ATMIC NETWORKS
YOUR GUIDE TO GLOBAL CERTIFIATION

"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks .com

Warning: Keep connected with our support team for latest updates

Question: 1

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

Question: 2

Your end users are located in close proximity to us-east1 and europe-west1. Their workloads need to communicate with each other. You want to minimize cost and increase network efficiency.

How should you design this topology?

A. Create 2 VPCs, each with their own regions and individual subnets. Create 2 VPN gateways to establish connectivity between these regions.

B. Create 2 VPCs, each with their own region and individual subnets. Use external IP addresses on the instances to establish connectivity between these regions.

C. Create 1 VPC with 2 regional subnets. Create a global load balancer to establish connectivity between the regions.

D. Create 1 VPC with 2 regional subnets. Deploy workloads in these subnets and have them communicate using private RFC1918 IP addresses.

Answer: D

Explanation:

https://cloud.google.com/vpc/docs/using-vpc#create-auto-network

We create one VPC network in auto mode that creates one subnet in each Google Cloud region automatically. So, region us-east1 and europe-west1 are in the same network and they can communicate using their internal IP address even though they are in different Regions. They take advantage of Google's global fiber network.

Question: 3

Your organization is deploying a single project for 3 separate departments. Two of these departments require network connectivity between each other, but the third department should remain in isolation. Your design should create separate network administrative domains between these departments. You want to minimize operational overhead.

How should you design the topology?

A. Create a Shared VPC Host Project and the respective Service Projects for each of the 3 separate departments.

B. Create 3 separate VPCs, and use Cloud VPN to establish connectivity between the two appropriate VPCs.

C. Create 3 separate VPCs, and use VPC peering to establish connectivity between the two appropriate VPCs.

D. Create a single project, and deploy specific firewall rules. Use network tags to isolate access between the departments.

Answer: C

Explanation:

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

Question: 4

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

Question: 5

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

How should you configure the Distribution VPC?

A. Create the Distribution VPC in auto mode. Peer both the VPCs via network peering.

B. Create the Distribution VPC in custom mode. Use the CIDR range 10.0.0.0/9. Create the necessary subnets, and then peer them via network peering.

C. Create the Distribution VPC in custom mode. Use the CIDR range 10.128.0.0/9. Create the necessary subnets,

and then peer them via network peering.

D. Rename the default VPC as "Distribution" and peer it via network peering.

Answer: B

Explanation:

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

Question: 6

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.

Question: 7

All the instances in your project are configured with the custom metadata enable-oslogin value set to FALSE and to block project-wide SSH keys. None of the instances are set with any SSH key, and no project-wide SSH keys have been configured. Firewall rules are set up to allow SSH sessions from any IP address range. You want to SSH into one instance.

What should you do?

A. Open the Cloud Shell SSH into the instance using gcloud compute ssh.

B. Set the custom metadata enable-oslogin to TRUE, and SSH into the instance using a third-party tool like putty or ssh.

C. Generate a new SSH key pair. Verify the format of the private key and add it to the instance. SSH into the instance using a third-party tool like putty or ssh.

D. Generate a new SSH key pair. Verify the format of the public key and add it to the project. SSH into the instance using a third-party tool like putty or ssh.

Answer: A

Explanation:

Question: 8

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 projects. Connect the VLAN attachment to the Shared VPC's host project.

C. Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects' Interconnects.

D. 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-same-attachment#using_with

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

Question: 9

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

Explanation:

Question: 10

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

Question: 11

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.

Question: 12

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.

Question: 13

You want to deploy a VPN Gateway to connect your on-premises network to GCP. You are using a non BGP-capable on-premises VPN device. You want to minimize downtime and operational overhead when your network grows. The device supports only IKEv2, and you want to follow Google- recommended practices.

What should you do?

A. • Create a Cloud VPN instance.• Create a policy-based VPN tunnel per subnet.• Configure the appropriate local and remote traffic selectors to match your local and remote networks.• Create the appropriate static routes.

B. • Create a Cloud VPN instance.• Create a policy-based VPN tunnel.• Configure the appropriate local and remote traffic selectors to match your local and remote networks.• Configure the appropriate static routes.

C. • Create a Cloud VPN instance.• Create a route-based VPN tunnel.• Configure the appropriate local and remote traffic selectors to match your local and remote networks.• Configure the appropriate static routes.

D. • Create a Cloud VPN instance.• Create a route-based VPN tunnel.• Configure the appropriate local and remote traffic selectors to 0.0.0.0/0.• Configure the appropriate static routes.

Answer: B

Explanation:

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

Question: 14

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

Question: 15

Your on-premises data center has 2 routers connected to your Google Cloud environment 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 a unique 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.
- BGP sessions are established between both on-premises routers and the Cloud Router.
- Only 1 of the on-premises router's routes are being added to the routing table.

What is the most likely cause of this problem?

A. The on-premises routers are configured with the same routes.

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. The ASNs being used on the on-premises routers are different.

Answer: D

Explanation:

https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp

Question: 16

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

Question: 17

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

Question: 18

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

How should you deploy this service in GCP?

A. Create a managed instance group from one of the images of the on-premises servers, and link this instance group to a target pool behind your load balancer.

B. Create a target pool, add all backend instances to this target pool, and deploy the target pool behind your load balancer.

C. Deploy a third-party virtual appliance as frontend to these servers that will accommodate the significant differences between these backend servers.

D. Use GCP's ECMP capability to load-balance traffic to the backend servers by installing multiple equal-priority static routes to the backend servers.

Answer: B

Explanation:

Question: 19

You decide to set up Cloud NAT. After completing the configuration, you find that one of your instances is not using the Cloud NAT for outbound NAT.

What is the most likely cause of this problem?

A. The instance has been configured with multiple interfaces.

B. An external IP address has been configured on the instance.

C. You have created static routes that use RFC1918 ranges.

D. The instance is accessible by a load balancer external IP address.

Answer: B

Explanation:

Question: 20

You want to set up two Cloud Routers so that one has an active Border Gateway Protocol (BGP) session, and the other one acts as a standby.

Which BGP attribute should you use on your on-premises router?

A. AS-Path

B. Community

C. Local Preference

D. Multi-exit Discriminator

Answer: D

Explanation:

Question: 21

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. Create a second tunnel on the existing Cloud VPN gateway that forwards the same IP range, but points at the new onpremises gateway IP.

D. Add a second Cloud VPN gateway in a different region than the existing VPN gateway. 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

Question: 22

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 registar.

D. Transfer ownership of the domain to a new registar.

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.

Question: 23

You have an application hosted on a Compute Engine virtual machine instance that cannot communicate with a resource outside of its subnet. When you review the flow and firewall logs, you do not see any denied traffic listed.

During troubleshooting you find:

- Flow logs are enabled for the VPC subnet, and all firewall rules are set to log.
- The subnetwork logs are not excluded from Stackdriver.
- The instance that is hosting the application can communicate outside the subnet.
- Other instances within the subnet can communicate outside the subnet.
- The external resource initiates communication.

What is the most likely cause of the missing log lines?

A. The traffic is matching the expected ingress rule.

B. The traffic is matching the expected egress rule.

C. The traffic is not matching the expected ingress rule.

D. The traffic is not matching the expected egress rule.

Answer: C

Explanation:

Question: 24

You have configured Cloud CDN using HTTP(S) load balancing as the origin for cacheable content. Compression is configured on the web servers, but responses served by Cloud CDN are not compressed.

What is the most likely cause of the problem?

A. You have not configured compression in Cloud CDN.

B. You have configured the web servers and Cloud CDN with different compression types.

C. The web servers behind the load balancer are configured with different compression types.

D. You have to configure the web servers to compress responses even if the request has a Via header.

Answer: D

Explanation:

If responses served by Cloud CDN are not compressed but should be, check that the web server software running on your instances is configured to compress responses. By default, some web server software will automatically disable compression for requests that include a Via header. The presence of a Via header indicates the request was forwarded by a proxy. HTTP proxies such as HTTP(S) load balancing add a Via header to each request as required by the HTTP specification. To enable compression, you may have to override your web server's default configuration to tell it to compress responses even if the request had a Via header.

Question: 25

You have a web application that is currently hosted in the us-central1 region. Users experience high latency when traveling in Asia. You've configured a network load balancer, but users have not experienced a performance improvement. You want to decrease the latency.

What should you do?

A. Configure a policy-based route rule to prioritize the traffic.

B. Configure an HTTP load balancer, and direct the traffic to it.

C. Configure Dynamic Routing for the subnet hosting the application.

D. Configure the TTL for the DNS zone to decrease the time between updates.

Answer: B

Explanation:

Question: 26

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: A,E

Explanation:

https://cloud.google.com/nat/docs/overview#interaction-pga Specifications

https://cloud.google.com/vpc/docs/configure-private-google-access#specifications

Question: 27

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 them. Use firewall rules to filter access between the specific networks.

B. Create 2 shared VPCs within the shared VPC Host Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.

C. Create 2 shared VPCs within the shared VPC Service Project, and create a Cloud VPN/Cloud Router between them. Use Flexible Route Advertisement (FRA) to filter access between the specific networks.

D. 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

Explanation:

Question: 28

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

Explanation:

Question: 29

You converted an auto mode VPC network to custom mode. Since the conversion, some of your Cloud Deployment Manager templates are no longer working. You want to resolve the problem.

What should you do?

A. Apply an additional IAM role to the Google API's service account to allow custom mode networks.

B. Update the VPC firewall to allow the Cloud Deployment Manager to access the custom mode networks.

C. Explicitly reference the custom mode networks in the Cloud Armor whitelist.

D. Explicitly reference the custom mode networks in the Deployment Manager templates.

Answer: D

Explanation:

Question: 30

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: D,E

Explanation:

Question: 31

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

Question: 32

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

Question: 33

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 Services. Create a VPC-native cluster and specify those ranges.

B. Create a subnet of size/28 with 2 secondary ranges of: /24 for Pods and /24 for Services. Create a VPC-native cluster and specify those ranges. When the services are ready to be deployed, resize the subnets.

C. Use gcloud container clusters create [CLUSTER NAME]--enable-ip-alias to create a VPC-native cluster.

D. Use gcloud container clusters create [CLUSTER NAME] to create a VPC-native cluster.

Answer: A

Explanation:

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

https://docs.netgate.com/pfsense/en/latest/book/network/understanding-cidr-subnet-mask- notation.html

Question: 34

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

Question: 35

Your company is working with a partner to provide a solution for a customer. Both your company and the partner organization are using GCP. There are applications in the partner's network that need access to some resources in your company's VPC. There is no CIDR overlap between the VPCs.

Which two solutions can you implement to achieve the desired results without compromising the security? (Choose two.)

A. VPC peering

B. Shared VPC

C. Cloud VPN

D. Dedicated Interconnect

E. Cloud NAT

Answer: AC

Explanation:

Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.

Question: 36

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. Wait 90 seconds. Re-enable Cloud CDN on the storage bucket.

Answer: B

Explanation:

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

Question: 37

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

Question: 38

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.

Question: 39

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

Question: 40

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/orelly/networking_2ndEd/fire/ch17_02.htm Besides, Google Cloud external TCP/UDP Network Load Balancing (after this referred to as Network Load Balancing) is a regional, non-proxied load balancer. Network Load Balancing distributes traffic among virtual machine (VM) instances in the same region in a Virtual Private Cloud (VPC) netw

Question: 41

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

Explanation:

Question: 42

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

Explanation:

Question: 43

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

Question: 44

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

What should you do?

A. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with no tag.

B. Create a more specific route than the system-generated subnet route, pointing the next hop to instance-B with a tag applied to instance-A.

C. Delete the system-generated subnet route and create a specific route to instance-B with a tag applied to instance-A.

D. Move instance-B to another VPC and, using multi-NIC, connect instance-B's interface to instance- A's network. Configure the appropriate routes to force traffic through to instance-A.

Answer: B

Explanation:

Question: 45

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

Question: 46

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

Explanation:

Question: 47

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

Question: 48

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

Question: 49

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. Verify the new feature in the new canary instance group, and then update the original instance group.

D. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template. 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#starting_a_canary_update

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

Question: 50

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

How should you provision your instances?

A. Create a single managed instance group, specify the desired region, and select Multiple zones for the location.

B. Create a managed instance group for each region, select Single zone for the location, and manually distribute instances across the zones in that region.

C. Create an unmanaged instance group in a single zone, and then create an HTTP load balancer for the instance group.

D. Create an unmanaged instance group for each zone, and manually distribute the instances across the desired zones.

Answer: A

Explanation:

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

Question: 51

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. Then edit the bucket setting and enable the private attribute.

C. Add an appropriate lifecycle rule on the storage bucket containing the two objects.

D. Add a Cache-Control entry with value private to the metadata of the object you don't want to be cached anymore. Invalidate all the previously cached copies.

Answer: D

Explanation:

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

Question: 52

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 trafficscrubbing 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

Question: 53

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

Your ISP is a Google Partner Interconnect provider.

Your on-premises VPN device's internet uplink and downlink speeds are 10 Gbps.

A test VPN connection between your on-premises gateway and GCP is performing at a maximum speed of 500 Mbps due to packet losses.

Most of the data transfer will be from GCP to the on-premises environment.

The application can burst up to 1.5 Gbps during peak transfers over the Interconnect.

Cost and the complexity of the solution should be minimal.

How should you provision the connectivity solution?

A. Provision a Partner Interconnect through your ISP.

B. Provision a Dedicated Interconnect instead of a VPN.

C. Create multiple VPN tunnels to account for the packet losses, and increase bandwidth using ECMP.

D. Use network compression over your VPN to increase the amount of data you can send over your VPN.

Answer: A

Explanation:

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

Question: 54

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 hours. The attack will stop when the application is ofline.

E. SSH into the backend compute engine instances, and view the auth logs and syslogs to further understand the nature of the attack.

Answer: BE

Explanation:

Question: 55

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

Question: 56

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

Question: 57

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.

Reference: https://cloud.google.com/vpc/docs/special-configurations

Question: 58

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

Reference: https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

Question: 59

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_with_cloud_storage_buckets

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

Question: 60

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

/fr/video

/en/video

/es/video

/../video

/fr/audio

/en/audio /es/audio

/../audio

Which solution should you recommend?

A. Rearrange the directory structure, create a URL map and leverage a path rule such as /video/* and /audio/*.

B. Rearrange the directory structure, create DNS hostname entries for video and audio and leverage a path rule such as /video/* and /audio/*.

C. Leave the directory structure as-is, create a URL map and leverage a path rule such as \/[a- z]{2}\/video and \/[a-z]{2}\/audio.

D. Leave the directory structure as-is, create a URL map and leverage a path rule such as /*/video and /*/ audio.

Answer: A

Explanation:

https://cloud.google.com/load-balancing/docs/url-map#configuring_url_maps

Path matcher constraints Path matchers and path rules have the following constraints: A path rule can only include a wildcard character (*) after a forward slash character (/). For example, /videos/* and /videos/hd/* are valid for path rules, but /videos* and /videos/hd* are not. Path rules do not use regular expression or substring matching. For example, path rules for either /videos/hd or /videos/hd/* do not apply to a URL with the path /video/hd-abcd. However, a path rule for /video/* does apply to that path. https://cloud.google.com/load-balancing/docs/url-map-concepts#pm- constraints

Question: 61

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.

Reference: https://cloud.google.com/interconnect/docs/how-to/direct-peering

Question: 62

You are configuring a new instance of Cloud Router in your Organization's Google Cloud environment to allow connection across a new Dedicated Interconnect to your data center Sales, Marketing, and IT each have a service project attached to the Organization's host project.

Where should you create the Cloud Router instance?

A. VPC network in all projects

B. VPC network in the IT Project

C. VPC network in the Host Project

D. VPC network in the Sales, Marketing, and IT Projects

Answer: C

Explanation:

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

Question: 63

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

Explanation:

Reference: https://geekflare.com/gcp-firewall-configuration/

Question: 64

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%20logs%20show,of%20ciphers%20must%20match%20exactly.&text=Make%20sure%20that%20you%20use,configure%20your%20peer%20VPN%20gateway.

Question: 65

You need to define an address plan for a future new GKE cluster in your 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: B

Explanation:

https://cloud.google.com/kubernetes-engine/docs/how-to/alias- ips#cluster_sizing_secondary_range_pods

Reference: https://cloud.google.com/kubernetes-engine/docs/how-to/alias-ips

https://cloud.google.com/kubernetes-engine/docs/how-to/flexible-pod-cidr

https://cloud.google.com/kubernetes-engine/docs/concepts/alias-ips#defaults_limits

Question: 66

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

Question: 67

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 (e. g., code-dev), and make the second project (e. g., data-dev) a service project.

D. 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.

E. 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

Explanation:

Question: 68

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

IP ranges for pods and services must be as small as possible.

The nodes and the master must not be reachable from the internet.

You must be able to use kubectl commands from on-premises subnets to manage the cluster.

How should you create the GKE cluster?

A. • Create a private cluster that uses VPC advanced routes.
- Set the pod and service ranges as /24.
- Set up a network proxy to access the master.

B. • Create a VPC-native GKE cluster using GKE-managed IP ranges.
- Set the pod IP range as /21 and service IP range as /24.
- Set up a network proxy to access the master.

C. • Create a VPC-native GKE cluster using user-managed IP ranges.
- Enable a GKE cluster network policy, set the pod and service ranges as /24.
- Set up a network proxy to access the master.
- Enable master authorized networks.

D. • Create a VPC-native GKE cluster using user-managed IP ranges.
- Enable privateEndpoint on the cluster master.
- Set the pod and service ranges as /24.
- Set up a network proxy to access the master.
- Enable master authorized networks.

Answer: D

Explanation:

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

Question: 69

You are creating an instance group and need to create a new health check for HTTP(s) load balancing.

Which two methods can you use to accomplish this? (Choose two.)

A. Create a new health check using the gcloud command line tool.

B. Create a new health check using the VPC Network section in the GCP Console.

C. Create a new health check, or select an existing one, when you complete the load balancer's backend configuration in the GCP Console.

D. Create a new legacy health check using the gcloud command line tool.

E. Create a new legacy health check using the Health checks section in the GCP Console.

Answer: AC

Explanation:

https://cloud.google.com/load-balancing/docs/health- checks#creating and modifying health checks

Question: 70

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

Explanation:

Question: 71

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

Explanation:

Reference: https://cloud.google.com/vpn/docs/concepts/overview

Question: 72

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

Explanation:

Question: 73

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#provisioning "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.

Question: 74

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

Explanation:

Question: 75

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

Explanation:

Reference: https://cloud.google.com/iam/docs/understanding-custom-roles

Question: 76

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".

Question: 77

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

Explanation:

Question: 78

You need to create a new VPC network that allows instances to have IP addresses in both the 10.1.1.0/24 network and the 172.16.45.0/24 network.

What should you do?

A. Configure global load balancing to point 172.16.45.0/24 to the correct instance.

B. Create unique DNS records for each service that sends traffic to the desired IP address.

C. Configure an alias-IP range of 172.16.45.0/24 on the virtual instances within the VPC subnet of 10.1.1.0/24.

D. Use VPC peering to allow traffic to route between the 10.1.0.0/24 network and the 172.16.45.0/24 network.

Answer: C

Explanation:

Question: 79

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-balancers-3697d76feeb1

Reference: https://cloud.google.com/load-balancing/docs/network

Question: 80

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

Explanation:

Question: 81

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

A. Configure Firewall Rules Logging. Use Firewall Insights to display the number of hits.

B. Configure Firewall Rules Logging. View the logs in Cloud Logging, and create a custom dashboard in Cloud Monitoring to display the number of hits.

C. Configure a firewall appliance from the Google Cloud Marketplace. Route all traffic through this appliance, and apply the firewall rules at this layer. Use the firewall appliance to display the number of hits.

D. Configure Packet Mirroring on the VPC. Apply a filter with an IP address list of the Denied Firewall rules. Configure an intrusion detection system (IDS) appliance as the receiver to display the number of hits.

Answer: A

Explanation:

Question: 82

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 advertisements. Leave all other options as their default settings.

D. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of

advertisements. Advertise all visible subnets to the Cloud Router.

Answer: C

Explanation:

Question: 83

You are configuring load balancing for a standard three-tier (web, application, and database) application. You have configured an external HTTP(S) load balancer for the web servers. You need to configure load balancing for the application tier of servers. What should you do?

A. Configure a forwarding rule on the existing load balancer for the application tier.

B. Configure equal cost multi-path routing on the application servers.

C. Configure a new internal HTTP(S) load balancer for the application tier.

D. Configure a URL map on the existing load balancer to route traffic to the application tier.

Answer: A

Explanation:

Question: 84

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 Logs. Create a sink in Cloud Logging to send filtered egress VPC Flow Logs to the IDS.

C. Create an internal TCP/UDP load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

D. Create an internal HTTP(S) load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

Answer: B

Explanation:

Question: 85

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. Clients should use this IP address to connect to the service.

B. 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/.

C. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS. Clients should use the name of the A record to connect to the service.

D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url https://[API_NAME]/[API_VERSION]/.

Answer: C

Explanation:

Question: 86

You recently deployed Cloud VPN to connect your on-premises data canter 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 Canter, 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

Explanation:

Question: 87

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 VPC. Connect your VPN gateways to the partner's gateways. Enable global dynamic routing in each VPC.

B. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC. Create one OpenVPN Access Server in each region of your partner's VPC. Connect your VPN gateway to your partner's servers.

C. Create one OpenVPN Access Server in each region of your VPC and your partner's VPC. Connect your servers to the partner's servers.

D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways with a pair of tunnels. Enable global dynamic routing in each VPC.

Answer: A

Explanation:

Question: 88

You need to create the network infrastructure to deploy a highly available web application in the us- east1 and us-west1 regions. The application runs on Compute Engine instances, and it does not require the use of a database. You want to follow Google-recommended practices. What should you do?

A. Create one VPC with one subnet in each region.

Create a regional network load balancer in each region with a static IP address.

Enable Cloud CDN on the load balancers.

Create an A record in Cloud DNS with both IP addresses for the load balancers.

B. Create one VPC with one subnet in each region.

Create a global load balancer with a static IP address.

Enable Cloud CDN and Google Cloud Armor on the load balancer.

Create an A record using the IP address of the load balancer in Cloud DNS.

C. Create one VPC in each region, and peer both VPCs.

Create a global load balancer.

Enable Cloud CDN on the load balancer.

Create a CNAME for the load balancer in Cloud DNS.

D. Create one VPC with one subnet in each region.

Create an HTTP(S) load balancer with a static IP address.

Choose the standard tier for the network.

Enable Cloud CDN on the load balancer.

Create a CNAME record using the load balancer's IP address in Cloud DNS.

Answer: C

Explanation:

Question: 89

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.

Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.

D. 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

Explanation:

Question: 90

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 VPCs. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network. Use Cloud Router's custom route advertisements to announce a default route to the on-premises locations.

D. Peer the two VPCs. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network. Use Cloud Router's custom route advertisements to announce the peered VPC network ranges to the on-premises locations.

Answer: A

Explanation:

Question: 91

You recently noticed a recurring daily spike in network usage in your Google Cloud project. You need to identify the virtual machine (VM) instances and type of traffic causing the spike in traffic utilization while minimizing the cost and management overhead required. What should you do?

A. Enable VPC Flow Logs and send the output to BigQuery for analysis.

B. Enable Firewall Rules Logging for all allowed traffic and send the output to BigQuery for analysis.

C. Configure Packet Mirroring to send all traffic to a VM. Use Wireshark on the VM to identity traffic utilization for each VM in the VPC.

D. Deploy a third-party network appliance and configure it as the default gateway. Use the third- party network appliance to identify users with high network traffic.

Answer: C

Explanation:

Question: 92

You need to enable Private Google Access for use by some subnets within your Virtual Private Cloud (VPC). Your security team set up the VPC to send all internet-bound traffic back to the on- premises

data center for inspection before egressing to the internet, and is also implementing VPC Service Controls in the environment for API-level security control. You have already enabled the subnets for Private Google Access. What

configuration changes should you make to enable Private Google Access while adhering to your security team's requirements?

A. Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range.

Create a custom route that points Google's restricted API address range to the default internet gateway as the next hop.

B. Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range.

Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.

C. Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record painting to Google's private AP address range.

Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.

D. Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record pointing to Google's private API address range.

Create a custom route that points Google's private API address range to the default internet gateway as the next hop.

Answer: C

Explanation:

Question: 93

You have deployed an HTTP(s) load balancer, but health checks to port 80 on the Compute Engine virtual machine instance are failing, and no traffic is sent to your instances. You want to resolve the problem. Which commands should you run?

A. gcloud compute instances add-access-config instance-1

B. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destinationranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS

C. gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS

D. gcloud compute health-checks update http health-check --unhealthy-threshold 10

Answer: A

Explanation:

Question: 94

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

Explanation:

Question: 95

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 onpremises 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

Explanation:

Question: 96

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 projects. Create a VPC firewall rule to open port 8080 in the Shared VPC for Dev. Create a firewall rule to deny traffic to port 8080 in the Shared VPC for Prod. Deploy VMs to those Shared VPCs.

C. 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.

D. 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

Explanation:

Question: 97

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.

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	rf-tunnel-a-to-b-rf-0	169 254 0 254	169.254 0.254	65502
router-2	if-lunnel-b-to-a-rf-0	169 254 0 254	169.254.0.254	65501

B.

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	rf-lunnel-a-to-b-rf-0	10 1 0 1	172 16.0.1	15052
router-2	rf-lunnel-b-to-a-rf-0	172 16 0 1	10101	15501

C.

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	rf-lunnel-a-lo-b-tf-0	169 254 20 1	169 254 20.2	65002
router-2	rf-lunnel-b-to-a-rf-0	169 254 20 2	169 254 20 1	65001

D.

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	if-tunnel-a-to-b-rf-0	172 16 0 254	10 10.254	16552
router-2	if-lunnel-b-to-a-rf-0	10 1 0 254	172 16 0 254	16551

A. Option A

B. Option B

C. Option C

D. Option D

Answer: C

Explanation:

Question: 98

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

Explanation:

Question: 99

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

Explanation:

Question: 100

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. Create service projects for Web, App, and Database.

B. Configure one VPC for Web, one VPC for App, and one VPC for Database. Configure HA VPN between each VPC.

C. Configure three Shared VPC host projects, each with a service project: one for Web, one for App, and one for Database.

D. Configure one VPC for Web, one VPC for App, and one VPC for Database. Use VPC Network Peering to connect all VPCs in a full mesh.

Answer: C

Explanation:

Question: 101

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. Create the relevant routes on the third-party appliances and VPC networks.

B. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network. Create separate VPC networks for on- premises and internet connectivity. Create the relevant routes on the third-party appliances and VPC networks.

C. Consolidate all existing projects' subnetworks into a single VPC. Create separate VPC networks for on-premises and internet connectivity. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network. Create the relevant routes on the third- party appliances and VPC networks.

D. Configure the third-party appliances with multiple interfaces. Create a hub VPC network for all projects, and create separate VPC networks for on-premises and internet connectivity. Create the relevant routes on the third-party appliances and VPC networks. Use VPC Network Peering to connect all projects' VPC networks to the hub VPC. Export custom routes from the hub VPC and import on all projects' VPC networks.

Answer: D

Explanation:

Question: 102

You have just deployed your infrastructure on Google Cloud. You now need to configure the DNS to meet the following requirements:

Your on-premises resources should resolve your Google Cloud zones.

Your Google Cloud resources should resolve your on-premises zones.

You need the ability to resolve ". internal" zones provisioned by Google Cloud.

What should you do?

A. Configure an outbound server policy, and set your alternative name server to be your on-premises DNS resolver. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

B. Configure both an inbound server policy and outbound DNS forwarding zones with the target as the on-premises DNS resolver. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.

C. Configure an outbound DNS server policy, and set your alternative name server to be your onpremises DNS resolver. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.

D. Configure Cloud DNS to DNS peer with your on-premises DNS resolver. Configure your onpremises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

Answer: A

Explanation:

Question: 103

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

Explanation:

Question: 104

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

A. Configure Private Google Access to privately access the Cloud Storage service using private IP addresses.

B. Configure a VPC Service Controls perimeter around project XYZ, and include storage.googleapis.com as a restricted service in the service perimeter.

C. Configure Cloud Storage with projectPrivate Access Control List (ACL) that gives permission to the project team based on their roles.

D. Configure Private Service Connect to privately access Cloud Storage from all VPCs under project XYZ.

Answer: C

Explanation:

Question: 105

You are maintaining a Shared VPC in a host project. Several departments within your company have infrastructure in different service projects attached to the Shared VPC and use Identity and Access Management (IAM) permissions to manage the cloud resources in those projects. VPC Network Peering is also set up between the Shared VPC and a common services VPC that is not in a service project. Several users are experiencing failed connectivity between certain instances in different Shared VPC service projects and between certain instances and the internet. You need to validate the network configuration to identify whether a misconfiguration is the root cause of the problem. What should you do?

A. Review the VPC audit logs in Cloud Logging for the affected instances.

B. Use Secure Shell (SSH) to connect to the affected Compute Engine instances, and run a series of PING tests to the other affected endpoints and the 8.8.8.8 IPv4 address.

C. Run Connectivity Tests from Network Intelligence Center to check connectivity between the affected endpoints in your network and the internet.

D. Enable VPC Flow Logs for all VPCs, and review the logs in Cloud Logging for the affected instances.

Answer: C

Explanation:

Question: 106

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

Explanation:

Question: 107

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

A. Enable firewall logs, and view the logs in Firewall Insights.

B. Enable HTTP(S) Load Balancing logging with sampling rate equal to 1, and view the logs in Cloud Logging.

C. Enable VPC Flow Logs, and view the logs in Cloud Logging.

D. Enable Google Cloud Armor audit logs, and view the logs on the Activity page in the Google Cloud Console.

Answer: A

Explanation:

Question: 108

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

A. Set up the engineer with Compute Shared VPC Admin IAM role at the folder level.

B. Set up the engineer with Compute Shared VPC Admin IAM role at the organization level.

C. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the folder level.

D. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the organization level.

Answer: B

Explanation:

Question: 109

You recently deployed Compute Engine instances in regions us-west1 and us-east1 in a Virtual Private Cloud (VPC) with default routing configurations. Your company security policy mandates that virtual machines (VMs) must not have public IP addresses attached to them. You need to allow your instances to fetch updates from the internet while preventing external access. What should you do?

A. Create a Cloud NAT gateway and Cloud Router in both us-west1 and us-east1.

B. Create a single global Cloud NAT gateway and global Cloud Router in the VPC.

C. Change the instances' network interface external IP address from None to Ephemeral.

D. Create a firewall rule that allows egress to destination 0.0.0.0/0.

Answer: A

Explanation:

Question: 110

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

A. Configure VPC Service Controls and create a secure perimeter. Define fine-grained perimeter controls and enforce that security posture across your Google Cloud services and projects.

B. Configure a Google Cloud Armor security policy in your project, and attach it to the backend service to secure the application.

C. Configure VPC firewall rules to protect the Compute Engine instances against distributed denial-of- service attacks.

D. Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at the organization level.

Answer: C

Explanation:

Question: 111

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

Explanation:

Question: 112

Your company has defined a resource hierarchy that includes a parent folder with subfolders for each department. Each department defines their respective project and VPC in the assigned folder and has the appropriate permissions to create Google Cloud firewall rules. The VPCs should not allow traffic to flow between them. You need to block all traffic from any source, including other VPCs, and delegate only the intra-VPC firewall rules to the respective departments. What should you do?

A. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 0.

B. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 1000.

C. Create two hierarchical firewall policies per department's folder with two rules in each: a high- priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to allow, and another lower-priority rule that blocks traffic from any other source.

D. Create two hierarchical firewall policies per department's folder with two rules in each: a high- priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the

action to goto_next, and another lower-priority rule that blocks traffic from any other source.

Answer: B

Explanation:

Question: 113

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

Explanation:

Question: 114

You are configuring an HA VPN connection between your Virtual Private Cloud (VPC) and onpremises 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

Explanation:

Question: 115

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

Explanation:

Question: 116

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

A. Configure a /28 primary IP address range for the node IP addresses. Configure a (25 secondary IP range for the Pods. Configure a /22 secondary IP range for the Services.

B. Configure a /28 primary IP address range for the node IP addresses. Configure a /25 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.

C. Configure a /28 primary IP address range for the node IP addresses. Configure a /28 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.

D. Configure a /28 primary IP address range for the node IP addresses. Configure a /24 secondary IP range for the Pads. Configure a /22 secondary IP range for the Services.

Answer: A

Explanation:

Question: 117

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

Explanation:

Question: 118

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

Explanation:

Question: 119

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

What should you do?

A. Use a 4-byte private ASN 4200000000-4294967294.

B. Use a 2-byte private ASN 64512-65535.

C. Use a public Google ASN 15169.

D. Use a public Google ASN 16550.

Answer: B

Explanation:

Question: 120

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

Explanation:

Question: 121

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. Add the tag to the application servers, and associate the service account with the database servers. 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

B. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com. Associate service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command:

gcloud compute firewall-rules create app-db-firewall-ru

- -allow TCP:3306 \
- -source-service-accounts sa-app@democloud-idp-demo.iam.gserviceaccount.com \
- -target-service-accounts sa-db@my-

project.iam.gserviceaccount.com

C. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com. Associate the service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command:

gcloud compute firewall-rules create app-db-firewall-ru

- -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

D. Create network tags app-server and db-server. Add the app-server tag to the application servers, and add the db-server tag to the database servers. 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

Explanation:

Question: 122

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 onpremises locations without address translation, but all RFC 1918 ranges are already in use in the onpremises 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

Explanation:

Question: 123

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

A. Create an allow on match ingress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.

B. Create an allow on match egress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.

C. Create an allow on match ingress firewall rule with the target tag “web-server” to allow all IP addresses for TCP ports 80 and 443.

D. Create an allow on match egress firewall rule with the target tag “web-server" to allow web server IP addresses for TCP ports 60 and 443.

Answer: C

Explanation:

Reference: https://cloud.google.com/load-balancing/docs/https

Question: 124

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. 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.

B. Keep the existing Dedicated Interconnect. Deploy a VLAN attachment to a Cloud Router in us- east4, and deploy another VLAN attachment to a Cloud Router in us-central1.

C. Order a new Dedicated Interconnect for a colocation facility closest to us-east4, and use VPC global routing to

access workloads in us-central1.

D. 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

Explanation:

Question: 125

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

A. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC.

Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target.

Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Configure VPC peering in the spoke VPCs to peer with the hub VPC.

B. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com

that points to 192.168.20.88.

Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke PCs, with the hub VPC as the target.

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. Associate the zone with the hub VPC.

Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target.

Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Create a hub-and-spoke VPN deployment in each spoke VPC to connect back to the on-premises network directly.

D. Create a private forwarding zone in Cloud DNS for 'corp altostrat.com' called corp-altostrat-com that points to

192. 168.20.88. Associate the zone with the hub VPC.

Create a private peering zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target.

Sat a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Create a hub and spoke VPN deployment in each spoke VPC to connect back to the hub VPC.

Answer: A

Explanation:

Question: 126

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

Direction	Action	Address range	Port	Priority
egress	deny	192 0 2 0/24	80	100
egress	deny	198 51 100 0 24	80	200
ingress	allow	203 0 113 0/24	80	300

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

Direction	Action	Address range	Port	Logging
egress	deny	192 0 2 42/32	80	true

You are using a new user account. You must assign the appropriate identity and Access Management (IAM) user roles to this new user account before updating the firewall rule. The new user account must be able to apply the update and view firewall logs. What should you do?

A. Assign the compute.securityAdmin and logging.viewer rule to the new user account. Apply the new firewall rule with a priority of 50.

B. Assign the compute.securityAdmin and logging.bucketWriter role to the new user account. Apply the new firewall rule with a priority of 150.

C. Assign the compute.orgSecurityPolicyAdmin and logging.viewer role to the new user account. Apply the new firewall rule with a priority of 50.

D. 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

Explanation:

Question: 127

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

Explanation:

Question: 128

Your company has provisioned 2000 virtual machines (VMs) in the private subnet of your Virtual Private Cloud (VPC) in the us-east1 region. You need to configure each VM to have a minimum of 128 TCP connections to a public repository so that users can download software updates and packages over the internet. You need to implement a Cloud NAT gateway so that the VMs are able to perform outbound NAT to the internet. You must ensure that all VMs can simultaneously connect to the public repository and download software updates and packages. Which two methods can you use to accomplish this? (Choose two.)

A. Configure the NAT gateway in manual allocation mode, allocate 2 NAT IP addresses, and update the minimum number of ports per VM to 256.

B. Create a second Cloud NAT gateway with the default minimum number of ports configured per VM to 64.

C. Use the default Cloud NAT gateway's NAT proxy to dynamically scale using a single NAT IP address.

D. Use the default Cloud NAT gateway to automatically scale to the required number of NAT IP addresses, and update the minimum number of ports per VM to 128.

E. Configure the NAT gateway in manual allocation mode, allocate 4 NAT IP addresses, and update the minimum number of ports per VM to 128.

Answer: A, B

Explanation:

Question: 129

You have the following routing design. You discover that Compute Engine instances in Subnet-2 in the asia-southeast1 region cannot communicate with compute resources on-premises. What should you do?

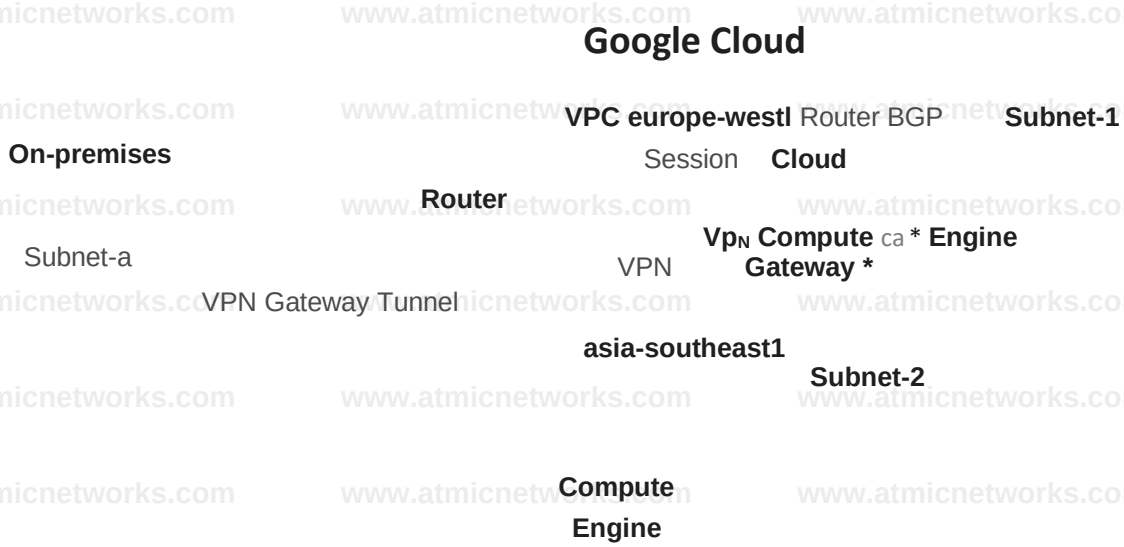

A. Configure a custom route advertisement on the Cloud Router.

B. Enable IP forwarding in the asia-southeast1 region.

C. Change the VPC dynamic routing mode to Global.

D. Add a second Border Gateway Protocol (BGP) session to the Cloud Router.

Answer: C

Explanation:

Question: 130

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

Explanation:

Question: 131

Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with onpremises connectivity already in place. You are deploying a new application using Google Kubernetes Engine (GKE), which must be accessible only from the same VPC network and on-premises locations. You must ensure that the GKE control plane is exposed to a predefined list of on-premises subnets through private connectivity only. What should you do?

A. Create a GKE private cluster with a private endpoint for the control plane. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Routers. Configure authorized networks to specify the desired on-premises subnets.

B. Create a GKE private cluster with a public endpoint for the control plane. Configure VPC Networking Peering export/import routes and custom route advertisements on the Cloud Routers.

C. Create a GKE private cluster with a private endpoint for the control plane. Configure authorized networks to specify the desired on-premises subnets.

D. Create a GKE public cluster. Configure authorized networks to specify the desired on-premises subnets.

Answer: C

Explanation:

Question: 132

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 zones. Create a global TCP load balancer. Add the Cloud Run endpoints to its backend service.

B. Deploy the Cloud Run services to multiple regions. Create serverless network endpoint groups (NEGs) that point to the services. Create a global HTTPS load balancer, and attach the serverless NEGs as backend services of the load balancer.

C. Deploy the Cloud Run services to multiple availability zones. Create Cloud Endpoints that point to the services. Create a global HTTPS load balancer, and attach the Cloud Endpoints to its backend

D. Deploy the Cloud Run services to multiple regions. Configure a round-robin A record in Cloud DNS.

Answer: B

Explanation:

Question: 133

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

Explanation:

Question: 134

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 appliances. Configure a route for 10.0.0.0/8 with the network load balancer as the next hop.

D. 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

Explanation:

Question: 135

You are responsible for enabling Private Google Access for the virtual machine (VM) instances in your Virtual Private Cloud (VPC) to access Google APIs. All VM instances have only a private IP address and need to access Cloud Storage. You need to ensure that all VM traffic is routed back to your onpremises data center for traffic scrubbing via your existing Cloud Interconnect connection. However, VM traffic to Google APIs should remain in the VPC. What should you do?

A. Delete the default route in your VPC.

Create a private Cloud DNS zone for googleapis.com, create a CNAME for *.googleapis.com to restricted googleapis.com, and create an A record for restricted googleapis com that resolves to the addresses in 199.36.153.4/30.

Create a static route in your VPC for the range 199.36.153.4/30 with the default internet gateway as the next hop.

B. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).

Create a public Cloud DNS zone with a CNAME for *.google.com to private googleapis com, create a

CNAME for * googleapis.com to private googleapis com, and create an A record for Private googleapis.com that resolves to the addresses in 199.36.153 8/30.

Create a static route in your VPC for the range 199 .36.153.8/30 with the default internet gateway as the next hop.

C. Configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP) with a lower priority (MED) than the default VPC route.

Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googieapis.com to private googleapis com,

and create an A record for private.googleapis.com that resolves to the addresses in 199 .36.153.8/30.

Create a static route in your VPC for the range 199.36. 153.8/30 with the default internet gateway as the next hop.

D. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).

Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googieapis.com to Private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30.

Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.

Answer: C

Explanation:

Question: 136

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.

Import the custom routes in the spokes.

B. 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. Import the custom routes in the spokes.

Delete the default internet gateway route of the spokes.

C. 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. Import the custom routes in the spokes.

D. 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

Explanation:

Question: 137

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 onpremises 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

Explanation:

Question: 138

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

Explanation:

Question: 139

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

Explanation:

Question: 140

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 onpremises 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 onpremises 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 onpremises 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 onpremises 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

Explanation:

Question: 141

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: allow

Target: VM B service account

Source ranges: VM A service account

Priority: 1000

B. Firewall rule direction: ingress

Action: allow

Target: specific VM B tag

Source ranges: VM A tag and VM A source IP address

Priority: 1000

C. Firewall rule direction: ingress

Action: allow

Target: VM A service account

Source ranges: VM B service account and VM B source IP address

Priority: 100

D. Firewall rule direction: ingress

Action: allow

Target: specific VM A tag

Source ranges: VM B tag and VM B source IP address

Priority: 100

Answer: D

Explanation:

Question: 142

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

A. Use the Network Intelligence Center Connectivity Tests to test the connectivity between the VPC

and the on-premises network.

B. Use Network Intelligence Center Network Topology to check the traffic flow, and replay the traffic from the time period when the connectivity issue occurred.

C. Configure VPC Flow Logs. Review the logs by filtering on the source and destination.

D. Configure a Compute Engine instance on the same VPC as the service running on Google Cloud to run a traceroute targeted at the on-premises service.

Answer: B

Explanation:

Question: 143

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

Explanation:

Question: 144

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

Explanation:

Question: 145

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. Create a VLAN attachment, a Cloud Router in us-west1, and a Border Gateway Protocol (BGP) session between your Cloud Router and your router.

B. Order a Direct Peering connection in the same metropolitan area. Configure a Border Gateway Protocol (BGP) session between Google and your router.

C. Configure HA VPN in us-west1. Configure a Border Gateway Protocol (BGP) session between your Cloud Router and your on-premises data center.

D. Order a Carrier Peering connection in the same metropolitan area. Configure a Border Gateway Protocol (BGP) session between Google and your router.

Answer: B

Explanation:

Question: 146

You suspect that one of the virtual machines (VMs) in your default Virtual Private Cloud (VPC) is under a denial-of-service attack. You need to analyze the incoming traffic for the VM to understand where the traffic is coming from. What should you do?

A. Enable Data Access audit logs of the VPC. Analyze the logs and get the source IP addresses from the subnetworks.get field.

B. Enable VPC Flow Logs for the subnet. Analyze the logs and get the source IP addresses from the connection field.

C. Enable VPC Flow Logs for the VPC. Analyze the logs and get the source IP addresses from the src_location field.

D. Enable Data Access audit logs of the subnet. Analyze the logs and get the source IP addresses from the networks.get field.

Answer: B

Explanation:

Question: 147

You are responsible for configuring firewall policies for your company in Google Cloud. Your security team has a strict set of requirements that must be met to configure firewall rules.

Always allow Secure Shell (SSH) from your corporate IP address.

Restrict SSH access from all other IP addresses.

There are multiple projects and VPCs in your Google Cloud organization. You need to ensure that other VPC firewall rules cannot bypass the security team's requirements. What should you do?

A. Configure a hierarchical firewall policy to the organization node to allow TCP port 22 for your corporate IP address with priority 0.

Configure a hierarchical firewall policy to the organization node to deny TCP port 22 for all IP addresses with priority 1.

B. Configure a VPC firewall rule to allow TCP port 22 for your corporate IP address with priority 0.

Configure a VPC firewall rule to deny TCP port 22 for all IP addresses with priority 1.

C. Configure a VPC firewall rule to allow TCP port 22 for your corporate IP address with priority 1.

Configure a VPC firewall rule to deny TCP port 22 for all IP addresses with priority 0.

D. Configure a hierarchical firewall policy to the organization node to allow TCP port 22 for your corporate IP address with priority 1

Configure a hierarchical firewall policy to the organization node to deny TCP port 22 for all IP addresses with priority 0.

Answer: A

Explanation:

Question: 148

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

Explanation:

Question: 149

You work for a university that is migrating to Google Cloud.

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 Dedicated Interconnect in the host project.

B. Use Shared VPC, and deploy the VLAN attachments in the service projects. Connect the VLAN attachment to the Shared VPC's host project.

C. Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects' Dedicated Interconnects.

D. Use standalone projects and deploy the VLAN attachments and Dedicated Interconnects in each of the individual projects.

Answer: A

Explanation:

Question: 150

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. Peer the serverless service VPC to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.

B. Create a serverless VPC access connector for each serverless service. Configure the connectors to allow traffic between the serverless services and your existing microservices.

C. Deploy your serverless services to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.

D. Create a serverless VPC access connector. Configure the serverless service to use the connector for communication to the microservices.

Answer: D

Explanation:

Question: 151

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 region. Configure the on-premises router to advertise routes with the same multi-exit discriminator (MED).

B. Configure an additional VLAN attachment of 10 Gbps in the same region. Configure the onpremises router to advertise routes with the same multi-exit discriminator (MED).

C. From the Google Cloud Console, modify the bandwidth of the VLAN attachment to 20 Gbps.

D. From the Google Cloud Console, request a new Dedicated Interconnect connection of 20 Gbps, and configure a VLAN attachment of 10 Gbps.

E. Configure Link Aggregation Control Protocol (LACP) on the on-premises router to use the 20-Gbps Dedicated Interconnect connection.

Answer: C, E

Explanation:

Question: 152

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

A. Use regional routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1

B. Use global routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1

C. Use regional routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us- west1 Cloud Router to a base priority of 1

D. Use global routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1

Answer: A

Explanation:

Question: 153

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

```
gcloud container clusters describe customer-l-cluster —zone us-centrall-c

cluster!pv4Cidr: 192.168.36.0/24 endpoint: 192.168.38.2 ipAllocationPolicy:
  cluster!pv4Cidr: 192.168.36.0/24
  cluster!pv4CidrBlock: 192.168.36.0/24 clusterSecondaryRangeName: customer-l-pods
  services!pv4Cidr: 192.168.37.0/24 services!p4CidrBlock: 192.168.37.0/24
  servicesSecondaryRangeName: customer-l-svc uselpAliases: true

masterAuthorizedNetworksConfig:

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

services!pv4Cidr: 192.162.37.0/24

subnetwork: customer-l-nodes zone: us-centrall-c
```

You have a virtual machine (VM) deployed in the same VPC in the subnetwork kubernetes- management with internal IP address 192.168.40 2/24 and no external IP address assigned. You need to communicate with the cluster master using kubectl. What should you do?

A. Add the network 192.168.40.0/24 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2.

B. Add the network 192.168.38.0/28 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2

C. Add the network 192.168.36.0/24 to the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 192.168.38.2

D. Add an external IP address to the VM, and add this IP address in the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 35.224.37.17.

Answer: A

Explanation:

Question: 154

You have the networking configuration shown. In the diagram Two VLAN attachments associated With two Dedicated

Interconnect connections terminate on the same Cloud Router (mycloudrouter). The Interconnect connections terminate on two separate on-premises routers. You advertise the same prefixes from the Border Gateway Protocol (BOP) sessions associated with each Of the VLAN attachments.

You notice an asymmetric traffic flow between the two Interconnect connections. Which of the following actions should you take to troubleshoot the asymmetric traffic flow?

A. From the Google Cloud console, navigate to the Hybrid Connectivity select the Cloud Router, and view BGP sessions.

B. From the Cloud CLI, run gcloud compute –protect_ID router get—status mycloudrouter —-region REGION and review the results.

C. From the Google Cloud console, navigate to Cloud Logging to view VPC Flow Logs and review the results

D. From the Cloud CLI. run gcloud compute routers describe mycloudrouter

--region REGION and review the results

Answer: A

Explanation:

The correct answer is B. From the Cloud CLI, run gcloud compute --project_ID router get-status mycloudrouter --

region REGION and review the results.

This command will show you the BGP session status, the advertised and learned routes, and the last error for each VLAN attachment. You can use this information to troubleshoot the asymmetric traffic flow and identify any issues with the BGP configuration or the Interconnect connections.

The other options are not correct because:

Option A will only show you the BGP session status, but not the advertised and learned routes or the last error for each VLAN attachment.

Option C will only show you the VPC Flow Logs, which are useful for monitoring and troubleshooting network performance and security issues within your VPC network, but not for your Interconnect connections.

Option D will only show you the basic information about the Cloud Router, such as its name, region, network, and BGP settings, but not the detailed status of each VLAN attachment.

Question: 155

You are in the process of deploying an internal HTTP(S) load balancer for your web server virtual machine (VM) Instances What two prerequisite tasks must be completed before creating the load balancer?

Choose 2 answers

A. Choose a region.

B. Create firewall rules for health checks

C. Reserve a static IP address for the load balancer

E. Determine the subnet mask for a proxy-only subnet.

E. Determine the subnet mask for Serverless VPC Access.

Answer: BC

Explanation:

The correct answer is B and C. You must create firewall rules for health checks and reserve a static IP address for the load balancer before creating the internal HTTP(S) load balancer.

The other options are not correct because:

Option A is not a prerequisite task. You can choose a region when you create the load balancer, but you do not need to do it beforehand.

Option D is not a prerequisite task. You can determine the subnet mask for a proxy-only subnet when you create the subnet, but you do not need to do it beforehand.

Option E is not related to the internal HTTP(S) load balancer. Serverless VPC Access is a feature that allows you to connect your serverless applications to your VPC network, but it is not required for the load balancer.

Question: 156

You want Cloud CDN to serve the https://www.example.com/images/spacetime.png static image file that is hosted in a private Cloud Storage bucket, You are using the VSE ORIG.-X_NZADERS cache mode You receive an HTTP 403 error when opening the file In your browser and you see that the HTTP response has a Cache-control: private, max-age=O header How should you correct this Issue?

A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role

B. Change the cache mode to cache all content.

C. Increase the default time-to-live (TTL) for the backend service.

D. Enable negative caching for the backend bucket

Answer: A

Explanation:

The correct answer is A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role.

This answer is based on the following facts:

Cloud CDN can serve private content from Cloud Storage buckets, but you need to grant the appropriate permissions to the Google-managed service account that represents your load balancer1.

The Storage Legacy Object Reader role grants read access to objects in a bucket2.

The Cache-control: private header indicates that the object is not publicly readable and requires authentication3.

The USE_ORIGIN_HEADERS cache mode instructs Cloud CDN to cache responses based on the CacheControl and Expires headers from the origin server4. Changing the cache mode, increasing the TTL, or enabling negative caching will not affect the 403 error.

Question: 157

You are deploying an application that runs on Compute Engine instances. You need to determine how to expose your application to a new customer You must ensure that your application meets the following requirements

- Maps multiple existing reserved external IP addresses to the Instance
- Processes IP Encapsulating Security Payload (ESP) traffic

What should you do?

A. Configure a target pool, and create protocol forwarding rules for each external IP address.

B. Configure a backend service, and create an external network load balancer for each external IP address

C. Configure a target instance, and create a protocol forwarding rule for each external IP address to be mapped to the instance.

D. Configure the Compute Engine Instances' network Interface external IP address from None to Ephemeral Add as many external IP addresses as required

Answer: C

Explanation:

The correct answer is C. Configure a target instance, and create a protocol forwarding rule for each external IP address to be mapped to the instance.

This answer is based on the following facts:

A target instance is a Compute Engine instance that handles traffic from one or more forwarding rules1. You can use

target instances to forward traffic to a single VM instance from one or more external IP addresses2.

A protocol forwarding rule specifies the IP protocol and port range for the traffic that you want to forward3. You can use protocol forwarding rules to forward traffic of any IP protocol, including ESP4.

The other options are not correct because:

Option A is not possible. You cannot create protocol forwarding rules for a target pool. A target pool is a group of instances that receives traffic from a network load balancer5.

Option B is not suitable. You do not need to create an external network load balancer for each external IP address. An external network load balancer distributes traffic among multiple backend instances based on the destination IP address and port. You can use a single load balancer with multiple forwarding rules to map multiple external IP addresses to the same backend service.

Option D is not feasible. You cannot add multiple external IP addresses to a single network interface of a Compute Engine instance. Each network interface can have only one external IP address that is either ephemeral or static. You can use alias IP ranges to assign multiple internal IP addresses to a single network interface, but not external IP addresses.

Question: 158

You are planning to use Terraform to deploy the Google Cloud infrastructure for your company The design must meet the following requirements

- Each Google Cloud project must represent an Internal project that your team Will work on
- After an internal project is finished, the infrastructure must be deleted
- Each Internal project must have Its own Google Cloud project owner to manage the Google Cloud resources-
- You have 10-100 projects deployed at a time,

While you are writing the Terraform code, you need to ensure that the deployment IS Simple, and the code IS reusable With

centralized management What should you doo

A. Create a Single pt0Ject and additional VPCs for each Internal project

B. Create a Single Project and Single VPC for each internal project

C. Create a single Shared VPC and attach each Google Cloud project as a service project

D. Create a Shared VPC and service project for each Internal project

Answer: C

Explanation:

The correct answer is C. Create a single Shared VPC and attach each Google Cloud project as a service project.

This answer is based on the following facts:

A Shared VPC allows you to share one or more VPC networks across multiple Google Cloud projects1. This simplifies the deployment and management of the network infrastructure, as you only need to create and maintain one VPC network for all your internal projects.

A Shared VPC consists of a host project that owns the VPC network and one or more service projects that use the VPC network2. You can attach and detach service projects as needed, depending on the lifecycle of your internal projects. You can also delete service projects without affecting the host project or other service projects.

A Shared VPC allows you to delegate administrative roles to different project owners3. You can grant the Shared VPC Admin role to the owner of the host project, who can manage the VPC network and its subnets. You can also grant the Service Project Admin role to the owners of the service projects, who can manage the Google Cloud resources in their own projects.

The other options are not correct because:

Option A is not suitable. Creating a single project and additional VPCs for each internal project will increase the complexity and cost of the network infrastructure. You will need to create and maintain multiple VPC networks, firewall rules, routes, and VPN tunnels. You will also have a limit on the number of VPC networks per project4.

Option B is not feasible. Creating a single project and single VPC for each internal project will not meet the requirement of having separate project owners for each internal project. You will have only one project owner who can manage all the Google Cloud resources in the same project.

Option D is not optimal. Creating a Shared VPC and service project for each internal project will not meet the requirement of having a simple and reusable code with centralized management. You will need to create and maintain multiple Shared VPCs, which will increase the complexity and cost of the network infrastructure. You will also have more Terraform code to write and manage for each Shared VPC.

Question: 159

Your company recently migrated to Google Cloud in a Single region. You configured separate Virtual Private Cloud (VPC) networks for two departments. Department A and Department B. Department A has requested access to resources that are part Of Department Bis VPC. You need to configure the traffic from private IP addresses to flow between the VPCs using multi-NIC virtual machines (VMS) to meet security requirements Your configuration also must

- Support both TCP and UDP protocols
- Provide fully automated failover
- Include health-checks

Require minimal manual Intervention In the client VMS

Which approach should you take?

A. Create the VMS In the same zone, and configure static routes With IP addresses as next hops.

B. Create the VMS in different zones, and configure static routes with instance names as next hops

C. Create an Instance template and a managed instance group. Configure a Single internal load balancer, and define a custom static route with the Internal TCP/UDP load balancer as the next hop

D. Create an instance template and a managed instance group. Configure two separate internal TCP/IJDP load balancers for each protocol (TCP!UDP), and configure the client VIVIS to use the internal load balancers' virtual IP addresses

Answer: D

Explanation:

The correct answer is D. Create an instance template and a managed instance group. Configure two separate internal

TCP/UDP load balancers for each protocol (TCP/UDP), and configure the client VMs to use the internal load balancers' virtual IP addresses.

This answer is based on the following facts:

Using multi-NIC VMs as network virtual appliances (NVAs) allows you to route traffic between different VPC networks1. You can use NVAs to implement custom network policies and security requirements.

Using an instance template and a managed instance group allows you to create and manage multiple identical NVAs2. You can also use health checks and autoscaling policies to ensure high availability and reliability of your NVAs.

Using internal TCP/UDP load balancers allows you to distribute traffic from client VMs to NVAs based on the protocol and port3. You can also use health checks and failover policies to ensure that only healthy NVAs receive traffic.

Configuring the client VMs to use the internal load balancers' virtual IP addresses allows you to simplify the routing configuration and avoid manual intervention4. You do not need to create static routes or update them when NVAs are added or removed.

The other options are not correct because:

Option A is not suitable. Creating the VMs in the same zone does not provide high availability or failover. Using static routes with IP addresses as next hops requires manual intervention when NVAs are added or removed.

Option B is not optimal. Creating the VMs in different zones provides high availability, but not failover. Using static routes with instance names as next hops requires manual intervention when NVAs are added or removed.

Option C is not feasible. Creating an instance template and a managed instance group provides high availability and reliability, but using a single internal load balancer does not support both TCP and UDP protocols. You cannot define a custom static route with an internal load balancer as the next hop.

Question: 160

You are designing an IP address scheme for new private Google Kubernetes Engine (GKE) clusters, Due to IP address exhaustion of the RFC 1918 address space in your enterprise, you plan to use privately used public IP space for the new dusters. You want to follow Google-recommended practices, What should you do after designing your IP scheme?

A. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters. Reuse the secondary address range for the pods across multiple private GKE clusters.

B. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters Reuse the secondary address range for the services across multiple private GKE clusters.

C. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster With the following options selected: --enab1e-ip-a1ias and --enable-private- nodes.

D. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected and – siable-default-snat, --enable-ip-alias, and –enable-private-nodes

Answer: D

Explanation:

The correct answer is D. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected: --disable-default-snat, -enable-ip-alias, and --enable-private-nodes.

This answer is based on the following facts:

Privately used public IP (PUPI) addresses are any public IP addresses not owned by Google that a customer can use privately on Google Cloud1. You can use PUPI addresses for GKE pods and services in private clusters to mitigate address exhaustion.

A private GKE cluster is a cluster that has no public IP addresses on the nodes2. You can use private clusters to isolate your workloads from the public internet and enhance security.

The --disable-default-snat option disables source network address translation (SNAT) for the cluster3. This option allows you to use PUPI addresses without conflicting with other public IP addresses on the internet.

The --enable-ip-alias option enables alias IP ranges for the cluster4. This option allows you to use separate subnet ranges for nodes, pods, and services, and to specify the size of those ranges.

The --enable-private-nodes option enables private nodes for the cluster5. This option ensures that the nodes have no public IP addresses and can only communicate with other Google Cloud resources in the same VPC network or peered networks.

The other options are not correct because:

Option A is not suitable. Creating RFC 1918 primary and secondary subnet IP ranges for the clusters does not solve the problem of address exhaustion. Re-using the secondary address range for pods across multiple private GKE clusters can cause IP conflicts and routing issues.

Option B is also not suitable. Creating RFC 1918 primary and secondary subnet IP ranges for the clusters does not solve the problem of address exhaustion. Re-using the secondary address range for services across multiple private GKE clusters can cause IP conflicts and routing issues.

Option C is not feasible. Creating privately used public IP primary and secondary subnet ranges for the clusters is a valid step, but creating a private GKE cluster with only --enable-ip-alias and --enableprivate-nodes options is not enough. You also need to disable default SNAT to avoid IP conflicts with other public IP addresses on the internet.

Question: 161

Your company runs an enterprise platform on-premises using virtual machines (VMS). Your internet customers have created tens of thousands of DNS domains panting to your public IP addresses allocated to the Vtvls Typically, your customers hard-code your IP addresses In their DNS records You are now planning to migrate the platform to Compute Engine and you want to use Bring your Own IP you want to minimize disruption to the Platform What Should you d0?

A. Create a VPC and request static external IP addresses from Google Cloud Assagn the IP addresses to the Compute Engine instances. Notify your customers of the new IP addresses so they can update their DNS

B. Verify ownership of your IP addresses. After the verification, Google Cloud advertises and provisions the IP prefix for you_ Assign the IP addresses to the Compute Engine Instances

C. Create a VPC With the same IP address range as your on-premises network Asson the IP addresses to the Compute Engine Instances.

D. Verify ownership of your IP addresses. Use live migration to import the prefix Assign the IP addresses to Compute Engine instances.

Answer: D

Explanation:

The correct answer is D because it allows you to use your own public IP addresses in Google Cloud without disrupting the platform or requiring your customers to update their DNS records. Option A is incorrect because it involves changing the IP addresses and notifying the customers, which can cause disruption and errors. Option B is incorrect because it does not use live migration, which is a feature that lets you control when Google starts advertising routes for

your prefix. Option C is incorrect because it does not involve bringing your own IP addresses, but rather using Google-provided IP addresses.

References:

Bring your own IP addresses

Professional Cloud Network Engineer Exam Guide

Bring your own IP addresses (BYOIP) to Azure with Custom IP Prefix

Question: 162

You ate planning to use Terraform to deploy the Google Cloud infrastructure for your company, The design must meet the following requirements

- Each Google Cloud project must represent an Internal project that your team Will work on
- After an Internal project is finished, the infrastructure must be deleted
- Each Internal project must have Its own Google Cloud project owner to manage the Google Cloud resources.
- You have 10—100 projects deployed at a time

While you are writing the Terraform code, you need to ensure that the deployment is simple and the code is reusable With
centralized management What should you do?

A. Create a Single project and additional VPCs for each internal project

B. Create a Single Shared VPC and attach each Google Cloud project as a service project

C. Create a Single project and Single VPC for each internal project

D. Create a Shared VPC and service project for each internal project

Answer: D

Explanation:

The correct answer is D because it meets the following requirements:

Each internal project has its own Google Cloud project, which can be easily created and deleted by Terraform using the google_project resource1.

Each internal project has its own Google Cloud project owner, which can be assigned by Terraform using the google_project_iam_member resource1.

The deployment is simple and the code is reusable with centralized management, because the Shared VPC allows you to connect multiple service projects to a single host project that contains the network resources2. This way, you can use Terraform modules to create and manage the network resources in the host project, and then reference them in the service projects3.

Option A is incorrect because it does not create separate Google Cloud projects for each internal project, which makes it harder to delete the infrastructure and assign project owners. Option B is incorrect because it does not create separate Google Cloud projects for each internal project, and also because it attaches the service projects to a Shared VPC, which is not recommended for shortlived projects2. Option C is incorrect because it does not use a Shared VPC, which means that each internal project has to create and manage its own network resources, which increases complexity and reduces reusability.

References:

google_project - Terraform Registry

Managing infrastructure as code with Terraform, Cloud Build, and GitOps | Google Cloud

Automating your automation by Creating Google Cloud Projects Automatically

Question: 163

Your team is developing an application that will be used by consumers all over the world. Currently, the application sits behind a global external application load balancer You need to protect the application from potential application-level attacks. What should you do?

A. Enable Cloud CDN on the backend service.

B. Create multiple firewall deny rules to block malicious users, and apply them to the global external application load balancer

C. Create a Google Cloud Armor security policy with web application firewall rules, and apply the security policy to the backend service.

D. Create a VPC Service Controls perimeter with the global external application load balancer as the protected service, and apply it to the backend service

Answer: C

Explanation:

The correct answer is C because it meets the requirement of protecting the application from potential application-level attacks. Google Cloud Armor security policies are sets of rules that match on attributes from Layer 3 to Layer 7 to protect externally facing applications1. Web application firewall (WAF) rules are predefined rules that detect and mitigate common web attacks such as crosssite scripting (XSS), SQL injection, remote file inclusion, and more2. By applying a Google Cloud
Armor security policy with WAF rules to the backend service, you can filter out malicious requests before they reach your application.

Option A is incorrect because Cloud CDN is a content delivery network that caches static content at the edge of Google's network, but it does not provide any protection against application-level attacks3. Option B is incorrect because firewall rules are applied at the VPC network level, not at the load balancer level4. Firewall rules also only match on Layer 3 and 4 attributes, not on Layer 7 attributes that are relevant for application-level attacks4. Option D is incorrect because VPC Service Controls perimeter is a feature that helps you secure your data from unauthorized access by users outside your organization, but it does not protect your application from external attacks.

References:

Security policy overview | Google Cloud Armor

Web application firewall (WAF) rules | Google Cloud Armor

Cloud CDN overview | Google Cloud

Using firewall rules | VPC

[VPC Service Controls overview | Google Cloud]

Question: 164

You are a network administrator at your company planning a migration to Google Cloud and you need to finish the migration as quickly as possible, To ease the transition, you decided to use the same architecture as your on-premises network' a hub-and-spoke model. Your on-premises architecture consists of over 50 spokes. Each spoke does not have connectivity to the other spokes, and all traffic IS sent through the hub for security reasons. You need to ensure that the Google Cloud architecture matches your on-premises architecture. You want to implement a solution that minimizes management overhead and cost, and uses default networking quotas and limits. What should you do?

A. Connect all the spokes to the hub with Cloud VPN.

B. Connect all the spokes to the hub with VPC Network Peering.

C. Connect all the spokes to the hub With Cloud VPN. Use a third-party network appliance as a default gateway to prevent connectivity between the spokes

D. Connect all the spokes to the hub with VPC Network Peering. Use a third-party network appliance as a default gateway to prevent connectivity between the spokes.

Answer: D

Explanation:

The correct answer is D because it meets the following requirements:

It matches the hub-and-spoke model of the on-premises network, where each spoke is a separate VPC network that is connected to a central hub VPC network.

It minimizes management overhead and cost, because VPC Network Peering is a simple and low-cost way to connect VPC networks without using any external IP addresses or VPN gateways1.

It uses default networking quotas and limits, because VPC Network Peering does not consume any quota or limit for VPN tunnels, external IP addresses, or forwarding rules2.

It prevents connectivity between the spokes, because VPC Network Peering is non-transitive by default, meaning that a spoke can only communicate with the hub, not with other spokes1. To enforce this restriction, a third-party network

appliance can be used as a default gateway in each spoke VPC network, which can filter out any traffic destined for other spokes3.

Option A is incorrect because it does not minimize cost, as Cloud VPN charges for egress traffic and requires external IP addresses for the VPN gateways4. Option B is incorrect because it does not prevent connectivity between the spokes, as VPC Network Peering allows direct communication between peered VPC networks by default1. Option C is incorrect because it does not minimize cost or use default quotas and limits, for the same reasons as option A.

References:

VPC Network Peering overview | VPC

Quotas and limits | VPC

Hub-and-spoke network architecture | Cloud Architecture Center

Cloud VPN overview | Google Cloud

Question: 165

Your company is planning a migration to Google Kubernetes Engine. Your application team informed you that they require a minimum of 60 Pods per node and a maximum of 100 Pods per node Which Pod per node CIDR range should you use?

A. /24

B. /25

C. /26

D. /28

Answer: B

Explanation:

To determine the Pod per node CIDR range, you need to calculate how many IP addresses are required for each node, and then choose the smallest CIDR range that can accommodate that number. A CIDR range of /n means that there are 2^(32-n) IP addresses available in that range. For example, a /24 range has 2^(32-24) = 256 IP addresses.

According to the question, the application team requires a minimum of 60 Pods per node and a maximum of 100 Pods

per node. Therefore, you need to choose a CIDR range that can provide at least 100 IP addresses per node, but not more than necessary. A /25 range has 2^(32-25) = 128 IP addresses, which is enough for 100 Pods per node. A /26 range has 2^(32-26) = 64 IP addresses, which is not enough for 60 Pods per node. A /24 range has 256 IP addresses, which is more than needed and wastes IP address space. A /28 range has 2^(32-28) = 16 IP addresses, which is far too small for any node.

Therefore, the best option is B. /25. This is also consistent with the Google Kubernetes Engine documentation, which states that each node is allocated a /24 range of IP addresses for Pods by default, but the maximum number of Pods per node is 1101. This means that there are approximately twice as many available IP addresses as possible Pods, which is similar to the ratio of 128 to 100 in the /25 range.

1: Configure maximum Pods per node | Google Kubernetes Engine (GKE) | Google Cloud

Question: 166

You are designing an IP address scheme for new private Google Kubernetes Engine (GKE) clusters.

Due to IP address exhaustion of the RFC 1918 address space In your enterprise, you plan to use privately used public IP space for the new clusters. You want to follow Google-recommended practices. What should you do after designing your IP scheme?

A. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters. Reuse the secondary address range for the pods across multiple private GKE clusters

B. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters Reuse the secondary address range for the services across multiple private GKE clusters

C. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected and

D. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster With the following options selected --disable-default-snat, —enable-ip-alias, and—enable-private-nodes

Answer: D

Explanation:

This answer follows the Google-recommended practices for using privately used public IP (PUPI) addresses for GKE Pod address blocks1. The benefits of this approach are:

It allows you to use any public IP addresses that are not owned by Google or your organization for your Pods, which can help mitigate address exhaustion in your enterprise.

It prevents any external traffic from reaching your Pods, as Google Cloud does not route PUPI addresses to the internet or to other VPC networks by default.

It enables you to use VPC Network Peering to connect your GKE cluster to other VPC networks that use different PUPI addresses, as long as you enable the export and import of custom routes for the peering connection.

It preserves the fully integrated network model of GKE, where Pods can communicate with nodes and other resources in the same VPC network without NAT.

The options that you need to select when creating a private GKE cluster with PUPI addresses are:

–disable-default-snat: This option disables source NAT for outbound traffic from Pods to destinations outside the cluster's VPC network. This is necessary to prevent Pods from using RFC 1918 addresses as their source IP addresses, which could cause conflicts with other networks that use the same address space2.

–enable-ip-alias: This option enables alias IP ranges for Pods and Services, which allows you to use separate subnet ranges for them. This is required to use PUPI addresses for Pods1.

–enable-private-nodes: This option creates a private cluster, where nodes do not have external IP addresses and can only communicate with the control plane through a private endpoint. This enhances the security and privacy of your cluster3.

Option A is incorrect because it does not use PUPI addresses for Pods, but rather RFC 1918 addresses. This does not solve the problem of address exhaustion in your enterprise. Option B is incorrect because it reuses the secondary address range for Services across multiple private GKE clusters, which could cause IP conflicts and routing issues. Option C is incorrect because it does not specify the options that are needed to create a private GKE cluster with PUPI addresses.

1: Configuring privately used public IPs for GKE | Kubernetes Engine | Google Cloud 2: Using Cloud NAT with GKE | Kubernetes Engine | Google Cloud 3: Private clusters | Kubernetes Engine | Google Cloud

Question: 167

You have the networking configuration shown in the diagram. A pair of redundant Dedicated Interconnect connections

(int-lgal and int-lga2) terminate on the same Cloud Router. The Interconnect connections terminate on two separate on-premises routers. You are advertising the same prefixes from the Border Gateway Protocol (BGP) sessions associated with the Dedicated Interconnect connections. You need to configure one connection as Active for both ingress and egress traffic. If the active Interconnect connection fails, you want the passive Interconnect connection to automatically begin routing all traffic Which two actions should you take to meet this requirement? (Choose Two)

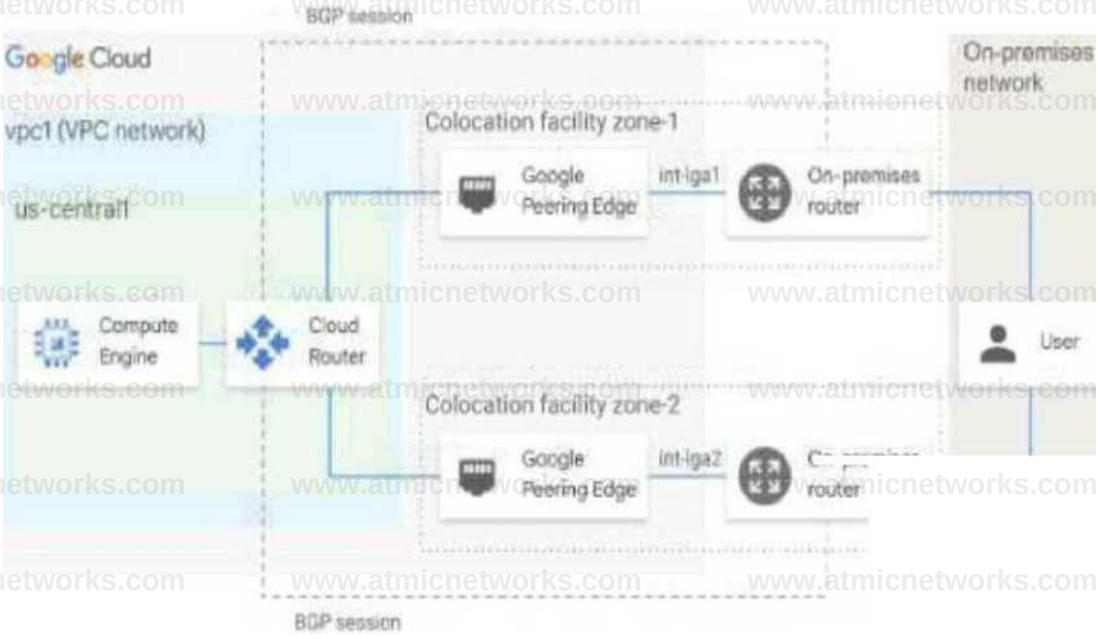

A. Configure the advertised route priority > 10,200 on the active Interconnect connection.

B. Advertise a lower MED on the passive Interconnect connection from the on-premises router

C. Configure the advertised route priority as 200 for the BGP session associated with the active Interconnect connection.

D. Configure the advertised route priority as 200 for the BGP session associated with the passive Interconnect connection.

E. Advertise a lower MED on the active Interconnect connection from the on-premises router

Answer: C, E

Explanation:

This answer meets the requirement of configuring one connection as Active for both ingress and egress traffic, and enabling automatic failover to the passive connection in case of failure. The reason is:

The advertised route priority is a value that Cloud Router uses to set the route priority when advertising routes to your on-premises router. The lower the value, the higher the priority1. By setting the advertised route priority as 200 for the active connection, you ensure that it has a higher priority than the

passive connection, which has the default value of 1001. This way, your on-premises router will prefer the routes from the active connection over the passive one for ingress traffic.

The MED (Multi-Exit Discriminator) is a value that your on-premises router uses to indicate its preference for receiving traffic from Cloud Router. The lower the value, the higher the preference2. By advertising a lower MED on the active connection from your on-premises router, you ensure that Cloud Router will prefer sending traffic to the active connection over the passive one for egress traffic.

If the active connection fails, Cloud Router will stop receiving routes from it and will start using the routes from the passive connection for egress traffic. Similarly, your on-premises router will stop receiving routes with priority 200 from the active connection and will start using the routes with priority 100 from the passive connection for ingress traffic. This achieves automatic failover without any manual intervention.

Option A is incorrect because setting the advertised route priority > 10,200 on the active connection would deprioritize it globally in your VPC network, which is not what you want1. Option B is incorrect because advertising a lower MED on the passive connection would make Cloud Router prefer sending traffic to it over the active one, which is not what you want2. Option D is incorrect because setting the advertised route priority as 200 for both connections would make them equally preferred by your on-premises router, which is not what you want1.

References:

Update the base route priority | Cloud Router | Google Cloud

Configuring BGP sessions | Cloud Router | Google Cloud

Question: 168

Your company's logo is published as an image file across multiple websites that are hosted by your company You have implemented Cloud CDN, however, you want to improve the performance of the cache hit ratio associated with this image file. What should you do?

A. Configure custom cache keys for the backend service that holds the image file, and clear the Host and Protocol checkboxes-

B. Configure Cloud Storage as a custom origin backend to host the image file, and select multi-region as the location type

C. Configure versioned IJRLs for each domain to serve users the •mage file before the cache entry expires

D. Configure the default time to live (TTL) as O for the image file.

Answer: A

Explanation:

This answer meets the requirement of improving the performance of the cache hit ratio associated with the image file. The reason is:

Custom cache keys allow you to control which parts of the request URL are used to build the cache key. The cache key is a unique identifier that Cloud CDN uses to store and retrieve cached content1.

By default, Cloud CDN uses the complete request URL, including the protocol (http or https) and the host (the domain name), to build the cache key. This means that if the same image file is requested from different domains or protocols, Cloud CDN will cache multiple copies of it, which reduces the cache hit ratio1.

By clearing the Host and Protocol checkboxes, you can tell Cloud CDN to ignore these parts of the request URL when building the cache key. This way, Cloud CDN will cache only one copy of the image file, regardless of which domain or protocol it is requested from, which improves the cache hit ratio1.

Option B is incorrect because configuring Cloud Storage as a custom origin backend does not affect the cache hit ratio. It only affects how Cloud CDN retrieves the content from the origin if it is not cached. Option C is incorrect because configuring versioned URLs for each domain does not improve the cache hit ratio. It actually worsens it, because it creates more variations of the request URL that Cloud CDN has to cache separately. Option D is incorrect because configuring the default TTL as 0 for the image file means that Cloud CDN will not cache it at all, which defeats the purpose of using Cloud CDN.

References:

Custom cache keys | Cloud CDN | Google Cloud

Question: 169

You are responsible for designing a new connectivity solution between your organization's on-premises data center and your Google Cloud Virtual Private Cloud (VPC) network Currently, there Is no end-to-end connectivity. You must ensure a service level agreement (SLA) of 99.99% availability What should you do?

A. Use one Dedicated Interconnect connection in a single metropolitan area. Configure one Cloud Router and enable global routing in the VPC.

B. Use a Direct Peering connection between your on-premises data center and Google Cloud. Configure Classic VPN with two tunnels and one Cloud Router.

C. Use two Dedicated Interconnect connections in a single metropolitan area. Configure one Cloud Router and enable global routing in the VPC.

D. Use HA VPN. Configure one tunnel from each Interface of the VPN gateway to connect to the corresponding interfaces on the peer gateway on-premises. Configure one Cloud Router and enable global routing in the VPC.

Answer: D

Explanation:

For Dedicated Interconnects: At least four Dedicated Interconnect connections, two connections in one metropolitan area (metro) and two connections in another metro. Connections that are in the same metro must be placed in different edge availability domains (metro availability zones) to achieve 99.99% availability.

For HA VPN:

HA VPN to peer VPN gateways Connect an HA VPN gateway to one or two separate peer VPN devices 99.99%

HA VPN between two Google Cloud networks Connect two Google Cloud VPC networks in a single region by using an HA VPN gateway in each network 99.99%

Question: 170

You need to create the technical architecture for hybrid connectivity from your data center to Google Cloud This will be managed by a partner. You want to follow Google-recommended practices for production-level applications. What should you do?

A. Ask the partner to install two security appliances in the data center. Configure one VPN connection

from each of these devices to Google

Cloud, and ensure that the VPN devices on-premises are in separate racks on separate power and cooling systems.

B. Configure two Partner Interconnect connections in one metropolitan area (metro). Make sure the Interconnect connections are placed in

different metro edge availability domains. Configure two VLAN attachments in a single region, and configure regional dynamic routing on

the VPC

C. Configure two Partner Interconnect connections in one metro and two connections in another metro Make sure the Interconnect

connections are placed in different metro edge availability domains. Configure two VLAN attachments in one region and two VLAN

attachments in another region, and configure global dynamic routing on the VPC

D. Configure two Partner Interconnect connections in one metro and two connections in another metro. Make sure the Interconnect connections are placed in different metro edge availability domains. Configure two VLAN attachments in one region and two VLAN attachments in another region, and configure regional dynamic routing on the VPC.

Answer: D

Explanation:

"Google's recommended practices for production-level applications" and then see overview of these 2 pages- https://cloud.google.com/network-connectivity/docs/interconnect/tutorials/production- level-overview and https://cloud.google.com/network- connectivity/docs/interconnect/tutorials/non-critical-overview .

Question: 171

You are designing a packet mirroring policy as pan of your network security architecture for your gaming workload. Your Infrastructure is located in the us-west2 region and deployed across several zones: us-west2-a. us-west2-b. and us-west2-c The Infrastructure Is running a web-based application on TCP ports 80 and 443 with other game servers that utilize the UDP protocol. You need to deploy packet mirroring policies and collector instances to monitor web application traffic while minimizing inter-zonal network egress costs.

Following Google-recommended practices, how should you deploy the packet mirroring policies and

collector instances?

A. Create three packet mirroring policies: one for each zone. Create three groups of collector instances: one group for each zone. Configure each policy to match traffic for Its zone based on instance-tags, and create a filter for TCP

traffic.

B. Create three packet mirroring policies: one for each zone. Create three groups of collector instances: one group for each zone. Configure each policy to match traffic for its zone based on subnets, and create a filter for TCP traffic

C. Create one packet mirroring policy for the us-west2 region. Create one group of collector instances for the us-west2 region Configure the packet mirroring policy to match traffic for web server instances based on instance-tags, and create a filter for TCP traffic.

D. Create three packet mirroring policies: one for each zone. Create one group of collector instances for the us-west2 region. Configure each packet mirroring policy to match traffic for its zone based on instance-tags, and create a filter for TCP traffic

Answer: D

Explanation:

Create Packet Mirroring Policies:

You need to create three packet mirroring policies, one for each zone (us-west2-a, us-west2-b, and us-west2-c). This ensures that each zone's traffic is mirrored appropriately without unnecessary cross-zone traffic.

Create Collector Instances:

Set up one group of collector instances for the us-west2 region. Having a single group of collector instances for the entire region minimizes the number of instances required and simplifies the management while keeping egress costs low since the collectors are within the same region.

Configuration of Policies:

Each packet mirroring policy should be configured to match traffic for its specific zone. Use instancetags to identify and match the relevant instances within each zone. This helps in correctly capturing the traffic from the appropriate sources.

Filter for TCP Traffic:

Create a filter for TCP traffic (ports 80 and 443). This step ensures that only the relevant web application traffic is mirrored, reducing the amount of data processed and improving efficiency.

Cost Efficiency:

By having packet mirroring policies specific to each zone and a regional collector group, you reduce inter-zonal network egress costs. The data remains within the same region, avoiding extra charges associated with cross-zone traffic.

References:

Google Cloud Packet Mirroring Documentation

Best Practices for Packet Mirroring

Cost Management in Google Cloud

This solution aligns with Google-recommended practices by ensuring efficient traffic capture, minimal inter-zonal costs, and streamlined management of the packet mirroring setup.

Question: 172

You have the following Shared VPC design VPC Flow Logs is configured for Subnet-1 In the host VPC.
You also want to monitor flow logs for Subnet-2. What should you do?

A. Configure a firewall rule to permit Subnet-2 IP addresses outbound in the host protect VPC.

B. Configure Packet Mirroring in both the host and service project VPCs.

C. Configure a VPC Flow Logs filter for Subnet-2 in the host project VPC.

D. Configure VPC Flow Logs in the service project VPC for Subnet-2.

Answer: D

Explanation:

Understanding VPC Flow Logs:

VPC Flow Logs is a feature that captures information about the IP traffic going to and from network interfaces in a VPC. It helps in monitoring and analyzing network traffic, ensuring security, and optimizing network performance.

Current Configuration:

According to the diagram, VPC Flow Logs is already configured for Subnet-1 in the host VPC. This means that traffic information for Subnet-1 is being captured and logged.

Requirement for Subnet-2:

The goal is to monitor flow logs for Subnet-2, which is in the service project VPC.

Correct Configuration for Subnet-2:

To monitor the flow logs for Subnet-2, you need to configure VPC Flow Logs within the service project VPC where Subnet-2 resides. This is because VPC Flow Logs must be configured in the same project and VPC where the subnet is located.

Implementation Steps:

Go to the Google Cloud Console.

Navigate to the service project where Subnet-2 is located.

Select the VPC network containing Subnet-2.

E. able VPC Flow Logs for Subnet-2 by editing the subnet settings and enabling the flow logs option.

Cost and Performance Considerations:

Enabling VPC Flow Logs may incur additional costs based on the volume of data logged. Ensure to review and understand the pricing implications.

Analyze and manage the data collected to avoid unnecessary logging and costs.

References:

Google Cloud VPC Flow Logs Documentation

Configuring VPC Flow Logs

Shared VPC Overview

By configuring VPC Flow Logs in the service project VPC for Subnet-2, you ensure that traffic data is correctly captured and monitored, adhering to Google Cloud's best practices.

Question: 173

Your multi-region VPC has had a long-standing HA VPN configured in "region 1" connected to your corporate network. You are planning to add two 10 Gbps Dedicated Interconnect connections and VLAN attachments in "region 2" to connect to the same corporate network. You need to plan for connectivity between your VPC and corporate network to ensure that traffic uses the Dedicated Interconnect connections as the primary path and the HA VPN as the secondary path. What should you do?

A. Enable regional dynamic routing mode on the VPC. Configure BGP associated with the HA VPN in "region 1" to use a base priority value of 100. Configure BGP associated with the VLAN attachments to use a base priority of 20000.

Configure your on-premises routers to use similar multi-exit discriminator (MED) values.

B. Enable global dynamic routing mode on the VPC. Configure BGP associated with the HA VPN in "region 1" to use a base priority value of 100. Configure BGP associated with the VLAN attachments to use a base priority of 20000. Configure your on-premises routers to use similar multi-exit discriminator (MED) values.

C. Enable regional dynamic routing mode on the VPC. Configure BGP associated with the HA VPN in "region 1" to use a base priority value of 20000. Configure BGP associated with the VLAN attachments to use a base priority of 100. Configure your on-premises routers to use similar multiexit discriminator (MED) values.

D. Enable global dynamic routing mode on the VPC. Configure BGP associated with the HA VPN in "region 1" to use a base priority value of 20000. Configure BGP associated with the VLAN attachments to use a base priority of 100. Configure your on-premises routers to use similar multiexit discriminator (MED) values.

Answer: B

Explanation:

For the Dedicated Interconnect to be the primary connection over the HA VPN, you should:

Enable global dynamic routing mode to allow the VPC to distribute routes dynamically across regions.

Set the BGP priority for the VLAN attachments associated with the Dedicated Interconnect to a lower base priority (e.g., 100) than the HA VPN's priority (e.g., 20000) to ensure it is preferred.

Setting up global dynamic routing with adjusted BGP priorities on both Interconnect and VPN will allow dynamic routing of traffic based on set preferences and path attributes, such as MED and priority levels. This setup ensures the Dedicated Interconnect, with a lower priority value, becomes the primary path for traffic, while the HA VPN, with a higher priority, serves as a backup.

Reference: Google Cloud - Cloud Interconnect

Reference: Google Cloud - HA VPN Overview

Question: 174

Your organization has a subset of applications in multiple regions that require internet access. You need to control internet access from applications to URLs, including hostnames and paths. The compute instances that run these applications have an associated secure tag. What should you do?

A. Deploy a Cloud NAT gateway. Use fully qualified domain name (FQDN) objects in the firewall policy rules to filter outgoing traffic to specific domains from machines that match the secure tag.

B. Deploy a single Secure Web Proxy instance with global access enabled. Apply a Secure Web Proxy policy to allow

access from machines that match the secure tag to the URLs defined in a URL list.

C. Deploy a Secure Web Proxy instance in each region. Apply a Secure Web Proxy policy to allow access from machines that match the secure tag to the URLs defined in a URL list.

D. Deploy a Cloud NAT gateway. Use fully qualified domain name (FQDN) objects in the firewall policy rules to filter outgoing traffic to specific domains from machines that match a service account.

Answer: B

Explanation:

To control internet access on a per-URL basis (including hostname and path), you should deploy Secure Web Proxy with global access enabled. The Secure Web Proxy will allow policy-based filtering of web traffic, allowing control over which URLs can be accessed based on the URL list defined in the policy. Unlike Cloud NAT, which does not support FQDN filtering, Secure Web Proxy is designed to provide such control, especially for scenarios with sensitive or controlled internet access requirements.

Reference: Google Cloud - Secure Web Proxy Overview

Reference: Google Cloud - Setting up URL filtering

Question: 175

You are troubleshooting connectivity issues between Google Cloud and a public SaaS provider. Connectivity between the two environments is through the public internet. Your users are reporting intermittent connection errors when using TCP to connect; however, ICMP tests show no failures. According to users, errors occur around the same time every day. You want to troubleshoot and gather information by using Google Cloud tools that are most likely to provide insights into what is occurring within Google Cloud. What should you do?

A. Create a Connectivity Test by using TCP, the source IP address of your test VM, and the destination IP address of the public SaaS provider. Review the live data plane analysis and take the next steps based on the test results.

B. Enable and review Cloud Logging on your Cloud NAT gateway. Look for logs with errors matching the destination IP address of the public SaaS provider.

C. Enable the Firewall insights API. Set the deny rule insights observation period to one day. Review the insights to assure there are no firewall rules denying traffic.

D. Enable and review Cloud Logging for Cloud Armor. Look for logs with errors matching the destination IP address of the public SaaS provider.

Answer: A

Explanation:

When troubleshooting connectivity issues, especially over public internet connections with intermittent errors, Connectivity Tests in Network Intelligence Center are crucial. This tool allows you to simulate the connectivity and understand the data plane status of Google Cloud resources. Since ICMP tests pass but TCP tests fail intermittently, using Connectivity Tests with TCP parameters will provide detailed insight into possible network issues like route misconfigurations, peering issues, or other transient problems affecting only specific protocols.

Reference: Google Cloud - Network Intelligence Center

Reference: Google Cloud - Troubleshooting with Connectivity Tests

Question: 176

Recently, your networking team enabled Cloud CDN for one of the external-facing services that is exposed through an external Application Load Balancer. The application team has already defined which content should be cached within the responses. Upon testing the load balancer, you did not observe any change in performance after the Cloud CDN enablement. You need to resolve the issue. What should you do?

A. Configure the CACHE_MAX_STATIC caching mode on Cloud CDN to ensure Cloud CDN caches

content depending on responses from the backends.

B. Configure the USE_ORIGIN_HEADERS caching mode on Cloud CDN to ensure Cloud CDN caches content based on response headers from the backends.

C. Configure the CACHE_ALL_STATIC caching mode on Cloud CDN to ensure Cloud CDN caches all static content as well as content defined by the backends.

D. Configure the FORCE_CACHE_ALL caching mode on Cloud CDN to ensure all appropriate content is cached.

Answer: B

Explanation:

When enabling Cloud CDN, for caching behavior to follow the application-defined caching headers, you need to configure the USE_ORIGIN_HEADERS caching mode. This setting ensures that the Cloud CDN respects the cache control headers specified by the backend, allowing the application-defined caching rules to dictate what content gets cached. This is often required when specific caching directives are already set by the application.

Reference: Google Cloud - Cloud CDN Caching Modes

Question: 177

Your organization has an on-premises data center. You need to provide connectivity from the onpremises data center to Google Cloud. Bandwidth must be at least 1 Gbps, and the traffic must not traverse the internet. What should you do?

A. Configure HA VPN by using high availability gateways and tunnels.

B. Configure Dedicated Interconnect by creating a VLAN attachment, activate the connection, and submit the pairing key to your service provider.

C. Configure Cross-Cloud Interconnect by creating a VLAN attachment, activate the connection, and then submit the pairing key to your service provider.

D. Configure Partner Interconnect by creating a VLAN attachment, submit the pairing key to your service provider, and activate the connection.

Answer: D

Explanation:

For private connectivity with at least 1 Gbps bandwidth and without using the public internet, Partner Interconnect is the suitable choice if you do not require the 10 Gbps minimum of Dedicated Interconnect. With Partner Interconnect, you create a VLAN attachment and work with a service provider that facilitates the connection between your on-premises network and Google Cloud. This solution supports connections as low as 50 Mbps and up to 10 Gbps.

Reference: Google Cloud - Choosing the right Interconnect product

Question: 178

Your organization is deploying a mission-critical application with components in different regions due to strict compliance requirements. There are latency issues between different applications that reside in us-central1 and us-east4. The application team suspects the Google Cloud network as the source of the excessive latency despite using the Premium Network Service Tier. You need to use Google-recommended practices with the least amount of effort to verify the inter-region latency by investigating network performance. What should you do?

A. Set up the Performance Dashboard in Network Intelligence Center. Select the traffic type (cross- zonal), the metric (latency - RTT), the time period, the desired regions (us-central1 and us-east4), and the network tier.

B. Enable VPC Flow Logs for the VPC. Identify major bottlenecks from the application level using Flow Analyzer.

C. Configure two Linux VMs in each zone for each region. Install the application, and run a load test using each zone

from different regions.

D. Configure a VM with a probe in Network Intelligence Center in each zone for each region. Choose the traffic type (cross-zonal), metric (latency - RTT), desired regions (us-central1 and us-east4), and the network tier.

Answer: A

Explanation:

The Performance Dashboard in the Network Intelligence Center provides a detailed view of network latency and performance metrics. For inter-region latency issues, you can quickly identify round-trip times (RTT) and latency using this tool by selecting the specific regions and network tiers, which allows you to diagnose any anomalies or patterns impacting performance.

Reference: Google Cloud - Network Intelligence Center Performance Dashboard

Question: 179

You are configuring the firewall endpoints as part of the Cloud Next Generation Firewall (Cloud NGFW) intrusion prevention service in Google Cloud. You have configured a threat prevention security profile, and you now need to create an endpoint for traffic inspection. What should you do?

A. Attach the profile to the VPC network, create a firewall endpoint within the zone, and use a firewall policy rule to apply the L7 inspection.

B. Create a firewall endpoint within the zone, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

C. Create a firewall endpoint within the region, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

D. Create a Private Service Connect endpoint within the zone, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

Answer: C

Explanation:

For Cloud NGFW in Google Cloud, firewall endpoints are typically created at the regional level, allowing you to associate these with your VPC network for Layer 7 traffic inspection. This regional setup ensures high availability and scales the inspection service across the network.

Reference: Google Cloud - Cloud NGFW

Question: 180

Your company's current network architecture has three VPC Service Controls perimeters:

One perimeter (PERIMETER_PROD) to protect production storage buckets

One perimeter (PERIMETER_NONPROD) to protect non-production storage buckets

One perimeter (PERIMETER_VPC) that contains a single VPC (VPC_ONE)

In this single VPC (VPC_ONE), the IP_RANGE_PROD is dedicated to the subnets of the production workloads, and the IP_RANGE_NONPROD is dedicated to subnets of non-production workloads.

Workloads cannot be created outside those two ranges. You need to ensure that production workloads can access only production storage buckets and non-production workloads can access only non-production storage buckets with minimal setup effort. What should you do?

A. Develop a design that uses the IP_RANGE_PROD and IP_RANGE_NONPROD perimeters to create two access levels, with each access level referencing a single range. Create two ingress access policies with each access policy referencing one of the two access levels. Update the PERIMETER_PROD and PERIMETER_NONPROD perimeters.

B. Develop a design that removes the PERIMETER_VPC perimeter. Update the PERIMETER_NONPROD perimeter to include the project containing VPC_ONE. Remove the PERIMETER_PROD perimeter.

C. Develop a design that creates a new VPC (VPC_NONPROD) in the same project as VPC_ONE. Migrate all the non-production workloads from VPC_ONE to the PERIMETER_NONPROD perimeter. Remove the PERIMETER_VPC perimeter. Update the PERIMETER_PROD perimeter to include VPC_ONE and the PERIMETER_NONPROD perimeter to include VPC_NONPROD.

D. Develop a design that removes the PERIMETER_VPC perimeter. Update the PERIMETER_PROD perimeter to include the project containing VPC_ONE. Remove the PERIMETER_NONPROD perimeter.

Answer: A

Explanation:

Using IP range-based access levels for VPC Service Controls allows segmentation of production and non-production resources within the same VPC. By creating separate access levels and ingress policies for each IP range, you ensure that only production subnets access production buckets and non-production subnets access non-production buckets, providing the required isolation.

Reference: Google Cloud - VPC Service Controls and Access Levels

Question: 181

Your organization recently exposed a set of services through a global external Application Load Balancer. After conducting some testing, you observed that responses would intermittently yield a non-HTTP 200 response. You need to identify the error. What should you do? (Choose 2 answers)

A. Access a VM in the VPC through SSH, and try to access a backend VM directly. If the request is successful from the VM, increase the quantity of backends.

B. Enable and review the health check logs. Review the error responses in Cloud Logging.

C. Validate the health of the backend service. Enable logging on the load balancer, and identify the error response in Cloud Logging. Determine the cause of the error by reviewing the statusDetails log field.

D. Delete the load balancer and backend services. Create a new passthrough Network Load Balancer. Configure a failover group of VMs for the backend.

E. Validate the health of the backend service. Enable logging for the backend service, and identify the error response in Cloud Logging. Determine the cause of the error by reviewing the statusDetails log field.

Answer: B, C

Explanation:

To identify errors with intermittent non-HTTP 200 responses:

Enable and review health check logs for your backend to identify potential issues with backend availability or connectivity (Option B).

Enable logging on the load balancer and review Cloud Logging, particularly the statusDetails field, to gather insights on error types and sources (Option C).

These steps allow for precise error identification by leveraging both health checks and detailed logging features available through Google Cloud's external load balancer diagnostics.

Reference: Google Cloud - HTTP(S) Load Balancing Troubleshooting

Question: 182

Your organization is developing a landing zone architecture with the following requirements:

No communication between production and non-production environments.

Communication between applications within an environment may be necessary.

Network administrators should centrally manage all network resources, including subnets, routes, and firewall rules.

Each application should be billed separately.

Developers of an application within a project should have the autonomy to create their compute resources.

Up to 1000 applications are expected per environment.

What should you do?

A. Create a design that has a Shared VPC for each project. Implement hierarchical firewall policies to apply micro-segmentation between VPCs.

B. Create a design where each project has its own VPC. Ensure all VPCs are connected by a Network Connectivity Center hub that is centrally managed by the network team.

C. Create a design that implements a single Shared VPC. Use VPC firewall rules with secure tags to enforce micro-segmentation between environments.

D. Create a design that has one host project with a Shared VPC for the production environment, another host project with a Shared VPC for the non-production environment, and a service project that is associated with the corresponding host project for each initiative.

Answer: D

Explanation:

Using separate Shared VPCs for production and non-production environments in different host projects (Option D) meets all requirements. This design allows network administrators to centrally manage resources within each Shared VPC while ensuring isolation between environments and separate billing. By associating service projects with each host project, developers can manage resources within their project without affecting the overall VPC network structure.

Reference: Google Cloud - Best Practices for Shared VPC

Question: 183

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

adhering to your security team's requirements?

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

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

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

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

Answer: D

Explanation:

For environments requiring API security controls, use restricted.googleapis.com as it restricts access to Google APIs and enforces VPC Service Controls. The custom DNS and routing configuration ensures compliance with security policies by directing all API traffic to restricted endpoints while maintaining Private Google Access.

Reference: Google Cloud - Private Google Access and DNS Configuration

Question: 184

You reviewed the user behavior for your main application, which uses an external global Application Load Balancer, and found that the backend servers were overloaded due to erratic spikes in client requests. You need to limit concurrent sessions and return an HTTP 429 "Too Many Requests" response back to the client while following Google-recommended practices. What should you do?

A. Create a Cloud Armor security policy, and apply the predefined Open Worldwide Application Security Project (OWASP) rules to automatically implement the rate limit per client IP address.

B. Configure the load balancer to accept only the defined amount of requests per client IP address,

increase the backend servers to support more traffic, and redirect traffic to a different backend to burst traffic.

C. Configure a VM with Linux, implement the rate limit through iptables, and use a firewall rule to send an HTTP 429 response to the client application.

D. Create a Cloud Armor security policy, and associate the policy with the load balancer. Configure the security policy's settings as follows: action: throttle, conform-action: allow, exceed-action: deny- 429.

Answer: D

Explanation:

To control traffic spikes and enforce rate limits, configure Cloud Armor with throttle and deny-429 actions. This allows you to set rate limits per client IP and ensures that excess traffic receives an HTTP 429 response, effectively controlling overload situations per Google best practices.

Reference: Google Cloud - Cloud Armor Rate Limiting

Question: 185

Your organization has a new security policy that requires you to monitor all egress traffic payloads from your virtual machines in the us-west2 region. 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. Create an internal HTTP(S) load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

C. Create an internal TCP/UDP load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.

D. Enable VPC Flow Logs. Create a sink in Cloud Logging to send filtered egress VPC Flow Logs to the IDS.

Answer: C

Explanation:

Packet Mirroring with an internal TCP/UDP load balancer allows for comprehensive monitoring of

egress traffic, which includes payloads. This is required for integration with an IDS for detailed inspection of traffic payloads, meeting the security policy needs for monitoring and detection.

Reference: Google Cloud - Packet Mirroring

Question: 186

You are configuring the final elements of a migration effort where resources have been moved from on-premises to Google Cloud. While reviewing the deployed architecture, you noticed that DNS resolution is failing when queries are being sent to the on-premises environment. You log in to a Compute Engine instance, try to resolve an on-premises hostname, and the query fails. DNS queries are not arriving at the on-premises DNS server. You need to use managed services to reconfigure Cloud DNS to resolve the DNS error. What should you do?

A. Validate that the Compute Engine instances are using the Metadata Service IP address as their resolver. Configure an outbound forwarding zone for the on-premises domain pointing to the onpremises DNS server. Configure Cloud Router to advertise the Cloud DNS proxy range to the onpremises network.

B. Validate that there is network connectivity to the on-premises environment and that the Compute Engine instances can reach other on-premises resources. If errors persist, remove the VPC Network Peerings and recreate the peerings after validating the routes.

C. Review the existing Cloud DNS zones, and validate that there is a route in the VPC directing traffic destined to the IP address of the DNS servers. Recreate the existing DNS forwarding zones to forward all queries to the on-premises DNS servers.

D. Ensure that the operating systems of the Compute Engine instances are configured to send DNS queries to the on-premises DNS servers directly.

Answer: A

Explanation:

To resolve DNS resolution issues for on-premises domains from Google Cloud, you should use Cloud DNS outbound forwarding zones. This setup forwards DNS requests for specific domains to onpremises DNS servers. Cloud Router is needed to advertise the range for the DNS proxy service back to the on-premises environment, ensuring that DNS queries from Compute Engine instances reach the on-premises DNS servers.

Reference: Google Cloud - DNS Forwarding

Question: 187

Your organization wants to seamlessly migrate a global external web application from Compute Engine to GKE. You need to deploy a simple, cloud-first solution that exposes both applications and sends 10% of the requests to the new application. What should you do?

A. Configure a global external Application Load Balancer with a Service Extension that points to an application running in a VM, which controls which requests go to each application.

B. Configure a global external Application Load Balancer with weighted traffic splitting.

C. Configure two separate global external Application Load Balancers, and use Cloud DNS geolocation routing policies.

D. Configure a global external Application Load Balancer with weighted request mirroring.

Answer: B

Explanation:

Weighted traffic splitting allows you to gradually route a percentage of traffic to the new GKE application while still serving the majority of requests through the Compute Engine instance. This gradual transition minimizes risks and ensures seamless traffic distribution during migration.

Reference: Google Cloud - Traffic Splitting for Load Balancers

Question: 188

Your organization has distributed geographic applications with significant data volumes. You need to create a design that exposes the HTTPS workloads globally and keeps traffic costs to a minimum. What should you do?

A. Deploy a regional external Application Load Balancer with Standard Network Service Tier.

B. Deploy a regional external Application Load Balancer with Premium Network Service Tier.

C. Deploy a global external proxy Network Load Balancer with Standard Network Service Tier.

D. Deploy a global external Application Load Balancer with Premium Network Service Tier.

Answer: D

Explanation:

The global external Application Load Balancer with Premium Network Service Tier provides optimized routing and lower latency for HTTPS workloads on a global scale. Premium tier minimizes costs by avoiding multiple regional configurations while ensuring reliable performance for global users.

Reference: Google Cloud - Network Service Tiers

Question: 189

Your organization wants to deploy HA VPN over Cloud Interconnect to ensure encryption in transit over the Cloud Interconnect connections. You have created a Cloud Router and two encrypted VLAN attachments that have a 5 Gbps capacity and a BGP configuration. The BGP sessions are operational. You need to complete the deployment of the HA VPN over Cloud Interconnect. What should you do?

A. Enable MACsec on Partner Interconnect.

B. Create an HA VPN gateway and associate the gateway with your two encrypted VLAN attachments. Configure the HA VPN Cloud Router, peer VPN gateway resources, and HA VPN tunnels. Use the same Cloud Router used for the Cloud Interconnect tier.

C. Create an HA VPN gateway and associate the gateway with your two encrypted VLAN attachments. Create a new dedicated HA VPN Cloud Router peer VPN gateway resources and HA VPN tunnels.

D. Enable MACsec for Cloud Interconnect on the VLAN attachments.

Answer: B

Explanation:

For secure traffic over Cloud Interconnect, you configure an HA VPN gateway to work with existing VLAN attachments and use the same Cloud Router. This setup integrates seamlessly, leveraging the established BGP sessions for VPN tunnel configurations.

Reference: Google Cloud - HA VPN over Cloud Interconnect

Question: 190

Your organization's security team recently discovered that there is a high risk of malicious activities originating from some of your VMs connected to the internet. These malicious activities are currently undetected when TLS communication is used. You must ensure that encrypted traffic to the internet is inspected. What should you do?

A. Enable Cloud Armor TLS inspection policy, and associate the policy with the backend VMs.

B. Use Cloud NGFW Enterprise. Create a firewall rule for egress traffic with the tls-inspect flag and associate the firewall rules with the VMs.

C. Configure a TLS agent on every VM to intercept TLS traffic before it reaches the internet. Configure Sensitive Data Protection to analyze and allow/deny the content.

D. Use Cloud NGFW Essentials. Create a firewall rule for egress traffic and enable VPC Flow Logs with the TLS inspect option. Analyze the output logs content and block the outputs that have malicious activities.

Answer: B

Explanation:

Cloud NGFW Enterprise provides TLS inspection to detect and manage threats within encrypted traffic. Configuring firewall rules for TLS inspection enables granular monitoring and filtering, ensuring secure internet traffic.

Reference: Google Cloud - Next-Generation Firewall TLS Inspection

Question: 191

Your organization has approximately 100 teams that need to manage their own environments. A central team must manage the network. You need to design a landing zone that provides separate projects for each team and ensure the solution can scale. What should you do?

A. Configure VPC Network Peering and peer one of the VPCs to the service project.

B. Configure Policy-based Routing for each team.

C. Configure a Shared VPC and create a VPC network in the host project.

D. Configure a Shared VPC, and create a VPC network in the service project.

Answer: C

Explanation:

Using a Shared VPC enables centralized network management and efficient resource access by service projects. This scalable setup supports isolated environments for each team while allowing the network team to manage network policies and resources in a host project.

Reference: Google Cloud - Shared VPC Overview

Question: 192

You are designing the architecture for your organization so that clients can connect to certain Google APIs. Your plan must include a way to connect to Cloud Storage and BigQuery. You also need to ensure the traffic does not traverse the internet. You want your solution to be cloud-first and require the least amount of configuration steps. What should you do?

A. Configure Private Google Access on the VPC resource. Create a default route to the internet.

B. Configure Private Google Access on the subnet resource. Create a default route to the internet.

C. Configure Cloud NAT and remove the default route to the internet.

D. Configure a global Secure Web Proxy and remove the default route to the internet.

Answer: B

Explanation:

Enabling Private Google Access on the subnet allows VMs to access Google APIs (like Cloud Storage and BigQuery) directly, without routing traffic over the internet. This approach is cloud-native and involves minimal setup, aligning with a cloud-first strategy.

Reference: Google Cloud - Private Google Access Overview

Question: 193

Your organization has a hub and spoke architecture with VPC Network Peering, and hybrid connectivity is centralized at the hub. The Cloud Router in the hub VPC is advertising subnet routes, but the on-premises router does not appear to be receiving any subnet routes from the VPC spokes.

You need to resolve this issue. What should you do?

A. Create custom learned routes at the Cloud Router in the hub to advertise the subnets of the VPC spokes.

B. Create custom routes at the Cloud Router in the spokes to advertise the subnets of the VPC spokes.

C. Create a BGP route policy at the Cloud Router, and ensure the subnets of the VPC spokes are being announced towards the on-premises environment.

D. Create custom routes at the Cloud Router in the hub to advertise the subnets of the VPC spokes.

Answer: A

Explanation:

Creating custom learned routes at the hub's Cloud Router is required for advertising VPC spokes' subnets to the on-premises environment. This centralizes route configuration and ensures that all spoke subnet routes are propagated to the hybrid network.

Reference: Google Cloud - Cloud Router Custom Routes

Question: 194

You are troubleshooting an application in your organization's Google Cloud network that is not functioning as expected. You suspect that packets are getting lost somewhere. The application sends packets intermittently at a low volume from a Compute Engine VM to a destination on your onpremises network through a pair of Cloud Interconnect VLAN attachments. You validated that the Cloud Next Generation Firewall (Cloud NGFW) rules do not have any deny statements blocking egress traffic, and you do not have any explicit allow rules. Following Google-recommended practices, you need to analyze the flow to see if packets are being sent correctly out of the VM to isolate the issue. What should you do?

A. Create a packet mirroring policy that is configured with your VM as the source and destined to a collector. Analyze the packet captures.

B. Enable VPC Flow Logs on the subnet that the VM is deployed in with sample_rate = 1.0, and run a query in Logs Explorer to analyze the packet flow.

C. Enable Firewall Rules Logging on your firewall rules and review the logs.

D. Verify the network/attachment/egress_dropped_packet.s_count Cloud Interconnect VLAN attachment metric.

Answer: B

Explanation:

Enabling VPC Flow Logs with sample_rate = 1.0 on the VM's subnet will give detailed information about network traffic flowing to and from your VM. You can then query this data in Logs Explorer to check whether packets are leaving the VM and reaching the intended destination. This is a recommended practice for troubleshooting such network issues.

Reference: Google VPC Flow Logs Documentation

Question: 195

You have recently taken over responsibility for your organization's Google Cloud network security configurations. You want to review your Cloud Next Generation Firewall (Cloud NGFW) configurations to ensure that there are no rules allowing ingress traffic to your VMs and services from the internet. You want to avoid manual work. What should you do?

A. Use Firewall Insights, and enable insights for overly permissive rules.

B. Review Network Analyzer insights on the VPC network category.

C. Export all your Cloud NGFW rules into a CSV file and search for 0.0.0.0/0.

D. Run Connectivity Tests from multiple external sources to confirm that traffic is not allowed to ingress to your most critical services in Google Cloud.

Answer: A

Explanation:

Using Firewall Insights and enabling insights for overly permissive rules helps automate the process of identifying firewall

rules that may allow unintended ingress from the internet. This is a quick and efficient method compared to manually searching through firewall configurations.

Question: 196

You are deploying an HA VPN within Google Cloud. You need to exchange routes dynamically between your on-premises gateway and Google Cloud. You have already created an HA VPN gateway and a peer VPN gateway resource.

What should you do?

A. Create a Cloud Router, add VPN tunnels, and then configure BGP sessions.

B. Create a second HA VPN gateway, add VPN tunnels, and enable global dynamic routing.

C. Create a Cloud Router, add VPN tunnels, and enable global dynamic routing.

D. Create a Cloud Router, add VPN tunnels, and then configure static routes to your subnet ranges.

Answer: A

Explanation:

To dynamically exchange routes between Google Cloud and your on-premises gateway, you need to create a Cloud Router and configure BGP sessions after adding VPN tunnels. BGP allows for dynamic route exchange, which is essential for establishing proper communication between the

environments.

Reference: Google Cloud HA VPN with BGP

Question: 197

Your organization is developing a landing zone architecture with the following requirements:

There should be no communication between production and non-production environments.

Communication between applications within an environment may be necessary.

Network administrators should centrally manage all network resources, including subnets, routes, and firewall rules.

Each application should be billed separately.

Developers of an application within a project should have the autonomy to create their compute resources.

Up to 1000 applications are expected per environment.

You need to create a design that accommodates these requirements. What should you do?

A. Create a design where each project has its own VPC. Ensure all VPCs are connected by a Network Connectivity Center hub that is centrally managed by the network team.

B. Create a design that implements a single Shared VPC. Use VPC firewall rules with secure tags to enforce micro-

segmentation between environments.

C. Create a design that has one host project with a Shared VPC for the production environment, another host project with a Shared VPC for the non-production environment, and a service project that is associated with the corresponding host project for each initiative.

D. Create a design that has a Shared VPC for each project. Implement hierarchical firewall policies to apply micro-segmentation between VPCs.

Answer: C

Explanation:

This design allows you to separate production and non-production environments while using Shared VPCs. Each environment has its own Shared VPC, and a service project is associated with each, allowing for separate billing and autonomy for developers. Centralized management of network resources is handled by the host projects.

Reference: Google Cloud Shared VPC Documentation

Question: 198

Your organization wants to set up hybrid connectivity with VLAN attachments that terminate in a single Cloud Router with 99.9% uptime. You need to create a network design for your on-premises router that meets those requirements and has an active/passive configuration that uses only one VLAN attachment at a time. What should you do?

A. Create a design that uses a BGP multi-exit discriminator (MED) attribute to influence the egress path from Google Cloud to the on-premises environment.

B. Create a design that uses the as_path BGP attribute to influence the egress path from Google Cloud to the on-premises environment.

C. Create a design that uses an equal-cost multipath (ECMP) with flow-based hashing on your onpremises devices.

D. Create a design that uses the local_pref BGP attribute to influence the egress path from Google Cloud to the on-premises environment.

Answer: A

Explanation:

The BGP multi-exit discriminator (MED) attribute is used in BGP configurations to influence the choice of path in an active/passive setup by prioritizing one path over another for egress traffic. This is ideal for a design that uses only one VLAN attachment at a time.

Reference: Google Cloud Interconnect BGP Configuration

Question: 199

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 Monitoring section of the Google Cloud console, use the Dashboard section to select a default dashboard for VPN usage.

B. In Network Intelligence Center, check for the number of packet drops on the VPN.

C. 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.

D. In the Google Cloud console, use Monitoring Query Language to create a custom alert for bandwidth utilization.

Answer: D

Explanation:

Using Monitoring Query Language (MQL) to create a custom alert for bandwidth utilization gives you flexibility and precision in setting thresholds. This helps you quickly determine when VPN traffic exceeds the limits, allowing for timely decisions about adding more links or transitioning to a Dedicated Interconnect.

Reference: Google Cloud Monitoring Documentation

Question: 200

You are troubleshooting connectivity issues between Google Cloud and a public SaaS provider. Connectivity between the two environments is through the public internet. Your users are reporting intermittent connection errors when using TCP to connect; however, ICMP tests show no failures. According to users, errors occur around the same time every day. You want to troubleshoot and gather information by using Google Cloud tools that are most likely to provide insights into what is occurring within Google Cloud. What should you do?

A. Enable the Firewall Insights API. Set the deny rule insights observation period to one day. Review the insights to assure there are no firewall rules denying traffic.

B. Enable and review Cloud Logging on your Cloud NAT gateway. Look for logs with errors matching the destination IP address of the public SaaS provider.

C. Create a Connectivity Test by using TCP, the source IP address of your test VM, and the destination IP address of the public SaaS provider. Review the live data plane analysis and take the next steps based on the test results.

D. Enable and review Cloud Logging for Cloud Armor. Look for logs with errors matching the destination IP address of the public SaaS provider.

Answer: C

Explanation:

Creating a Connectivity Test using TCP in Network Intelligence Center allows you to simulate the connection to the public SaaS provider and receive real-time data plane analysis. This will help determine whether there are any issues with the network path for the specific TCP connection.

Reference: Google Cloud Connectivity Tests Documentation

Question: 201

Your organization has approximately 100 teams that need to manage their own environments. A central team must manage the network. You need to design a landing zone that provides separate projects for each team. You must also make sure the solution can scale. What should you do?

A. Configure VPC Network Peering, and peer one of the VPCs to the service project.

B. Configure a Shared VPC, and create a VPC network in the service project.

C. Configure a Shared VPC, and create a VPC network in the host project.

D. Configure Policy-based Routing for each team.

Answer: C

Explanation:

A Shared VPC allows the central networking team to manage the VPC network while individual teams can manage their resources in service projects. This solution provides scalability by allowing for multiple service projects under the same Shared VPC, and it allows the network team to maintain control over the network resources.

Reference: Google Cloud Shared VPC Architecture

Question: 202

Your organization has a hub and spoke architecture with VPC Network Peering, and hybrid connectivity is centralized at the hub. The Cloud Router in the hub VPC is advertising subnet routes, but the on-premises router does not appear to be receiving any subnet routes from the VPC spokes. You need to resolve this issue. What should you do?

A. Create custom routes at the Cloud Router in the hub to advertise the subnets of the VPC spokes. B. Create custom learned routes at the Cloud Router in the hub to advertise the subnets of the VPC spokes.

C. Create custom routes at the Cloud Router in the spokes to advertise the subnets of the VPC spokes.

D. Create a BGP route policy at the Cloud Router, and ensure the subnets of the VPC spokes are being announced towards the on-premises environment.

Answer: D

Explanation:

Creating a BGP route policy at the Cloud Router ensures that the subnets of the VPC spokes are properly advertised to the on-premises environment. This allows the on-premises router to receive and use those routes. Without the correct BGP policies, route advertisement may not happen as expected.

Reference: Google Cloud BGP Route Configuration

Question: 203

There are two established Partner Interconnect connections between your on-premises network and Google Cloud. The VPC that hosts the Partner Interconnect connections is named "vpc-a" and contains three VPC subnets across three regions, Compute Engine instances, and a GKE cluster. Your on-premises users would like to resolve records hosted in a Cloud DNS private zone following Google-recommended practices. You need to implement a solution that allows your on-premises users to resolve records that are hosted in Google Cloud. What should you do?

A. Associate the private zone to "vpc-a." Create an outbound forwarding policy and associate the policy to "vpc-a." Configure the on-premises DNS servers to forward queries for the private zone to the entry point addresses created when the policy was attached to "vpc-a."

B. Configure a DNS proxy service inside one of the GKE clusters. Expose the DNS proxy service in GKE as an internal load balancer. Configure the on-premises DNS servers to forward queries for the private zone to the IP address of the internal load balancer.

C. Use custom route advertisements to announce 169.254.169.254 via BGP to the on-premises environment. Configure the on-premises DNS servers to forward DNS requests to 169.254.169.254. D. Associate the private zone to "vpc-a." Create an inbound forwarding policy and associate the policy to "vpc-a." Configure the on-premises DNS servers to forward queries for the private zone to the entry point addresses created when the policy was attached to "vpc-a."

Answer: A

Explanation:

Associating the private zone to "vpc-a" and creating an outbound forwarding policy allows DNS queries to be forwarded from on-premises to Google Cloud DNS. The on-premises DNS servers will forward queries to the entry points created when the forwarding policy was applied to "vpc-a," enabling proper name resolution.

Reference: Google Cloud DNS Outbound Forwarding

Question: 204

You configured a single IPSec Cloud VPN tunnel for your organization to a third-party customer. You confirmed that the VPN tunnel is established; however, the BGP session status states that BGP is not configured. The customer has provided you with their BGP settings:

Local BGP address: 169.254.11.1/30

Local ASN: 64515

Peer BGP address: 169.254.11.2

Peer ASN: 64517

Base MED: 1000

MD5 Authentication: Disabled

You need to configure the local BGP session for this tunnel based on the settings provided by the customer. You already associated the Cloud Router with the Cloud VPN Tunnel. What settings should you use for the BGP session?

A. Peer ASN: 64517
Advertised Route Priority (MED): 100
Local BGP IP: 169.254.11.2
Peer BGP IP: 169.254.11.1
MD5 Authentication: Disabled

B. Peer ASN: 64515
Advertised Route Priority (MED): 100
Local BGP IP: 169.254.11.2
Peer BGP IP: 169.254.11.1
MD5 Authentication: Disabled

C. Peer ASN: 64515
Advertised Route Priority (MED): 1000
Local BGP IP: 169.254.11.2
Peer BGP IP: 169.254.11.1
MD5 Authentication: Enabled

D. Peer ASN: 64515
Advertised Route Priority (MED): 100
Local BGP IP: 169.254.11.1
Peer BGP IP: 169.254.11.2
MD5 Authentication: Disabled

Answer: A

Explanation:
The correct configuration requires setting the Peer ASN as 64517 (as this is the ASN of the third-party customer). The local and peer BGP IP addresses should also be set correctly based on the provided information, and MD5 authentication should be disabled. The route priority should be set to 100 to reflect standard behavior.

Reference: Google Cloud VPN BGP Configuration

Question: 205

You are configuring the firewall endpoints as part of the Cloud Next Generation Firewall (Cloud NGFW) intrusion prevention service in Google Cloud. You have configured a threat prevention security profile, and you now need to create an endpoint for traffic inspection. What should you do?

A. Create a Private Service Connect endpoint within the zone, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

B. Create a firewall endpoint within the region, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

C. Create a firewall endpoint within the zone, associate the endpoint to the VPC network, and use a firewall policy rule to apply the L7 inspection.

D. Attach the profile to the VPC network, create a firewall endpoint within the zone, and use a firewall policy rule to apply the L7 inspection.

Answer: C

Explanation:

To apply Layer 7 (L7) inspection for intrusion prevention, you must create a firewall endpoint within the zone where the traffic inspection is required. This endpoint is then associated with the VPC network, and a firewall policy rule is applied for the L7 inspection.

Reference: Google Cloud NGFW Setup

Question: 206

Your company's current network architecture has two VPCs that are connected by a dual-NIC instance that acts as a bump-in-the-wire firewall between the two VPCs. Flows between pairs of subnets across the two VPCs are working correctly. Suddenly, you receive an alert that none of the flows between the two VPCs are working anymore. You need to troubleshoot the problem. What should you do? (Choose 2 answers)

A. Verify that the dual-NIC instance has not been added to a backend service.

B. Verify that a public IP address has not been assigned to any network interface of the dual-NIC instance.

C. Use Cloud Logging to verify that there were no modifications to the VPC firewall rules or policies that were applied to the two network interfaces of the dual-NIC instance.

D. Verify that a VPC Service Controls perimeter has not been enabled for the project that contains the two VPCs and the dual-NIC instance.

E. Verify that the dual-NIC instance has the --can-ip-forward attribute enabled.

Answer: C, E

Explanation:

You should check Cloud Logging to see if any firewall rules or policies were modified, as these could block traffic between the VPCs. Additionally, the --can-ip-forward attribute must be enabled for the dual-NIC instance to allow forwarding traffic between the interfaces.

Reference: Google Cloud Dual-NIC Instance Configuration

Question: 207

Your team deployed two applications in GKE that are exposed through an external Application Load Balancer. When

queries are sent to www.mountkirkgames.com/sales and www.mountkirkgames.com/get-an-analysis, the correct pages are displayed. However, you have received complaints that www.mountkirkgames.com yields a 404 error. You need to resolve this error. What should you do?

A. Review the Ingress YAML file. Define the default backend. Reapply the YAML.

B. Review the Ingress YAML file. Add a new path rule for the * character that directs to the base service. Reapply the YAML.

C. Review the Service YAML file. Define a default backend. Reapply the YAML.

D. Review the Service YAML file. Add a new path rule for the * character that directs to the base service. Reapply the YAML.

Answer: A

Explanation:

The 404 error is occurring because there is no default backend defined for requests to the root URL. Defining the default backend in the Ingress YAML file ensures that requests to www.mountkirkgames.com are routed to the correct service.

Reference: Google Cloud GKE Ingress Setup

Question: 208

Your organization recently created a sandbox environment for a new cloud deployment. To have parity with the production environment, a pair of Compute Engine instances with multiple network interfaces (NICs) were deployed. These Compute Engine instances have a NIC in the Untrusted VPC (10.0.0.0/23) and a NIC in the Trusted VPC (10.128.0.0/9). A HA VPN tunnel has been established to the on-premises environment from the Untrusted VPC. Through this pair of VPN tunnels, the onpremises environment receives the route advertisements for the Untrusted and Trusted VPCs. In return, the on-premises environment advertises a number of CIDR ranges to the Untrusted VPC. However, when you tried to access one of the test services from the on-premises environment to the Trusted VPC, you received no response. You need to configure a highly available solution to enable the on-premises users to connect to the services in the Trusted VPC. What should you do?

A. Add both multi-NIC VMs to a new unmanaged instance group, named nva-uig.

Create an internal passthrough Network Load Balancer in the Untrusted VPC, named ilb-untrusted, with the nva-uig unmanaged instance group designated as the backend.

Create a custom static route in the Untrusted VPC for destination 10.123.0.0/9 and the next hop ilb- untrusted.

Create an internal passthrough Network Load Balancer in the Trusted VPC, named ilb-trusted, with the nva-uig unmanaged instance group designated as the backend.

Create a custom static route in the Trusted VPC for destination 0.0.0.0/0 and the next hop ilb-trusted.

B. Add both multi-NIC VMs to a new unmanaged instance group, named nva-uig.

Create an internal passthrough Network Load Balancer in the Untrusted VPC, named ilb-untrusted, with the nva-uig unmanaged instance group designated as the backend.

Create a custom static route in the Untrusted VPC for destination 10.128.0.0/9 and the next hop ilb- untrusted.

Create an internal passthrough Network Load Balancer in the Trusted VPC, named ilb-trusted, with the nva-uig unmanaged instance group designated as the backend.

Create a custom static route in the Trusted VPC for destination 10.0.0.0/23 and the next hop ilb- trusted.

C. Add both multi-NIC VMs to a new unmanaged instance group, named nva-uigO.

Create an internal passthrough Network Load Balancer in the Untrusted VPC, named ilb-untrusted, with the nva-uigO as backend.

Create a custom static route in the Untrusted VPC for destination 10.128.0.0/9 and the next hop ilb- untrusted.

Add both multi-NIC VMs to a new unmanaged instance group, named nva-uigl.

Create an internal passthrough Network Load Balancer in the Trusted VPC, named ilb-trusted, with the nva-uigl as backend.

Create a custom static route in the Trusted VPC for destination 0.0.0.0/0 and the next hop ilb-trusted.

D. Add both multi-NIC VMs to a new unmanaged instance group, named nva-uig.

Create two custom static routes in the Untrusted VPC for destination 10.128.0.0/9 and set each of the VMs' NIC as the next hop.

Create two custom static routes in the Trusted VPC for destination 10.0.0.0/23 and set each of the VMs' NIC as the next hop.

Answer: B

Explanation:

The solution requires creating internal passthrough load balancers for both VPCs, with custom static routes pointing to each load balancer. This ensures connectivity between the on-premises environment and the Trusted VPC via the Untrusted VPC.

Reference: Google Cloud Internal Load Balancer Setup

Question: 209

Your frontend application VMs and your backend database VMs are all deployed in the same VPC but across different subnets. Global network firewall policy rules are configured to allow traffic from the frontend VMs to the backend VMs. Based on a recent compliance requirement, this traffic must now be inspected by network virtual appliances (NVAs)

firewalls that are deployed in the same VPC. The NVAs are configured to be full network proxies and will source NAT-allowed traffic. You need to configure VPC routing to allow the NVAs to inspect the traffic between subnets. What should you do?

A. Place your NVAs behind an internal passthrough Network Load Balancer named ilb1. Add global network firewall policy rules to allow traffic through your NVAs. Create a custom static route with the destination IP range of the backend VM subnet, frontend instance tag, and the next hop of ilb1. Add a frontend network tag to your frontend VMs.

B. Create your NVA with multiple interfaces. Configure NIC0 for NVA in the backend subnet. Configure NIC1 for NVA in the frontend subnet. Place your NVAs behind an internal passthrough Network Load Balancer named ilb1. Add global network firewall policy rules to allow traffic through your NVAs. Create a custom static route with the destination IP range of the backend VM subnet, frontend instance tag, and the next hop of ilb1. Add a frontend network tag to your frontend VMs. C. Place your NVAs behind an internal passthrough Network Load Balancer named ilb1. Add the global network firewall policy rules to allow traffic through your NVAs. Create a policy-based route (PBR) with the source IP range of the backend VM subnet, destination IP range of the frontend VM subnet, and the next hop of ilb1. Scope the PBR to the VMs with the backend network tag. Add a backend network tag to your backend servers.

D. Place your NVAs behind an internal passthrough Network Load Balancer named ilb1. Add global network firewall policy rules to allow traffic through your NVAs. Create a policy-based route (PBR) with the source IP range of the frontend VM subnet, destination IP range of the backend VM subnet, and the next hop of ilb1. Scope the PBR to the VMs with the frontend network tag. Add a frontend network tag to your frontend servers.

Answer: D

Explanation:
The correct solution requires creating a policy-based route (PBR) to force the traffic from the frontend subnet to the backend subnet through the NVA. The PBR should be scoped to the frontend VMs, with the next hop being the passthrough load balancer (ilb1) behind which the NVAs reside. This ensures that all traffic is inspected by the NVA before reaching the backend.

Reference: Google Cloud Policy-based Routing Documentation

Question: 210

Your organization wants to deploy HA VPN over Cloud Interconnect to ensure encryption-in-transit over the Cloud Interconnect connections. You have created a Cloud Router and two encrypted VLAN attachments that have a 5 Gbps capacity and a BGP configuration. The BGP sessions are operational.
You need to complete the deployment of the HA VPN over Cloud Interconnect. What should you do?

A. Create an HA VPN gateway and associate the gateway with your two encrypted VLAN attachments. Configure the HA VPN Cloud Router, peer VPN gateway resources, and HA VPN tunnels. Use the same encrypted Cloud Router used for the Cloud Interconnect tier.

B. Enable MACsec for Cloud Interconnect on the VLAN attachments.

C. Enable MACsec on Partner Interconnect.

D. Create an HA VPN gateway and associate the gateway with your two encrypted VLAN attachments. Create a new dedicated HA VPN Cloud Router, peer VPN gateway resources, and HA VPN tunnels.

Answer: A

Explanation:

The correct approach is to create an HA VPN gateway and associate it with the encrypted VLAN attachments. The same Cloud Router used for BGP sessions with Cloud Interconnect can be used for the HA VPN. This configuration ensures encryption of the traffic passing over the Cloud Interconnect links.

Reference: Google Cloud HA VPN over Cloud Interconnect

Question: 211

Your organization recently re-architected your cloud environment to use Network Connectivity Center. However, an error occurred when you tried to add a new VPC named vpc-dev as a spoke. The error indicated that there was an issue with an existing spoke and the IP space of a VPC named vpc- pre-prod. You must complete the migration quickly and efficiently. What should you do?

A. Remove the conflicting VPC spoke for vpc-pre-prod from the set of VPC spokes in Network Connectivity Center. Add the VPC spoke for vpc-dev. Add the previously removed vpc-pre-prod as a VPC spoke.

B. Delete the VMs associated with the conflicting subnets, then delete the conflicting subnets in vpc- dev. Recreate the subnets with a new IP range and redeploy the previously deleted VMs in the new subnets. Add the VPC spoke for vpc-dev.

C. Exclude the conflicting IP range by using the --exclude-export-ranges flag when creating the VPC spoke for vpc-dev.

D. Exclude the conflicting IP range by using the --exclude-export-ranges flag in the hub when attaching the VPC spoke for vpc-dev.

Answer: A

Explanation:

The most efficient way to resolve the conflict is to temporarily remove the conflicting vpc-pre-prod spoke, add the vpc-dev spoke, and then re-add vpc-pre-prod. This ensures that the migration happens quickly without the need to change IP ranges or delete resources.

Reference: Google Network Connectivity Center Documentation

Question: 212

Your organization has resources in two different VPCs, each in different Google Cloud projects, and requires connectivity between the resources in the two VPCs. You have already determined that there is no IP address overlap; however, one VPC uses privately used public IP (PUPI) ranges. You would like to enable connectivity between these resources by using a lower cost and higher performance method. What should you do?

A. Create an HA VPN between the two VPCs that includes the PUPI ranges in the custom route advertisements of the Cloud Router. Create the necessary ingress VPC firewall rules that target the specific resources by using IP ranges as the source filter.

B. Create a VPC Network Peering connection between the two VPCs that allows the export and import of custom routes for public IP addresses. Create the necessary ingress VPC firewall rules that target the specific resources by using service accounts as the source filter.

C. Create a VPC Network Peering connection between the two VPCs that allows the export and import of subnet routes with public IP addresses. Create the necessary ingress VPC firewall rules that target the specific resources by using IP ranges as the source filter.

D. Create a VPC Network Peering connection between the two VPCs that allows the export and import of subnet routes with public IP addresses. Create the necessary ingress VPC firewall rules that target the specific resources by using network tags as the source filter.

Answer: C

Explanation:

VPC Network Peering is the most cost-effective and high-performance method for connecting two VPCs. Since one VPC uses privately used public IP (PUPI) ranges, you need to configure peering to allow the export and import of subnet routes with public IP addresses. Firewall rules can be used to control traffic between the resources.

Reference: Google Cloud VPC Peering Documentation

Question: 213

You have several VMs across multiple VPCs in your cloud environment that require access to internet endpoints. These VMs cannot have public IP addresses due to security policies, so you plan to use Cloud NAT to provide outbound internet access. Within your VPCs, you have several subnets in each region. You want to ensure that only specific subnets have access to the internet through Cloud NAT. You want to avoid any unintentional configuration issues caused by other administrators and align to Google-recommended practices. What should you do?

A. Deploy Cloud NAT in each VPC and configure a custom source range that includes the allowed subnets. Configure Cloud NAT rules to only permit the allowed subnets to egress through Cloud NAT. B. Create a firewall rule in each VPC at priority 500 that targets all instances in the network and denies egress to the internet (0.0.0.0/0). Create a firewall rule at priority 300 that targets all instances in the network, has a source filter that maps to the allowed subnets, and allows egress to the internet (0.0.0.0/0). Deploy Cloud NAT and configure all primary and secondary subnet source ranges.

C. Create a firewall rule in each VPC at priority 500 that targets all instances in the network and denies egress to the internet (0.0.0.0/0). Create a firewall rule at priority 300 that targets all instances in the network, has a source filter that maps to the allowed subnets, and allows egress to the internet (0.0.0.0/0). Deploy Cloud NAT and configure a custom source range that includes the allowed subnets.

D. Create a constraints/compute.restrictCloudNATUsage organizational policy constraint. Attach the constraint to a folder that contains the associated projects. Configure the allowedValues to only contain the subnets that should have internet access. Deploy Cloud NAT and select only the allowed subnets.

Answer: D

Explanation:

Using an organizational policy with the restrictCloudNATUsage constraint allows you to limit Cloud NAT usage to specific subnets, ensuring that only the necessary subnets can access the internet. This method aligns with Google-recommended practices for controlling Cloud NAT configurations across multiple VPCs and regions.

Reference: Google Cloud NAT Documentation

Question: 214

Your organization recently exposed a set of services through a global external Application Load Balancer. After conducting some testing, you observed that responses would intermittently yield a non-HTTP 200 response. You need to identify the error. What should you do? (Choose 2 answers)

A. Delete the load balancer and backend services. Create a new passthrough Network Load Balancer. Configure a failover group of VMs for the backend.

B. Access a VM in the VPC through SSH and try to access a backend VM directly. If the request is successful from the VM, increase the quantity of backends.

C. Enable and review the health check logs. Review the error responses in Cloud Logging.

D. Validate the health of the backend service. Enable logging for the backend service and identify the error response in Cloud Logging. Determine the cause of the error by reviewing the statusDetails log field.

E. Validate the health of the backend service. Enable logging on the load balancer and identify the error response in Cloud Logging. Determine the cause of the error by reviewing the statusDetails log field.

Answer: C, E

Explanation:
To troubleshoot the intermittent non-HTTP 200 responses, you should enable and review health check logs and log the backend service's responses in Cloud Logging. Reviewing the statusDetails field helps identify the cause of the error. Enabling logging on the load balancer and backend service provides visibility into the issue.

Reference: Google Cloud Load Balancer Logging

Question: 215

(You are managing the security configuration of your company's Google Cloud organization. The Operations team needs specific permissions on both a Google Kubernetes Engine (GKE) cluster and a Cloud SQL instance. Two predefined Identity and Access Management (IAM) roles exist that contain a subset of the permissions needed by the team. You need to configure the necessary IAM permissions for this team while following Google-recommended practices. What should you do?)

A. Grant the team the two predefined IAM roles.

B. Create a custom IAM role that combines the permissions from the two relevant predefined roles.

C. Create a custom IAM role that includes only the required permissions from the predefined roles.

D. Grant the team the IAM roles of Kubernetes Engine Admin and Cloud SQL Admin.

Answer: C

Explanation:

Granting more permissions than necessary violates the principle of least privilege, a fundamental security best practice. While option A grants the necessary permissions (as subsets exist in two predefined roles), it might also grant more permissions than the Operations team strictly requires for their tasks on GKE and Cloud SQL. Option D is too broad; 'Admin' roles grant extensive permissions
that likely exceed the specific needs.

Google Cloud's best practices strongly recommend adhering to the principle of least privilege. Creating a custom role allows you to precisely define the set of permissions the Operations team needs for their specific tasks on the GKE cluster and the Cloud SQL instance, without granting any unnecessary permissions. This minimizes the potential blast radius in case of accidental or malicious actions.

Google Cloud Documentation References:

IAM best practices: https://cloud.google.com/iam/docs/best-practices - This document explicitly recommends granting the minimum necessary permissions.

Creating and managing custom roles: https://cloud.google.com/iam/docs/creating-managing- custom-roles - This explains how to create roles tailored to specific job functions.

Understanding roles: https://cloud.google.com/iam/docs/understanding-roles - This outlines the concepts of predefined and custom roles and their use cases.

Question: 216

(Your digital media company stores a large number of video files on-premises. Each video file ranges from 100 MB to 100 GB. You are currently storing 150 TB of video data in your on-premises network, with no room for expansion. You need to migrate all infrequently accessed video files older than one year to Cloud Storage to ensure that on-premises storage remains available for new files. You must also minimize costs and control bandwidth usage. What should you do?)

A. Create a Cloud Storage bucket. Establish an Identity and Access Management (IAM) role with write permissions to the bucket. Use the gsutil tool to directly copy files over the network to Cloud Storage.

B. Set up a Cloud Interconnect connection between the on-premises network and Google Cloud. Establish a private endpoint for Filestore access. Transfer the data from the existing Network File System (NFS) to Filestore.

C. Use Transfer Appliance to request an appliance. Load the data locally, and ship the appliance back to Google for ingestion into Cloud Storage.

D. Use Storage Transfer Service to move the data from the selected on-premises file storage systems to a Cloud Storage bucket.

Answer: D

Explanation:

Let's analyze each option:

A . Using gsutil: While gsutil can transfer data to Cloud Storage, for 150 TB of infrequently accessed data, direct transfer over the network might be slow and consume significant bandwidth, potentially impacting other network operations. It also lacks built-in mechanisms for filtering files based on age.

B . Using Cloud Interconnect and Filestore: Cloud Interconnect provides a dedicated connection, but Filestore is a fully managed NFS service primarily designed for high-performance file sharing for applications running in Google Cloud. Migrating 150 TB of infrequently accessed data to Filestore would be cost-inefficient compared to Cloud Storage and doesn't directly address the requirement of moving older than one year files.

C . Using Transfer Appliance: Transfer Appliance is suitable for very large datasets (petabytes) or when network connectivity is poor or unreliable. While it addresses bandwidth concerns, it involves a physical appliance and might be an overkill for 150 TB of data, especially if network connectivity is reasonable.

D . Using Storage Transfer Service: Storage Transfer Service is specifically designed for moving large amounts of data between online storage systems, including on-premises file systems and Cloud Storage. It offers features like filtering by file age, scheduling transfers, and bandwidth control, directly addressing all the requirements of the question : migrating infrequently accessed files older than one year to Cloud Storage, minimizing costs (by using appropriate Cloud Storage classes for infrequent access), and controlling bandwidth usage.

Google Cloud Documentation References:

Storage Transfer Service Overview: https://cloud.google.com/storage-transfer- service/docs/overview - This page details the capabilities and use cases of Storage Transfer Service, including transferring from on-premises.

Storage Transfer Service for on-premises data: https://cloud.google.com/storage-transfer- service/docs/on-prem-overview - This specifically covers transferring data from on-premises file systems.

Cloud Storage Classes: https://cloud.google.com/storage/docs/storage-classes - Understanding the different storage classes (Standard, Nearline, Coldline, Archive) is crucial for cost optimization of infrequently accessed data. Storage Transfer Service can be configured to move data to a costeffective class like Nearline or Coldline.

Question: 217

(You are developing an internet of things (IoT) application that captures sensor data from multiple devices that have already been set up. You need to identify the global data storage product your company should use to store this data. You must ensure that the storage solution you choose meets your requirements of sub-millisecond latency. What should you do?)

A. Store the IoT data in Spanner. Use caches to speed up the process and avoid latencies.

B. Store the IoT data in Bigtable.

C. Capture IoT data in BigQuery datasets.

D. Store the IoT data in Cloud Storage. Implement caching by using Cloud CDN.

Answer: B

Explanation:

Let's evaluate each option based on the requirement of sub-millisecond latency for globally stored IoT data:

A . Spanner with Caching: While Spanner offers strong consistency and global scalability, the base latency might not consistently be sub-millisecond for all read/write operations globally. Introducing caching adds complexity and doesn't guarantee sub-millisecond latency for all initial reads or cache misses.

B . Bigtable: Bigtable is a highly scalable NoSQL database service designed for low-latency, high- throughput workloads. It excels at storing and retrieving large volumes of time-series data, which is typical for IoT sensor data. Its architecture is optimized for single-key lookups and scans, providing consistent sub-millisecond latency, making it a strong candidate for this use case.

C . BigQuery: BigQuery is a fully managed, serverless data warehouse designed for analytical queries on large datasets. While it's excellent for analyzing IoT data in batch, it's not optimized for the low- latency, high-throughput ingestion and retrieval required for real-time IoT applications with submillisecond latency needs.

D . Cloud Storage with Cloud CDN: Cloud Storage is object storage and is not designed for low-latency transactional workloads. Cloud CDN is a content delivery network that caches content closer to users for faster delivery, but it's not suitable for the primary storage of rapidly incoming IoT sensor data requiring sub-millisecond write latency.

Google Cloud Documentation References:

Cloud Bigtable Overview: https://cloud.google.com/bigtable/docs/overview - This document highlights Bigtable's suitability for low-latency and high-throughput applications, including IoT. It mentions its ability to handle massive amounts of data with consistent performance.

Spanner Overview: https://cloud.google.com/spanner/docs/overview - While Spanner offers low latency, Bigtable is generally preferred for extremely high-throughput, low-latency use cases like raw sensor data ingestion due to its optimized architecture for such workloads.

BigQuery Overview: https://cloud.google.com/bigquery/docs/introduction - This emphasizes BigQuery's analytical capabilities rather than low-latency operational workloads.

Cloud Storage Overview: https://cloud.google.com/storage/docs/overview - This describes Cloud Storage as object storage, not ideal for sub-millisecond latency reads and writes required for realtime IoT data.

Question: 218

(You are managing an application deployed on Cloud Run. The development team has released a new version of the application. You want to deploy and redirect traffic to this new version of the application. To ensure traffic to the new version of the application is served with no startup time, you want to ensure that there are two idle instances available for incoming traffic before adjusting the traffic flow. You also want to minimize administrative overhead. What should you do?)

A. Ensure the checkbox "Serve this revision immediately" is unchecked when deploying the new revision. Before changing the traffic rules, use a traffic simulation tool to send load to the new revision.

B. Configure service autoscaling and set the minimum number of instances to 2.

C. Configure revision autoscaling for the new revision and set the minimum number of instances to 2.

D. Configure revision autoscaling for the existing revision and set the minimum number of instances to 2.

Answer: C

Explanation:

Let's analyze each option to find the one that meets the requirements of no startup time for new traffic, two idle instances, and minimal administrative overhead:

A . Unchecking "Serve this revision immediately" and using a traffic simulation tool: Unchecking "Serve this revision immediately" does prevent the new revision from receiving traffic immediately.

However, manually using a traffic simulation tool adds administrative overhead. It also doesn't guarantee that two idle instances will be ready before traffic is shifted; you would need to monitor and adjust traffic manually based on the simulation.

B . Configuring service autoscaling and setting the minimum number of instances to 2: Service-level autoscaling applies to all revisions of the service. Setting the minimum instances at the service level would ensure at least two instances are running across all active revisions, not specifically for the new revision before traffic shift.

C . Configuring revision autoscaling for the new revision and setting the minimum number of instances to 2: This is the correct approach. By configuring revision autoscaling specifically for the new revision and setting the minimum number of instances to 2, Cloud Run will ensure that at least two instances of the new version are running and ready to serve traffic before you redirect any traffic to it. This eliminates startup latency when you do shift traffic. It also minimizes administrative overhead as Cloud Run manages the instance scaling based on this configuration.

D . Configuring revision autoscaling for the existing revision and setting the minimum number of instances to 2: This would ensure the existing version has at least two idle instances, which doesn't directly address the requirement of having idle instances ready for the new version before traffic redirection.

Google Cloud Documentation References:

Cloud Run Autoscaling: https://cloud.google.com/run/docs/configuring/min-instances - This document explains how to configure minimum and maximum instances for Cloud Run services and revisions. It clarifies that you can set minimum instances at the revision level to ensure instances are always ready.

Cloud Run Traffic Management: https://cloud.google.com/run/docs/managing/traffic - This describes how to deploy new revisions and gradually shift traffic between them. Combining minimum instances on the new revision with traffic splitting allows for zero-downtime deployments with prewarmed instances.

Question: 219

(You need to migrate multiple PostgreSQL databases from your on-premises data center to Google Cloud. You want to significantly improve the performance of your databases while minimizing changes to your data schema and application code. You expect to exceed 150 TB of data per geographical region. You want to follow Google-recommended practices and minimize your operational costs. What should you do?)

A. Migrate your data to AlloyDB.

B. Migrate your data to Spanner.

C. Migrate your data to Firebase.

D. Migrate your data to Bigtable.

Answer: A

Explanation:

Let's analyze each option based on the requirements: PostgreSQL compatibility, significant performance improvement, minimal schema/code changes, handling large data volumes, Google- recommended practices, and cost minimization:

A . Migrate your data to AlloyDB: AlloyDB for PostgreSQL is a fully managed, PostgreSQL-compatible database service that offers significant performance improvements over standard PostgreSQL due to its architectural optimizations. It is designed to handle large data volumes and minimizes the need for schema and application code changes as it's wire-compatible with PostgreSQL. This aligns well with the requirements for performance improvement, minimal changes, large data, and being a Google-recommended option for PostgreSQL workloads.

B . Migrate your data to Spanner: Spanner is a globally distributed, horizontally scalable database with strong consistency. While it offers excellent scalability and performance, it's not directly PostgreSQL-compatible. Migrating to Spanner would likely require significant schema and application code changes due to differences in data modeling and SQL dialect.

C . Migrate your data to Firebase: Firebase is a suite of mobile and web development tools, with its primary database offering being Firestore (a NoSQL document database) and Realtime Database. These are not PostgreSQL-compatible and

would require substantial changes to the data model and application code.

D . Migrate your data to Bigtable: Bigtable is a highly scalable NoSQL wide-column store. It's not compatible with PostgreSQL and requires a completely different data model and application logic.

Therefore, AlloyDB is the most suitable option as it provides PostgreSQL compatibility for minimal migration effort, significant performance improvements, scalability for large data volumes, and is a recommended Google Cloud database service for PostgreSQL workloads.

Google Cloud Documentation References:

AlloyDB for PostgreSQL Overview: https://cloud.google.com/alloydb/docs/overview - This document highlights AlloyDB's PostgreSQL compatibility, performance benefits, scalability, and suitability for

migrating existing PostgreSQL workloads.

Spanner Overview: https://cloud.google.com/spanner/docs/overview - This emphasizes Spanner's unique features and differences from traditional relational databases like PostgreSQL.

Firebase Documentation: https://firebase.google.com/docs - This outlines the features of Firebase, including Firestore and Realtime Database, highlighting their NoSQL nature and incompatibility with PostgreSQL.

Cloud Bigtable Overview: https://cloud.google.com/bigtable/docs/overview - This describes Bigtable as a NoSQL database, emphasizing its differences from relational databases like PostgreSQL.

Question: 220

(You are deploying an application to Google Kubernetes Engine (GKE). The application needs to make API calls to a private Cloud Storage bucket. You need to configure your application Pods to authenticate to the Cloud Storage API, but your organization policy prevents the usage of service account keys. You want to follow Google-recommended practices. What should you do?)

A. Create the GKE cluster and deploy the application. Request a security exception to create a Google service account key. Set the constraints/iam.serviceAccountKeyExpiryHours organization policy to 8 hours.

B. Create the GKE cluster and deploy the application. Request a security exception to create a Google service account key. Set the constraints/iam.serviceAccountKeyExpiryHours organization policy to 24 hours.

C. Create the GKE cluster with Workload Identity Federation. Configure the default node service account to access the bucket. Deploy the application into the cluster so the application can use the node service account permissions. Use Identity and Access Management (IAM) to grant the service account access to the bucket.

D. Create the GKE cluster with Workload Identity Federation. Create a Google service account and a Kubernetes ServiceAccount, and configure both service accounts to use Workload Identity

Federation. Attach the Kubernetes ServiceAccount to the application Pods and configure the Google service account to access the bucket with Identity and Access Management (IAM).

Answer: D

Explanation:

Create a Google Service Account: You create a dedicated Google service account specifically for your application's interaction with the private Cloud Storage bucket. This allows you to grant precise IAM

permissions to this service account on the bucket (e.g., roles/storage.objectViewer or roles/storage.objectCreator).

Create a Kubernetes ServiceAccount: You create a Kubernetes ServiceAccount within your GKE cluster. This is the identity that your application Pods will assume within the cluster.

Configure Workload Identity Federation: You establish a trust relationship between the Kubernetes ServiceAccount and the Google service account using Workload Identity Federation. This involves configuring IAM policies that allow the Kubernetes ServiceAccount to impersonate the Google service account.

Annotate Pods with the Kubernetes ServiceAccount: You associate the created Kubernetes ServiceAccount with your application Pods. When the application in these Pods makes a call to the Cloud Storage API, the Workload Identity agent running on the GKE nodes automatically exchanges the Kubernetes ServiceAccount token for a short-lived Google Cloud access token for the associated Google service account.

This approach offers several security advantages and aligns with Google's recommended practices:

Principle of Least Privilege: The Google service account is granted only the necessary permissions to access the specific Cloud Storage bucket.

No Service Account Keys to Manage: You avoid the security risks associated with creating, storing, and rotating service account keys.

Auditable Authentication: All API calls are attributed to the specific Google service account, providing better auditability.

Simplified Management: Workload Identity Federation automates the credential management process for your application.

Google Cloud Documentation References:

Workload Identity: https://cloud.google.com/kubernetes-engine/docs/how-to/workload-identity 1 - This is the primary documentation explaining how to use Workload Identity to allow applications in GKE to access Google Cloud services securely without using service account keys.

Question: 221

Your company uses Compute Engine instances that are exposed to the public internet. Each compute instance has a single network interface with a single public IP address. You need to block any connection attempt that originates from internet clients with IP addresses that belong to the bgp_asn_toblock BGP ASN. What should you do?

A. Q Create a new Cloud Armor edge security policy, and use the —network-src-asns parameter.

B. Q Create a new Cloud Armor network edge security policy, and use the —network-src-asns parameter.

C. O Create a new firewall policy ingress rule, and use the —network-src-asns parameter.

D. Q Create a new Cloud Armor backend security policy, and use the —network-src-asns parameter.

Answer: B

Explanation:

To block traffic based on BGP ASN, you need to use Cloud Armor network edge security policies. Backend security policies protect backend services (like those behind an external HTTP(S) Load Balancer), but for traffic directly to Compute Engine instances with public IPs, you need a network edge security policy. Firewall policies operate at the VPC level and do not have the capability to filter based on BGP ASN. The —network-src-asns parameter is specifically used with Cloud Armor network edge security policies to filter traffic based on the source ASN.

Exact Extract:

"Cloud Armor network edge security policies can protect external IP addresses of virtual machine (VM) instances and load balancers. They support rules that allow or deny traffic based on various criteria, including source autonomous system numbers (ASNs) using the --network-src-asns parameter."

"Network edge security policies are designed for traffic destined for external IP addresses of Google Cloud resources that are not behind an external HTTP(S) load balancer, such as Compute Engine instances with public IP addresses."Reference: Google Cloud Armor Documentation - Network edge security policies overview

Question: 222

You've received reports of latency between two application VMs which run in two different regions of your Google Cloud VPC network. There is typically about 8ms of latency, but now there is approximately 17ms of latency. You've eliminated application issues as a root cause, and you suspect that the latency may be a Google Cloud platform issue. You need to confirm this hypothesis using Google-recommended practices. What should you do?

A. Q Use Network Intelligence Center Performance Dashboard to view the inter-region packet loss for your VPC.

B. O Install and run tcpdump on both instances, and calculate the latency between the two instances by comparing the timestamps in the packet captures.

C. Q Use Network Intelligence Center Performance Dashboard to view inter-region latency for the Google Cloud network.

D. Q Use Network Intelligence Center Connectivity Tests, run a test between the two VMs, and review the inter-region latency in the test results.

Answer: C

Explanation:

When diagnosing latency issues that are suspected to be a Google Cloud platform issue, the Network

Intelligence Center Performance Dashboard is the recommended tool. It provides visibility into network performance metrics across Google Cloud, including inter-region latency, which can help confirm if the increased latency is due to a platform-wide issue rather than specific to your VPC or application. While Connectivity Tests are useful for verifying reachability and basic network configurations, and tcpdump is for in-instance packet analysis, neither provides the broad, platformlevel visibility needed to assess general Google Cloud network performance trends and baselines. Exact Extract: "The Performance Dashboard in Network Intelligence Center helps you visualize network performance metrics for your Google Cloud network, including inter-region latency and packet loss. It provides insights into the health of the Google Cloud network and can help you identify if performance degradations are due to the underlying platform." "The Performance Dashboard provides an aggregated view of network performance across Google Cloud, allowing you to compare your current latency with historical trends and Google Cloud's overall network performance."Reference: Google Cloud Network Intelligence Center Documentation - Performance Dashboard

Question: 223

You are configuring an Application Load Balancer. The backend resides in your on-premises data center and is connected by Dedicated Interconnect. You need to ensure the load balancer can reference these on-premises resources. You do not want the traffic to traverse the internet at all. What should you do?

A. Q Configure a Private Service Connect network endpoint group (NEG) as a backend service as part of the load balancer. Ensure firewalls are opened for the client source IPs.
B. Q Configure a zonal network endpoint group (NEG) as a backend service as part of the load balancer. Ensure firewalls are opened for the client source IPs.
C. Q Configure an internet network endpoint group (NEG) as a backend service as part of the load balancer. Ensure firewalls are opened for the proxy-only subnet.
D. Q Configure a hybrid network endpoint group (NEG) as a backend service as part of the load balancer. Ensure firewalls are opened for the proxy-only subnet.

Answer: D

Explanation:

To connect an Application Load Balancer to on-premises resources via Dedicated Interconnect without traversing the internet, you must use a hybrid network endpoint group (NEG). Hybrid NEGs are specifically designed for connecting Google Cloud load balancers to on-premises or other cloud environments via hybrid connectivity solutions like Cloud Interconnect or Cloud VPN. The proxy-only subnet is essential for the load balancer's proxies to communicate with the backends.

Exact Extract:
"A hybrid NEG enables you to use an external HTTP(S) Load Balancer with backends that are outside of Google Cloud, such as your on-premises data centers or other cloud providers. This is typically used in conjunction with hybrid connectivity solutions like Cloud Interconnect or Cloud VPN." "When using an external HTTP(S) Load Balancer with hybrid NEGs, you must configure a proxy-only subnet in the region where the load balancer is deployed. This subnet is used by the load balancer's proxies to reach your backends."Reference: Google Cloud Load Balancing Documentation - Hybrid

NEG overview

Question: 224

You are deploying GKE clusters in your organization's Google Cloud environment. The pods in these clusters need to egress directly to the internet for a majority of their communications. You need to deploy the clusters and associated networking features using the most cost-efficient approach, and following Google-recommended practices. What should you do?

A. Q Deploy the GKE cluster with public cluster nodes. Do not deploy Cloud NAT or Secure Web Proxy for the cluster.
B. Q Deploy the GKE cluster with private cluster nodes. Deploy Secure Web Proxy, and configure the pods to use Secure Web Proxy as an HTTP(S) proxy.
C. Q Deploy the GKE cluster with public cluster nodes. Deploy Secure Web Proxy, and configure the pods to use Secure Web Proxy as an HTTP(S) proxy.
D. Q Deploy the GKE cluster with private cluster nodes. Deploy Cloud NAT for the primary subnet of the cluster.

Answer: A

Explanation:

For GKE pods that need to egress directly to the internet for most of their communications, the most cost-efficient and straightforward approach is to deploy a GKE cluster with public cluster nodes.
Public nodes have external IP addresses, allowing pods to directly reach the internet. This eliminates the need for additional services like Cloud NAT or Secure Web Proxy for outbound internet access, which would incur extra costs and management overhead.
Exact Extract:
"Public clusters have nodes with external IP addresses, allowing them to directly initiate connections to the internet. This is the simplest configuration for clusters that require direct internet egress for their workloads."
"When using public clusters, Cloud NAT is not required for outbound internet connectivity from the nodes or pods, as they can use their external IP addresses. This can reduce operational overhead and cost compared to private clusters that need NAT."Reference: Google Kubernetes Engine Documentation - Cluster network configuration, Public clusters vs Private clusters

Question: 225

E. u manage two VPCs: VPC1 and VPC2, each with resources spread across two regions. You connected the VPCs with HA VPN in both regions to ensure redundancy. You've observed that when one VPN gateway fails, workloads that are located within the same region but different VPCs lose communication with each other. After further debugging, you notice that VMs in VPC2 receive traffic but their replies never get to the VMs in VPC1. You need to quickly fix the issue. What should you do?

A. Q Enable regional dynamic routing mode in VPC2.
B. Q Enable global dynamic routing mode in VPC1.
C. Q Enable global dynamic routing mode in VPC2.
D. Q Enable regional dynamic routing mode in VPC1.

Answer: C

Explanation:

The problem description indicates that VMs in VPC2 receive traffic but their replies don't reach VPC1, especially when a VPN gateway fails. This strongly suggests an asymmetric routing issue, where VPC2's routing table might not be aware of all necessary routes to send return traffic to VPC1, particularly in a multi-region setup with failover. By default, VPC networks are in regional dynamic routing mode, meaning they only learn routes from Cloud Routers in the same region. To ensure that routes learned from one region (where the active VPN tunnel might be) are available globally across the VPC, you need to enable global dynamic routing mode in the VPC that is experiencing the return traffic issue (VPC2 in this case). This allows VPC2 to learn and apply routes from Cloud Routers in all regions, ensuring that even if a VPN tunnel fails in one region, the routes learned from the active tunnel in another region are still available for return traffic.

Exact Extract:

"A VPC network's dynamic routing mode controls whether routes learned by Cloud Routers in one region are available to VMs in other regions. By default, VPC networks are in regional dynamic routing mode, which means Cloud Routers in a region only advertise routes to and learn routes from other Cloud Routers in the same region. This can lead to asymmetric routing issues in multi-region deployments."

"To ensure that routes learned from Cloud Routers are propagated to all regions within a VPC network, you must set the dynamic routing mode to global."Reference: Google Cloud VPC Documentation - Dynamic routing mode

Question: 226

Your company's on-premises office is connected to Google Cloud using HA VPN. The security team will soon enable VPC Service Controls. You need to create a plan with minimal configuration adjustments, so clients at the office will still be able to privately call the Google APIs and be protected by VPC Service Controls. What should you do?

A. Create a design with a DNS configuration that resolves the Google APIs to 199.36.153.4/30; advertise 199.36.153.4/30 from Google Cloud to the onpremises routers; add an access level to authorize the on-premises network to access the APIs.

B. Create a design with a DNS configuration that resolves the Google APIs to 199.36.153.8/30; advertise 199.36.153.8/30 from Google Cloud to the onpremises routers.

C. Create a design with a DNS configuration that resolves the Google APIs to 199.36.153.8/30; advertise 199.36.153.8/30 from Google Cloud to the onpremise routers: add an access level to authorize the on-premises network to access the APIs.

D. Create a design with a DNS configuration that resolves the Google APIs to 199.36.153.4/30; advertise 199.36.153.4/30 from Google Cloud to the onpremises routers.

Answer: C

Explanation:

When integrating on-premises networks with VPC Service Controls for private access to Google APIs, the recommended approach involves using Private Google Access for Hybrid Connectivity and configuring DNS resolution to the restricted.googleapis.com domain. This domain resolves to the 199.36.153.8/30 IP address range. It's crucial to advertise this range from Google Cloud to your onpremises routers so that on-premises clients can route traffic to the Google APIs privately. Additionally, to allow your on-premises network to access the APIs within the VPC Service Controls perimeter,

you must define an access level that includes the IP address range of your on-premises network.

Exact Extract:

"To enable private access to Google APIs and services from on-premises networks protected by a VPC Service Controls perimeter, you must configure Private Google Access for Hybrid Connectivity." "For on-premises hosts, configure your DNS to resolve *.googleapis.com to restricted.googleapis.com. The restricted.googleapis.com domain resolves to the IP address range 199.36.153.8/30."

"You must advertise the 199.36.153.8/30 range from your Cloud Routers to your on-premises routers through BGP."

"To allow on-premises networks to access protected resources within a VPC Service Controls perimeter, you must define an access level that includes the on-premises network's IP ranges. This access level is then applied to the service perimeter."Reference: Google Cloud VPC Service Controls Documentation - Private connectivity to Google APIs, Private Google Access for Hybrid Connectivity

Question: 227

You are troubleshooting an issue where your organization's Cloud HA VPN is disconnected from your on-premises router for approximately 10 seconds before reestablishing the tunnel. The issue regularly occurs every few hours. You notice that the HA VPN logs show an entry of Received SA_DELETE when this issue occurs. You need to resolve this issue and prevent future VPN downtime from impacting your production applications. What should you do?

A. Q Update the pre-shared key (PSK) of the on-premises router's VPN tunnel configuration to match the PSK of the Cloud HA VPN.

B. Q Update the on-premises router's BGP router ID to reflect the link-local IP peer address assigned by Cloud Router.

C. Q Update the on-premises router's Phase 1 and Phase 2 lifetime IKE parameters to match the values in the Cloud HA VPN documentation.

D. Q Update the on-premises router's Diffie-Hellman groups and cipher proposal list to match the values in the Cloud HA VPN documentation.

Answer: C

Explanation:

The SA_DELETE message in VPN logs, especially when followed by a re-establishment, is a strong indicator that the Security Association (SA) lifetimes for Phase 1 (IKE SA) or Phase 2 (IPsec SA) are mismatched between the Google Cloud HA VPN and the on-premises router. When one side's SA expires, it sends an SA_DELETE message to the peer, which then triggers a rekeying process. If the lifetimes don't match, one side might prematurely terminate the SA, leading to brief disconnections.

Ensuring that the IKE parameters (specifically Phase 1 and Phase 2 lifetimes) match the recommended values in the Cloud HA VPN documentation is crucial for stable tunnel operation. Exact Extract:

"If your VPN tunnel frequently disconnects and reconnects, and you see SA_DELETE messages in your logs, it often indicates a mismatch in the Phase 1 (IKE SA) and Phase 2 (IPsec SA) lifetimes configured on your on-premises VPN gateway and the Cloud VPN gateway."

"For optimal stability, ensure that the IKE lifetime and ESP/IPsec lifetime parameters on your onpremises VPN device exactly match the recommended values provided in the Cloud VPN documentation. Mismatched lifetimes can cause tunnels to rekey prematurely or fall out of sync, leading to temporary disconnections."Reference: Google Cloud VPN Documentation - Troubleshooting VPN issues, Recommended IKE and IPsec settings

Question: 228

Your organization is migrating workloads from AWS to Google Cloud. Because a particularly critical workload will take longer to migrate, you need to set up Google Cloud CDN and point it to the existing application at AWS. What should you do?

A. Create a hybrid NEG that points to the existing IP of the application.
- Map the NEG to a passthrough Network Load Balancer as a target pool.
- Enable Cloud CDN on the target pool.

B. Create an internet NEG that points to the existing FQDN of the application.
- Map the NEG to an Application Load Balancer as a backend service.
- Enable Cloud CDN on the backend service.

C. Create a hybrid NEG that points to the existing IP of the application.
- Map the NEG to an Application Load Balancer as a backend service.
- Enable Cloud CDN on the backend service.

D. Create an internet NEG that points to the existing FQDN of the application.
- Map the NEG to a passthrough Network Load Balancer as a backend service.
- Enable Cloud CDN on the backend service.

Answer: B

Explanation:

To configure Cloud CDN for an application hosted outside of Google Cloud (e.g., in AWS), you need to use an internet network endpoint group (NEG). An internet NEG allows you to point to external endpoints using their FQDN or IP address. Cloud CDN works with external HTTP(S) Load Balancers, and you enable CDN on the backend service associated with the load balancer. A Network Load Balancer (passthrough) does not support Cloud CDN.

Exact Extract:

"To enable Cloud CDN for content hosted outside of Google Cloud, you must use an external HTTP(S) Load Balancer with an internet network endpoint group (NEG)."

"An internet NEG specifies one or more external endpoints that can be reached by an external HTTP(S) Load Balancer. You can specify endpoints using an IP address and port, or a fully qualified domain name (FQDN) and port."

"Cloud CDN is enabled on the backend service of an external HTTP(S) Load Balancer."Reference:

Google Cloud CDN Documentation - Caching external content, Internet NEGs overview

Question: 229

Your company acquired a new division. The new division's network team requires complete control over their networking infrastructure. You need to extend your existing Google Cloud network infrastructure, that consists of a single VPC, to allow workloads from all divisions to communicate with each other. You want to avoid incurring extra costs and granting unnecessary permissions to the new division's networking team. What should you do?

A. Q • Create a new project for the new division's network team.
- Create a new VPC within the new project.
- Establish a VPC peering between your existing VPC and the new division's VPC.
- Grant roles/compute. networkAdmin on the newly created project to the new division's network team group.

B. O * Create a new project for the new division's network team.

- Create a new VPC within the new project.
- Establish a VPC peering between your existing VPC and the new division's VPC.
- Create a new subnet dedicated to the new division's workloads.
- Grant roles/compute .networkuser on the new project to the new division's network team group.

C. O • Create a new project for the new division's network team.

- Create a new VPC within the new project.
- Establish a VPN connection between your existing VPC and the new division's VPC.
- Grant roles/compute .networkAdmin on the newly created project to the new division's network team group.

D. Q • Ensure that the project hosting the existing network infrastructure is enabled as a host project.

- Create a new subnet dedicated to the new division's workloads in the existing VPC.
- Grant roles/compute. networkuser on the newly created subnet to the new division's network team group.

Answer: A

Explanation:

The requirement for the new division's network team to have "complete control over their networking infrastructure" while allowing communication between divisions and avoiding unnecessary permissions points directly to VPC Network Peering. This approach allows each division to manage its own VPC independently (in its own project), provides full control to the new division's network team within their project, and enables secure, private communication between the VPCs without traversing the public internet. Granting roles/compute.networkAdmin on their newly created project ensures they have the necessary control over their dedicated VPC. Using Shared VPC (option D) would centralize network administration under your existing project, which goes against the requirement of the new division having "complete control." VPN (option C) would incur additional costs and introduce more complexity than VPC peering for intra-Google Cloud connectivity. Option B is flawed because creating a subnet in the new VPC isn't directly relevant to granting permissions on the new project for VPC peering setup, and networkuser role on the new project alone wouldn't give complete network control.

Exact Extract:

"VPC Network Peering allows you to connect two VPC networks so that resources in each network can communicate with each other using internal IP addresses. Traffic stays within Google's network." "Each side of a VPC Network Peering connection is configured independently. This means that each network administrator retains full control over their own network, including routes, firewalls, and network services."

"VPC Network Peering is ideal for scenarios where different organizations or divisions want to maintain separate network administrative domains while still allowing their resources to communicate privately."Reference: Google Cloud VPC Network Peering Documentation - Overview, Use cases

Question: 230

You are responsible for connectivity between AWS. Google Cloud, and an on-premises data center. Soon, the application team will deploy a data replication service that will move approximately 900 TB of data between Google Cloud and AWS daily. This data is sensitive and must be encrypted in transit. Your data center already has connections to both AWS and Google Cloud through 10 Gbps circuits. You need to configure additional connectivity between these environments and ensure the highest performance and lowest latency to meet business requirements. You also need to keep the existing connectivity topology to the on-premises data center the same. What should you do?

A.(Q) • Deploy Cross-Cloud Interconnect connections between AWS and Google Cloud with 100 Gbps circuits.

- Create VLAN attachments in your VPC, configuring IPsec encryption on both sides of the connection.
- Use Cloud Router and BGP to exchange dynamic routes between AWS and Google Cloud.

B. Q • Deploy Dedicated Interconnect connections between Google Cloud and your on-premises data center with 100

Gbps circuits from Google Cloud to your on-premises data center.

- Deploy an AWS Direct Connect 100 Gbps circuit from AWS to your on-premises data center.
- Create VLAN attachments in your VPC, configuring IPsec encryption on both sides of the connection.
- Use Cloud Router and BGP to exchange dynamic routes between AWS. Google Cloud, and the on premises data center.
- Remove the obsolete 10 Gbps circuits on Google Cloud and AWS.

C. Q • Deploy Dedicated Interconnect connections between Google Cloud and your on-premises data center with 100 Gbps circuits.

- Deploy an AWS Direct Connect 100 Gbps circuit from AWS to your on-premises data center as well. • Create VLAN attachments in your VPC.
- Use Cloud Router and BGP to exchange dynamic routes between AWS, Google Cloud, and the on premises data center.
- Remove the obsolete 10 Gbps circuits on Google Cloud and AWS.

D. Q • Deploy Cross-Cloud Interconnect connections between AWS and Google Cloud with 100 Gbps circuits.

- Enable MACsec for Cloud Interconnect on the circuits, and create VLAN attachments in your VPC.
- Use Cloud Router and BGP to exchange dynamic routes between AWS and Google Cloud.

Answer: A

Explanation:

The core requirement is to move a massive amount of sensitive data (900 TB daily) directly between Google Cloud and AWS with highest performance, lowest latency, and in-transit encryption, while maintaining existing on-premises connectivity.

Option A directly addresses this by recommending Cross-Cloud Interconnect with 100 Gbps circuits between AWS and Google Cloud. Cross-Cloud Interconnect is designed for high-throughput, low- latency connectivity between different cloud providers. The crucial part for sensitive data and encryption is "configuring IPsec encryption on both sides of the connection," as Cross-Cloud Interconnect itself provides a private path but not inherent encryption. Cloud Router and BGP are essential for dynamic route exchange. This option focuses on the direct cloud-to-cloud path for the high volume data transfer.

Options B and C involve upgrading the existing connections to the on-premises data center and routing all traffic through it. While this could work, it adds an unnecessary hop and likely higher latency for direct cloud-to-cloud traffic, making it less optimal for "highest performance and lowest latency" between clouds. Additionally, removing existing 10Gbps circuits is not necessary and might impact the existing topology if not done carefully.

Option D suggests MACsec, which provides Layer 2 encryption. While good for physical security, for data replication services with sensitive data, IPsec (Layer 3 encryption) is more commonly used and flexible for end-to-end encryption across a routed network, and is typically preferred for data integrity and confidentiality over an IP network. Also, MACsec requires specific hardware support and is typically implemented at the interconnect termination points, not necessarily end-to-end for an application. Given the sensitive nature of the data and the large volume, IPsec provides the necessary transport-level encryption.

Exact Extract:

"Cross-Cloud Interconnect enables direct connectivity between your Google Cloud VPC networks and other cloud provider networks. It provides high-bandwidth, low-latency connections, ideal for large- scale data transfers between clouds."

"For sensitive data, you can implement IPsec VPN tunnels over Cross-Cloud Interconnect connections to provide encryption in transit. This ensures data confidentiality and integrity over the dedicated interconnect."

"Cloud Router dynamically exchanges routes between your Google Cloud VPC network and your other cloud network over the Cross-Cloud Interconnect connection using BGP."Reference: Google Cloud Cross-Cloud Interconnect Documentation - Overview, Encryption options for hybrid connectivity

Question: 231

You are creating a new GKE standard cluster. You need to configure the cluster to ensure that pods can reach other VMs in Google Cloud in the 192.168.0.0/24 subnet using the source IP of the GKE nodes. What should you do?

A. Q Set a GKE pod IP address range that fits in 10.0.0.0/8. Configure the —disable-def ault-snat. flag.
B. Q Set a GKE pod IP address range that fits in 10.0.0.0/8. Do not configure the —disable-def ault- snat flag.
C. Q Set a GKE pod IP address range that does not fit in 10.0.0.0/8. Do not configure the —disable- default-snat flag.
D. Q Set a GKE pod IP address range that does not fit in 10.0.0.0/8. Configure the —disable-default- snat flag.

Answer: A

Explanation:

By default, GKE uses SNAT (Source Network Address Translation) for pod egress traffic to destinations outside the cluster's IP ranges but within RFC 1918 private IP ranges (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16). This means that traffic from pods leaving the cluster for these private IP destinations will have their source IP address translated to the node's IP address.

To ensure pods can reach VMs in the 192.168.0.0/24 subnet using the source IP of the GKE nodes, you want the default SNAT behavior to apply to this destination. The default SNAT rule applies when the destination is an RFC 1918 address and the source is a pod IP that is not within the same RFC 1918 range as the destination (e.g., if your pods are in a 10.x.x.x range and the destination is 192.168.x.x).

Therefore, you should:

Set a GKE pod IP address range that fits in 10.0.0.0/8: This ensures that the pod IPs are within an RFC 1918 range different from 192.168.0.0/24.

Do NOT configure the --disable-default-snat flag: If you disable default SNAT, pods would use their own IP addresses as source IPs, which might not be routable to the 192.168.0.0/24 subnet unless specific routes are configured. The goal is to use the node's IP.

The combination of having pod IPs in a different RFC 1918 range and not disabling default SNAT ensures that GKE performs SNAT, making the node's IP the source for traffic destined for the 192.168.0.0/24 subnet.

Exact Extract:

"By default, GKE performs SNAT (Source Network Address Translation) for egress traffic from pods to destinations outside the cluster's IP address ranges but within the private IP address ranges defined in RFC 1918 (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16). When SNAT occurs, the source IP address of the egress packets is the node's IP address instead of the pod's IP address."

"The --disable-default-snat flag, when used, disables this default SNAT behavior. If you want traffic to use the node's IP as the source when reaching internal RFC 1918 destinations, do not set this flag."Reference: Google Kubernetes Engine Documentation - IP masquerade agent, Private IP addresses for GKE Pods and Services

Question: 232

You are implementing a VPC architecture for your organization by using a Network Connectivity Center hub and spoke topology:

- There is one Network Connectivity Center hybrid spoke to receive on-premises routes.
- There is one VPC spoke that needs to be added as a Network Connectivity Center spoke.

Your organization has limited routable IP space fortheir cloud environment (192.168.0.0/20). The Network Connectivity Center spoke VPC is connected to on-premises with a Cloud Interconnect connection in the us-east4 region. The on-premises IP range is 172.16.0.0/16. You need to reach onpremises resources from multiple Google Cloud regions (us-

westl, europe-centrall, and asia- southeastl) and minimize the IP addresses being used. What should you do?

A. O 1. Configure a Private NAT gateway and NAT subnet in us-westl (192.168.1.0/24), europe- centrall (192.168.2.0/24) and asia-southeastl (192.168.3.0/24).
2. Add the VPC as a spoke and configure an export include policy to advertise only 192.168.1.0/24, 192.168.2.0/24, and 192.168.3.0/24 to the hub.
3. Enable global dynamic routing to allow resources in us-westl, us-centrall and asia-southeastl to reach the on-premises location through us-east4.

B. Q 1. Configure a Private NAT gateway instance in us-westl (192.168.1.0/24), europe-centrall (192.168.2.0/24), and asia-southeastl (192.168.3.0/24).
2. Add the VPC as a spoke and configure an export exclude policy on the VPC spoke to advertise only the NAT subnets 192.168.1.0/24, 192.168.2.0/24, and 192.168.3.0/24 to the hub.
3. Enable global dynamic routing to allow resources in us-westl, us-centrall, and asia-southeastl to reach the on-premises location through us-east4.

C. Q 1. Configure a Private NAT gateway instance in us-east4 (192.168.1.0/24).
2. Add the VPC as a spoke and configure an export include policy on the VPC spoke to advertise 192.168.1.0/24 to the hub.
3. Enable global dynamic routing to allow resources in us-westl, us-centrall and asia-southeast l to reach the on-premises location through us-east 4.

D. O 1- Configure a Private NAT gateway instance in us-westl (172.16.1.0/24), europe-centrall (172.16.2.0/24), and asia-southeastl (172.16.3.0/24).
2. Add the VPC as a spoke and configure an export include policy on the VPC spoke to advertise only the NAT subnets 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24 to the hub.
3. Enable global dynamic to allow resources in us-westl, us-centrall, and asia-southeastl to reach the on-premises location through us-east4.

Answer: C

Explanation:

The key requirements are: limited IP space (192.168.0.0/20), reaching on-premises (172.16.0.0/16) from multiple Google Cloud regions (us-west1, europe-central1, asia-southeast1), and minimizing IP addresses used. The Cloud Interconnect connection to on-premises is in us-east4.

Minimize IP addresses and centralized NAT: Since all traffic to on-premises will traverse the Cloud Interconnect in us-east4, it's most efficient to configure a single Private NAT gateway instance in us- east4. This allows resources from other regions to egress to on-premises through this single NAT gateway, using a minimal NAT subnet (192.168.1.0/24 in this case), thus conserving the limited 192.168.0.0/20 IP space.

Network Connectivity Center Spoke Export Policy: The VPC spoke needs to advertise the NAT subnet to the Network Connectivity Center hub. An export include policy is used to specify which routes (in this case, the 192.168.1.0/24 NAT subnet) should be advertised to the hub.

Global Dynamic Routing: To allow resources in us-west1, europe-central1, and asia-southeast1 to reach the on-premises location through the us-east4 Cloud Interconnect and NAT gateway, the VPC containing these resources (the spoke VPC) must have global dynamic routing enabled. This ensures that routes learned in one region (like the on-premises routes via us-east4) are available to VMs in all other regions of that VPC.

Options A and B configure Private NAT gateways in multiple regions, which consumes more IP addresses than necessary

given that the Cloud Interconnect is only in us-east4. Option D uses 172.16.x.x for NAT subnets, which clashes with the on-premises IP range and the requirement to use the 192.168.0.0/20 space for cloud.

Exact Extract:

"Private NAT allows instances with private IP addresses in one VPC network to connect to onpremises or other cloud networks through a NAT IP address in a different region or network." "To allow VMs in multiple regions to reach a central destination through a NAT gateway located in a specific region, you must configure global dynamic routing on the VPC network. This ensures that

routes to the NAT gateway's subnet are propagated across all regions."

"When using Network Connectivity Center spokes, you can use export policies to control which routes are advertised from a spoke to the hub. An include policy specifies the exact prefixes to advertise."Reference: Google Cloud Private NAT Documentation, Network Connectivity Center Documentation - Spoke policies, VPC Network Documentation - Dynamic routing mode

Question: 233

Your organization is running out of private IPv4 IP addresses. You need to create a new design pattern to reduce IP usage in your Google Kubernetes Engine clusters. Each new GKE cluster should have a unique /24 range of routable RFC1918 IP addresses. What should you do?

A. Q Configure NAT by using the IP masquerading agent in the GKE cluster.
B. Q Share the primary and secondary ranges between multiple clusters.
C. Q Use dual stack IPv4/IPv6 clusters, and assign IPv6 ranges for Pods and Services.
D. Q Configure the secondary ranges outside the RFC1918 space, or use privately used public IPs.

Answer: C

Explanation:

The most effective long-term solution to address IPv4 address exhaustion in GKE clusters, while still ensuring routability and unique ranges per cluster, is to transition to dual-stack IPv4/IPv6 clusters and leverage IPv6 for Pods and Services. This allows you to conserve IPv4 addresses for critical use cases while providing a vast address space with IPv6 for pods and services, significantly reducing the pressure on your private IPv4 ranges. Google Cloud GKE fully supports dual-stack networking.

Exact Extract:

"Dual-stack clusters enable you to assign both IPv4 and IPv6 addresses to Pods and Services. This approach helps conserve IPv4 address space by shifting a significant portion of the network communication to IPv6, particularly for internal cluster communication or communication with other IPv6-enabled services."

"When you enable dual-stack networking, GKE assigns an IPv6 address range to your Pods and can also assign IPv6 addresses to Services. This significantly expands the available addressing capacity within your cluster."Reference: Google Kubernetes Engine Documentation - Dual-stack networking