

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

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

Answer: BC

Explanation:

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

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

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

Answer: A

NEW QUESTION 5

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 6

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

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

Answer: D

Explanation:

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

NEW QUESTION 7

You need to enable Cloud CDN for all the objects inside a storage bucket. You want to ensure that all the object in the storage bucket can be served by the CDN. What should you do in the GCP Console?

- A. Create a new cloud storage bucket, and then enable Cloud CDN on it.
- B. Create a new TCP load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- C. Create a new SSL proxy load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- D. Create a new HTTP load balancer, select the storage bucket as a backend, enable Cloud CDN on the backend, and make sure each object inside the storage bucket is shared publicly.

Answer: D

Explanation:

https://cloud.google.com/load-balancing/docs/https/adding-backend-buckets-to-load-balancers#using_cloud_cdn Cloud CDN needs HTTP(S) Load Balancers and Cloud Storage bucket has to be shared publicly.
<https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket>

NEW QUESTION 8

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 9

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 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 10

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 15

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

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

Answer: AE

Explanation:

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

NEW QUESTION 18

Your organization's security policy requires that all internet-bound traffic return to your on-premises data center through HA VPN tunnels before egressing to the internet, while allowing virtual machines (VMs) to leverage private Google APIs using private virtual IP addresses 199.36.153.4/30. You need to configure the routes to enable these traffic flows. What should you do?

- A. Configure a custom route 0.0.0.0/0 with a priority of 500 whose next hop is the default internet gateway. Configure another custom route 199.36.153.4/30 with priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.
- B. Configure a custom route 0.0.0.0/0 with a priority of 1000 whose next hop is the internet gateway. Configure another custom route 199.36.153.4/30 with a priority of 500 whose next hop is the VPN tunnel back to the on-premises data center.
- C. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 1000. Configure a custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the default internet gateway.
- D. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 500. Configure another custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.

Answer: A

NEW QUESTION 19

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 21

You want to use Cloud Interconnect to connect your on-premises network to a GCP VPC. You cannot meet Google at one of its point-of-presence (POP) locations, and your on-premises router cannot run a Border Gateway Protocol (BGP) configuration.
Which connectivity model should you use?

- A. Direct Peering
- B. Dedicated Interconnect
- C. Partner Interconnect with a layer 2 partner
- D. Partner Interconnect with a layer 3 partner

Answer: D

Explanation:

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview>

For Layer 3 connections, your service provider establishes a BGP session between your Cloud Routers and their edge routers for each VLAN attachment. You don't need to configure BGP on your on-premises router. Google and your service provider automatically set the correct configurations.

<https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#connectivity-type>

NEW QUESTION 23

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 27

You have a storage bucket that contains the following objects:

- folder-a/image-a-1.jpg
- folder-a/image-a-2.jpg
- folder-b/image-b-1.jpg
- folder-b/image-b-2.jpg

Cloud CDN is enabled on the storage bucket, and all four objects have been successfully cached. You want to remove the cached copies of all the objects with the prefix folder-a, using the minimum number of commands.

What should you do?

- A. Add an appropriate lifecycle rule on the storage bucket.
- B. Issue a cache invalidation command with pattern /folder-a/*.
- C. Make sure that all the objects with prefix folder-a are not shared publicly.
- D. Disable Cloud CDN on the storage bucket.
- E. Wait 90 second
- F. Re-enable Cloud CDN on the storage bucket.

Answer: B

Explanation:

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

NEW QUESTION 32

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 35

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 40

You are designing a hybrid cloud environment for your organization. Your Google Cloud environment is interconnected with your on-premises network using Cloud HA VPN and Cloud Router. The Cloud Router is configured with the default settings. Your on-premises DNS server is located at 192.168.20.88 and is protected by a firewall, and your Compute Engine resources are located at 10.204.0.0/24. Your Compute Engine resources need to resolve on-premises private hostnames using the domain corp.altostrat.com while still

resolving Google Cloud hostnames. You want to follow Google-recommended practices. What should you do?

- A. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24. Set a custom route advertisement on the Cloud Router for 10.204.0.0/24
- B. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 35.199.192.0/19. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.
- C. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24. Modify the /etc/resolv.conf file on your Compute Engine instances to point to 192.168.20.88
- D. Create a private zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com. Configure DNS Server Policies and create a policy with Alternate DNS servers to 192.168.20.88. Configure your on-premises firewall to accept traffic from 35.199.192.0/19. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Answer: D

NEW QUESTION 42

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 Dev
- D. Create a firewall rule to deny traffic to port 8080 in the Shared VPC for Prod
- 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 47

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 51

You are configuring a new HTTP application that will be exposed externally behind both IPv4 and IPv6 virtual IP addresses, using ports 80, 8080, and 443. You will have backends in two regions: us-west1 and

us-east1. You want to serve the content with the lowest-possible latency while ensuring high availability and autoscaling, and create native content-based rules using the HTTP hostname and request path. The IP addresses of the clients that connect to the load balancer need to be visible to the backends. Which configuration should you use?

- A. Use Network Load Balancing
- B. Use TCP Proxy Load Balancing with PROXY protocol enabled
- C. Use External HTTP(S) Load Balancing with URL Maps and custom headers
- D. Use External HTTP(S) Load Balancing with URL Maps and an X-Forwarded-For header

Answer: D

NEW QUESTION 54

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 59

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 64

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 65

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 68

After a network change window one of your company's applications stops working. The application uses an on-premises database server that no longer receives any traffic from the application. The database server IP address is 10.2.1.25. You examine the change request, and the only change is that 3 additional VPC subnets were created. The new VPC subnets created are 10.1.0.0/16, 10.2.0.0/16, and 10.3.1.0/24/ The on-premises router is advertising 10.0.0.0/8. What is the most likely cause of this problem?

- A. The less specific VPC subnet route is taking priority.
- B. The more specific VPC subnet route is taking priority.
- C. The on-premises router is not advertising a route for the database server.
- D. A cloud firewall rule that blocks traffic to the on-premises database server was created during the change.

Answer: B

NEW QUESTION 70

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 74

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 79

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 81

You are designing a shared VPC architecture. Your network and security team has strict controls over which routes are exposed between departments. Your Production and Staging departments can communicate with each other, but only via specific networks. You want to follow Google-recommended practices. How should you design this topology?

- A. Create 2 shared VPCs within the shared VPC Host Project, and enable VPC peering between the
- B. Use firewall rules to filter access between the specific networks.
- C. Create 2 shared VPCs within the shared VPC Host Project, and create a Cloud VPN/Cloud Router between the
- D. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.
- E. Create 2 shared VPCs within the shared VPC Service Project, and create a Cloud VPN/Cloud Router between the
- F. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.
- G. Create 1 VPC within the shared VPC Host Project, and share individual subnets with the Service Projects to filter access between the specific networks.

Answer: D

NEW QUESTION 84

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

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

Answer: D

Explanation:

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

NEW QUESTION 87

You want to configure a NAT to perform address translation between your on-premises network blocks and GCP. Which NAT solution should you use?

- A. Cloud NAT
- B. An instance with IP forwarding enabled
- C. An instance configured with iptables DNAT rules
- D. An instance configured with iptables SNAT rules

Answer: A

NEW QUESTION 90

You want to create a service in GCP using IPv6. What should you do?

- A. Create the instance with the designated IPv6 address.
- B. Configure a TCP Proxy with the designated IPv6 address.
- C. Configure a global load balancer with the designated IPv6 address.
- D. Configure an internal load balancer with the designated IPv6 address.

Answer: C

Explanation:

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

NEW QUESTION 92

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 94

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 97

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 101

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

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

Answer: B

NEW QUESTION 105

You successfully provisioned a single Dedicated Interconnect. The physical connection is at a colocation facility closest to us-west2. Seventy-five percent of your

workloads are in us-east4, and the remaining twenty-five percent of your workloads are in us-central1. All workloads have the same network traffic profile. You need to minimize data transfer costs when deploying VLAN attachments. What should you do?

- A. Keep the existing Dedicated interconnect
- B. Deploy a VLAN attachment to a Cloud Router in us-west2, and use VPC global routing to access workloads in us-east4 and us-central1.
- C. Keep the existing Dedicated Interconnect
- D. Deploy a VLAN attachment to a Cloud Router in us-east4, and deploy another VLAN attachment to a Cloud Router in us-central1.
- E. Order a new Dedicated Interconnect for a colocation facility closest to us-east4, and use VPC global routing 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 108

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

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

Answer: C

NEW QUESTION 111

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 116

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 121

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 125

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

Which next hop should you choose?

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

Answer: C

Explanation:

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

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

NEW QUESTION 130

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 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

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

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

Answer: B

NEW QUESTION 144

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 146

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 151

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 154

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 158

You work for a university that is migrating to GCP. These are the cloud requirements:

- On-premises connectivity with 10 Gbps
- Lowest latency access to the cloud
- Centralized Networking Administration Team

New departments are asking for on-premises connectivity to their projects. You want to deploy the most cost-efficient interconnect solution for connecting the campus to Google Cloud.

What should you do?

- A. Use Shared VPC, and deploy the VLAN attachments and Interconnect in the host project.
- B. Use Shared VPC, and deploy the VLAN attachments in the service project
- C. Connect the VLAN attachment to the Shared VPC's host project.
- D. Use standalone projects, and deploy the VLAN attachments in the individual project
- E. Connect the VLAN attachment to the standalone projects' Interconnects.
- F. Use standalone projects and deploy the VLAN attachments and Interconnects in each of the individual projects.

Answer: A

Explanation:

<https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnects-other-projects>

Using Cloud Interconnect with Shared VPC You can use Shared VPC to share your VLAN attachment in a project with other VPC networks. Choosing Shared VPC

is preferable if you need to create many projects and would like to prevent individual project owners from managing their connectivity back to your on-premises network. In this scenario, the host project contains a common Shared VPC network usable by VMs in service projects. Because VMs in the service projects use this network, Service Project Admins don't need to create other VLAN attachments or Cloud Routers in the service projects. In this scenario, you must create VLAN attachments and Cloud Routers for a Cloud Interconnect connection only in the Shared VPC host project. The combination of a VLAN attachment and its associated Cloud Router are unique to a given Shared VPC network.

<https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiple-networks-access-sa>

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

NEW QUESTION 163

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 167

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

What should you do?

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

Answer: D

Explanation:

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

NEW QUESTION 168

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 169

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.co
- 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@democloud-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.co
- 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 172

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 174

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 178

Your organization has a single project that contains multiple Virtual Private Clouds (VPCs). You need to secure API access to your Cloud Storage buckets and BigQuery datasets by allowing API access only from resources in your corporate public networks. What should you do?

- A. Create an access context policy that allows your VPC and corporate public network IP ranges, and then attach the policy to Cloud Storage and BigQuery.
- B. Create a VPC Service Controls perimeter for your project with an access context policy that allows your corporate public network IP ranges.
- C. Create a firewall rule to block API access to Cloud Storage and BigQuery from unauthorized networks.
- D. Create a VPC Service Controls perimeter for each VPC with an access context policy that allows your corporate public network IP ranges.

Answer: B

NEW QUESTION 179

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 184

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 187

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

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

Answer: A

Explanation:

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

NEW QUESTION 189

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 192

You have provisioned a Dedicated Interconnect connection of 20 Gbps with a VLAN attachment of 10 Gbps. You recently noticed a steady increase in ingress traffic on the Interconnect connection from the on-premises data center. You need to ensure that your end users can achieve the full 20 Gbps throughput as quickly as possible. Which two methods can you use to accomplish this? (Choose two.)

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

Answer: CE

NEW QUESTION 194

.....

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/>