Google.Professional Cloud Network Engineer.vJun-2024.by.LorencyProjin.90q

Number: Professional Cloud Network Engineer
Passing Score: 800
Time Limit: 120
File Version: 21.0

Certification: Professional Cloud Network Engineer
Certification Full Name: Professional Cloud Network Engineer

Exam A

QUESTION 1
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.

Correct Answer: A
Section:
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-agke-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-masknotation.html

QUESTION 2
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.

Correct Answer: B
Section:

QUESTION 3
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.

Correct Answer: B
Section:

QUESTION 4
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

Correct Answer: D
Section:

QUESTION 5
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.

Correct Answer: C
Section:
Explanation:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#redundancyoptions

QUESTION 6
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

Correct Answer: B
Section:
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 7
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

Correct Answer: C
Section:
Explanation:
Reference: https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnectsother-projects

QUESTION 8
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.

Correct Answer: C
Section:
Explanation:
Reference: https://geekflare.com/gcp-firewall-configuration/

QUESTION 9
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.

Correct Answer: C
Section:
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 10
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.

Correct Answer: C
Section:

QUESTION 11
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.

Correct Answer: D
Section:
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 12
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

Correct Answer: D
Section:
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 13
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

Correct Answer: A, C
Section:
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 14
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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html

QUESTION 15
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

Correct Answer: A, B
Section:
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 16
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.

Correct Answer: B, C
Section:
Explanation:
https://cloud.google.com/dns/docs/best-practices

QUESTION 17
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.

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/network-connectivity/docs/router/support/troubleshooting#ecmp

QUESTION 18
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.

Correct Answer: D, E
Section:
Explanation:
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/dedicated/retrievingloas

QUESTION 19
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.

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/armor/docs/security-policy-concepts#preview_mode

QUESTION 20
You have a web application that is currently hosted in the us-central1 region. Users experience high latency when traveling in Asi a. 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.

Correct Answer: B
Section:

QUESTION 21
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.

Correct Answer: A, E
Section:
Explanation:
https://cloud.google.com/nat/docs/overview#interaction-pga Specifications
https://cloud.google.com/vpc/docs/configure-private-google-access#specifications

QUESTION 22
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 Googlerecommended 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.

Correct Answer: D
Section:

QUESTION 23
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.

Correct Answer: A
Section:
Explanation:

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

QUESTION 24
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

Correct Answer: D
Section:
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 25
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

Correct Answer: B
Section:

QUESTION 26
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

Correct Answer: A
Section:

QUESTION 27
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.

Correct Answer: A
Section:
Explanation:
Overview By creating and managing SSH keys, you can let users access a Linux instance through thirdparty tools. An SSH key consists of the following files: A public SSH key file that is applied to instancelevel 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 28
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.

Correct Answer: B
Section:

QUESTION 29
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.

Correct Answer: D
Section:
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 30
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.

Correct Answer: B
Section:

QUESTION 31
You have created an HTTP(S) load balanced service. You need to verify that your backend instancesare 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.

Correct Answer: C
Section:
Explanation:
https://cloud.google.com/load-balancing/docs/health-check-concepts#contentbased_health_checks

QUESTION 32
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.

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/interconnect/docs/how-to/dedicated/creating-vlan-attachments

QUESTION 33
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.

Correct Answer: B, D
Section:

QUESTION 34
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.

Correct Answer: B
Section:

QUESTION 35
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	if-tunnel-a-to-b-if-0	169.254.0.254	169.254.0.254	65502
router-2	if-tunnel-b-to-a-if-0	169.254.0.254	169.254.0.254	65501

B.

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

C.

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

D.

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

Correct Answer: C
Section:

QUESTION 36
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.

Correct Answer: A
Section:

QUESTION 37
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.

Correct Answer: D
Section:
Explanation:
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instancegroups#starting_a_canary_update
https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instancegroups

QUESTION 38
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 forthe instance group.
D. Create an unmanaged instance group for each zone, and manually distribute the instances across the desired zones.

Correct Answer: A
Section:
Explanation:
https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances

QUESTION 39
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.

Correct Answer: D
Section:
Explanation:

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

QUESTION 40
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.

Correct Answer: A
Section:
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 41
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.

Correct Answer: A
Section:
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 42
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 offline.
E. SSH into the backend compute engine instances, and view the auth logs and syslogs to further understand the nature of the attack.

Correct Answer: B, E
Section:

QUESTION 43
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.

Correct Answer: A, C
Section:

QUESTION 44
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

Correct Answer: D
Section:
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 45
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.

Correct Answer: A

Section:

QUESTION 46
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.

Correct Answer: D
Section:

QUESTION 47
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.

Correct Answer: A
Section:

QUESTION 48
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 policyfilter for egress traffic.

Correct Answer: B
Section:

QUESTION 49
You are developing an HTTP API hosted on a Compute Engine virtual machine instance that must beinvoked only by multiple clients within the same Virtual Private Cloud (VPC). You want clients to beable 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 forwardingrule. 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 forwardingrule. Then, define an A record in Cloud DNS. Clients should use the name of the A record to connectto 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]/.

Correct Answer: B
Section:

QUESTION 50
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.

Correct Answer: C
Section:

QUESTION 51
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.

Correct Answer: D
Section:

QUESTION 52
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.

Correct Answer: D, E
Section:

QUESTION 53
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

Correct Answer: B
Section:
Explanation:
The correct answer is B. /25.
This answer is based on the following facts:
The Pod per node CIDR range determines the size of the IP address range that is assigned to each node for Pods1. The Pods that run on a node are allocated IP addresses from the node's assigned CIDR range1.
The size of the CIDR range corresponds to the maximum number of Pods per node. For example, a /24 CIDR range allows up to 256 IP addresses, but the default maximum number of Pods per node for Standard clusters is 1102. A /25 CIDR range allows up to 128 IP addresses, which is enough for 100 Pods per node.
The other options are not correct because:
Option A is too large. A /24 CIDR range allows more IP addresses than needed for 100 Pods per node. This could result in inefficient use of the IP address space and limit the number of nodes that can be created in the cluster.
Option C is too small. A /26 CIDR range allows only 64 IP addresses, which is not enough for 60 Pods per node. This could result in insufficient capacity to schedule Pods on the nodes.
Option D is also too small. A /28 CIDR range allows only 16 IP addresses, which is far below the minimum requirement of 60 Pods per node. This could result in Pod scheduling failures and poor performance.

QUESTION 54
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. Re-use 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 Re-use 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

Correct Answer: D
Section:
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 --enable-private-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 55

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.

Correct Answer: B
Section:

QUESTION 56
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].

Correct Answer: A
Section:
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 57
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.

Correct Answer: A
Section:

QUESTION 58
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.

Correct Answer: A
Section:

QUESTION 59
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.

Correct Answer: A
Section:

QUESTION 60
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.

Correct Answer: A
Section:
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/networkconnectivity/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 61
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

Correct Answer: B
Section:
Explanation:
https://cloud.google.com/kubernetes-engine/docs/how-to/aliasips#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 62
You have created a firewall with rules that only allow traffic over HTTP, HTTPS, and SSH ports. Whiletesting, you specifically try to reach the server over multiple ports and protocols; however, you donot 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.

Correct Answer: D
Section:
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-ruleslogging#specifications

QUESTION 63
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.

Correct Answer: B
Section:

QUESTION 64
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.

Correct Answer: C
Section:

QUESTION 65
Your company has 10 separate Virtual Private Cloud (VPC) networks, with one VPC per project in a single region in Google Cloud. Your security team requires each VPC network to have private connectivity to the main on-premises location via a Partner Interconnect connection in the same region. To optimize cost and operations, the same connectivity must be shared with all projects. You must ensure that all traffic between different projects, on-premises locations, and the internet can be inspected using the same third-party appliances. What should you do?

A. Configure the third-party appliances with multiple interfaces and specific Partner Interconnect VLAN attachments per project. 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 thirdparty 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.

Correct Answer: D
Section:

QUESTION 66
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.

Correct Answer: A
Section:

QUESTION 67
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.

Correct Answer: C
Section:

QUESTION 68
You have a Cloud Storage bucket in Google Cloud project XYZ. The bucket contains sensitive dat a. 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.

Correct Answer: C
Section:

QUESTION 69
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.

Correct Answer: C
Section:

QUESTION 70

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.

Correct Answer: B
Section:

QUESTION 71
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 CloudLogging.
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.

Correct Answer: A
Section:

QUESTION 72
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.

Correct Answer: B
Section:

QUESTION 73
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.

Correct Answer: A
Section:

QUESTION 74
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-ofservice attacks.
D. Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at theorganization level.

Correct Answer: C
Section:

QUESTION 75
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 onpremises data center.

Correct Answer: A
Section:

QUESTION 76
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 highpriority 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 highpriority 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.

Correct Answer: B
Section:

QUESTION 77
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.

Correct Answer: A
Section:

QUESTION 78
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.

Correct Answer: B
Section:

QUESTION 79
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"

Correct Answer: C
Section:

QUESTION 80
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.

Correct Answer: B
Section:

QUESTION 81
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.

Correct Answer: C
Section:

QUESTION 82
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.

Correct Answer: D
Section:

QUESTION 83
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

Correct Answer: D
Section:

QUESTION 84
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

Correct Answer: C
Section:
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 85
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

Correct Answer: D
Section:
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 86
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.

Correct Answer: B
Section:

QUESTION 87
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.

Correct Answer: C
Section:

QUESTION 88
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 do?

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.

Correct Answer: D
Section:
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.
Bring your own IP addresses
Professional Cloud Network Engineer Exam Guide
Bring your own IP addresses (BYOIP) to Azure with Custom IP Prefix

QUESTION 89
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.

Correct Answer: D
Section:
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 90
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?

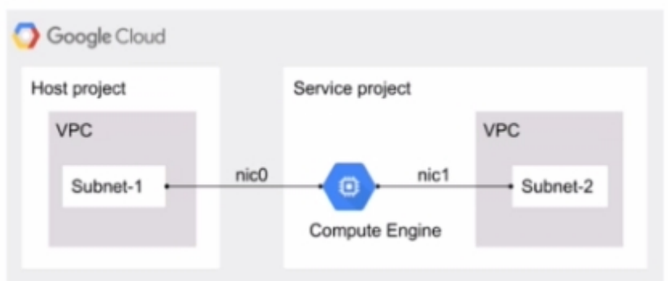

A. Configure a firewall rule to permit Subnet-2 IP addresses outbound in the host protect VPC.
B. Configure Packet Mirroring in both the host and service project VPCs.
C. Configure a VPC Flow Logs filter for Subnet-2 in the host project VPC.
D. Configure VPC Flow Logs in the service project VPC for Subnet-2.

Correct Answer: D
Section:
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.
Enable 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.
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.