Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer

NEW QUESTION 1
You need to define an address plan for a future new Google Kubernetes Engine (GKE) cluster in your Virtual Private Cloud (VPC). This will be a VPC-native cluster, and the default Pod IP range allocation will be used. You must pre-provision all the needed VPC subnets and their respective IP address ranges before cluster creation. The cluster will initially have a single node, but it will be scaled to a maximum of three nodes if necessary. You want to allocate the minimum number of Pod IP addresses. Which subnet mask should you use for the Pod IP address range?

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

Answer: A

NEW QUESTION 2
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 resolve
B. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.
C. Configure both an inbound server policy and outbound DNS forwarding zones with the target as the on-premises DNS resolve
D. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
E. Configure an outbound DNS server policy, and set your alternative name server to be your on-premises DNS resolve
F. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
G. Configure Cloud DNS to DNS peer with your on-premises DNS resolve
H. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

Answer: A

NEW QUESTION 3
You have ordered Dedicated Interconnect in the GCP Console and need to give the Letter of Authorization/Connecting Facility Assignment (LOA-CFA) to your cross-connect provider to complete the physical connection.
Which two actions can accomplish this? (Choose two.)

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

Answer: DE

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

NEW QUESTION 4
You converted an auto mode VPC network to custom mode. Since the conversion, some of your Cloud Deployment Manager templates are no longer working.
You want to resolve the problem.
What should you do?

A. Apply an additional IAM role to the Google API's service account to allow custom mode networks.
B. Update the VPC firewall to allow the Cloud Deployment Manager to access the custom mode networks.
C. Explicitly reference the custom mode networks in the Cloud Armor whitelist.
D. Explicitly reference the custom mode networks in the Deployment Manager templates.

Answer: D

NEW QUESTION 5
You create a Google Kubernetes Engine private cluster and want to use kubectl to get the status of the pods. In one of your instances you notice the master is not responding, even though the cluster is up and running.
What should you do to solve the problem?

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

Answer: D

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

NEW QUESTION 6
You are responsible for configuring firewall policies for your company in Google Cloud. Your security team has a strict set of requirements that must be met to configure firewall rules.

Always allow Secure Shell (SSH) from your corporate IP address. Restrict SSH access from all other IP addresses.
There are multiple projects and VPCs in your Google Cloud organization. You need to ensure that other VPC firewall rules cannot bypass the security team's requirements. What should you do?

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

Answer: A

NEW QUESTION 7
You created a new VPC network named Dev with a single subnet. You added a firewall rule for the network Dev to allow HTTP traffic only and enabled logging. When you try to log in to an instance in the subnet via Remote Desktop Protocol, the login fails. You look for the Firewall rules logs in Stackdriver Logging, but you do not see any entries for blocked traffic. You want to see the logs for blocked traffic.
What should you do?

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

Answer: D

Explanation:
Ingress packets in VPC Flow Logs are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs. We want to see the logs for blocked traffic so we have to look for them in firewall logs.
https://cloud.google.com/vpc/docs/flow-logs#key_properties

NEW QUESTION 8
You need to enable Private Google Access for use by some subnets within your Virtual Private Cloud (VPC). Your security team set up the VPC to send all internet-bound traffic back to the on- premises data center for inspection before egressing to the internet, and is also implementing VPC Service Controls in the environment for API-level security control. You have already enabled the subnets for Private Google Access. What configuration changes should you make to enable Private Google Access while adhering to your security team's requirements?

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

Answer: C

NEW QUESTION 9
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 addresse
B. Configure a (25 secondary IP range for the Pod
C. Configure a /22 secondary IP range for the Services.
D. Configure a /28 primary IP address range for the node IP addresse
E. Configure a /25 secondary IP range for the Pod
F. Configure a /21 secondary IP range for the Services.
G. Configure a /28 primary IP address range for the node IP addresse
H. Configure a /28 secondary IP range for the Pod
I. Configure a /21 secondary IP range for the Services.
J. Configure a /28 primary IP address range for the node IP addresse
K. Configure a /24 secondary IP range for the Pad
L. Configure a /22 secondary IP range for the Services.

Answer: A

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

A. Tune TCP parameters on the on-premises servers.
B. Compress files using utilities like tar to reduce the size of data being sent.
C. Remove the -m flag from the gsutil command to enable single-threaded transfers.

D. Use the perfdiag parameter in your gsutil command to enable faster performance: gsutil perfdiag gs://[BUCKET NAME].

Answer: A

Explanation:
https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid
https://cloud.google.com/blog/products/gcp/5-steps-to-better-gcp-network-performance?hl=ml

NEW QUESTION 10
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 manuallydistribute instances across the zones in that region.
C. Create an unmanaged instance group in a single zone, and then create an HTTP load balancer for the instance group.
D. Create an unmanaged instance group for each zone, and manually distribute the instances across the desired zones.

Answer: A

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

NEW QUESTION 11
You want to use Partner Interconnect to connect your on-premises network with your VPC. You already have an Interconnect partner.
What should you first?

A. Log in to your partner's portal and request the VLAN attachment there.
B. Ask your Interconnect partner to provision a physical connection to Google.
C. Create a Partner Interconnect type VLAN attachment in the GCP Console and retrieve the pairing key.
D. Run gcloud compute interconnect attachments partner update <attachment> / -- region <region>--admin-enabled.

Answer: B

Explanation:
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview?hl=En#provisionin "To provision a Partner Interconnect connection with a service provider, you start by connecting your
on-premises network to a supported service provider. Work with the service provider to establish connectivity.

NEW QUESTION 16
You have configured a service on Google Cloud that connects to an on-premises service via a Dedicated Interconnect. Users are reporting recent connectivity issues. You need to determine whether the traffic is being dropped because of firewall rules or a routing decision. What should you do?

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

Answer: B

NEW QUESTION 20
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 perimete
B. Define fine-grained perimeter controls and enforce that security posture across your Google Cloud services and projects.
C. Configure a Google Cloud Armor security policy in your project, and attach it to the backend service to secure the application.
D. Configure VPC firewall rules to protect the Compute Engine instances against distributed denial-of-service attacks.
E. Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at the organization level.

Answer: C

NEW QUESTION 25
You recently deployed Compute Engine instances in regions us-west1 and us-east1 in a Virtual Private Cloud (VPC) with default routing configurations. Your company security policy mandates that virtual machines (VMs) must not have public IP addresses attached to them. You need to allow your instances to fetch updates from the internet while preventing external access. What should you do?

A. Create a Cloud NAT gateway and Cloud Router in both us-west1 and us-east1.
B. Create a single global Cloud NAT gateway and global Cloud Router in the VPC.
C. Change the instances' network interface external IP address from None to Ephemeral.
D. Create a firewall rule that allows egress to destination 0.0.0.0/0.

Answer: A

NEW QUESTION 29
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.

Answer: A

NEW QUESTION 31
You want to use Cloud Interconnect to connect your on-premises network to a GCP VPC. You cannot meet Google at one of its point-of-presence (POP) locations, and your on-premises router cannot run a Border Gateway Protocol (BGP) configuration.
Which connectivity model should you use?

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

Answer: D

Explanation:
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview
For Layer 3 connections, your service provider establishes a BGP session between your Cloud Routers and their edge routers for each VLAN attachment. You don't need to configure BGP on your on-premises router. Google and your service provider automatically set the correct configurations.
https://cloud.google.com/network-connectivity/docs/interconnect/concepts/partner-overview#connectivity-type

NEW QUESTION 35
You have created a firewall with rules that only allow traffic over HTTP, HTTPS, and SSH ports. While testing, you specifically try to reach the server over multiple ports and protocols; however, you do not see any denied connections in the firewall logs. You want to resolve the issue.
What should you do?

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

Answer: D

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

NEW QUESTION 36
You have applications running in the us-west1 and us-east1 regions. You want to build a highly available VPN that provides 99.99% availability to connect your applications from your project to the cloud services provided by your partner's project while minimizing the amount of infrastructure required. Your partner's services are also in the us-west1 and us-east1 regions. You want to implement the simplest solution. What should you do?

A. Create one Cloud Router and one HA VPN gateway in each region of your VPC and your partner's VP
B. Connect your VPN gateways to the partner's gateway
C. Enable global dynamic routing in each VPC.
D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VP
E. Create one OpenVPN Access Server in each region of your partner's VP
F. Connect your VPN gateway to your partner's servers.
G. Create one OpenVPN Access Server in each region of your VPC and your partner's VP
H. Connect your servers to the partner's servers.
I. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VP
J. Connect your VPN gateways to the partner's gateways with a pair of tunnel
K. Enable global dynamic routing in each VPC.

Answer: A

NEW QUESTION 37
You created a new VPC for your development team. You want to allow access to the resources in this VPC via SSH only.
How should you configure your firewall rules?

A. Create two firewall rules: one to block all traffic with priority 0, and another to allow port 22 with priority 1000.

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

Answer: C

NEW QUESTION 38
You are migrating to Cloud DNS and want to import your BIND zone file. Which command should you use?

A. gcloud dns record-sets import ZONE_FILE --zone MANAGED_ZONE
B. gcloud dns record-sets import ZONE_FILE --replace-origin-ns --zone MANAGED_ZONE
C. gcloud dns record-sets import ZONE_FILE --zone-file-format --zone MANAGED_ZONE
D. gcloud dns record-sets import ZONE_FILE --delete-all-existing --zone MANAGED ZONE

Answer: C

Explanation:
https://cloud.google.com/sdk/gcloud/reference/dns/record-sets/import

NEW QUESTION 43
Your company is running out of network capacity to run a critical application in the on-premises data center. You want to migrate the application to GCP. You also want to ensure that the Security team does not lose their ability to monitor traffic to and from Compute Engine instances.
Which two products should you incorporate into the solution? (Choose two.)

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

Answer: AB

Explanation:
A: Using VPC Flow Logs VPC Flow Logs records a sample of network flows sent from and received by VM instances, including instances used as GKE nodes. These logs can be used for network monitoring, forensics, real-time security analysis, and expense optimization. https://cloud.google.com/vpc/docs/using-flow-logs
(B): Firewall Rules Logging overview Firewall Rules Logging allows you to audit, verify, and analyze the effects of your firewall rules. For example, you can determine if a firewall rule designed to deny traffic is functioning
as intended. Firewall Rules Logging is also useful if you need to determine how many connections are affected by a given firewall rule. You enable Firewall Rules Logging individually for each firewall rule whose connections you need to log. Firewall Rules Logging is an option for any firewall rule, regardless of the action (allow or deny) or direction (ingress or egress) of the rule.
https://cloud.google.com/vpc/docs/firewall-rules-logging

NEW QUESTION 47
Your company offers a popular gaming service. Your instances are deployed with private IP addresses, and external access is granted through a global load balancer. You believe you have identified a potential malicious actor, but aren't certain you have the correct client IP address. You want to identify this actor while minimizing disruption to your legitimate users.
What should you do?

A. Create a Cloud Armor Policy rule that denies traffic and review necessary logs.
B. Create a Cloud Armor Policy rule that denies traffic, enable preview mode, and review necessary logs.
C. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to disabled, and review necessary logs.
D. Create a VPC Firewall rule that denies traffic, enable logging and set enforcement to enabled, and review necessary logs.

Answer: B

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

NEW QUESTION 49
Your company is working with a partner to provide a solution for a customer. Both your company and the partner organization are using GCP. There are applications in the partner's network that need access to some resources in your company's VPC. There is no CIDR overlap between the VPCs.
Which two solutions can you implement to achieve the desired results without compromising the security? (Choose two.)

A. VPC peering
B. Shared VPC
C. Cloud VPN
D. Dedicated Interconnect
E. Cloud NAT

Answer: AC

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

NEW QUESTION 53
Your end users are located in close proximity to us-east1 and europe-west1. Their workloads need to communicate with each other. You want to minimize cost

and increase network efficiency.
How should you design this topology?

A. Create 2 VPCs, each with their own regions and individual subnet
B. Create 2 VPN gateways to establish connectivity between these regions.
C. Create 2 VPCs, each with their own region and individual subnet
D. Use external IP addresses on the instances to establish connectivity between these regions.
E. Create 1 VPC with 2 regional subnet
F. Create a global load balancer to establish connectivity between the regions.
G. Create 1 VPC with 2 regional subnet
H. Deploy workloads in these subnets and have them communicate using private RFC1918 IP addresses.

Answer: D

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

NEW QUESTION 56
All the instances in your project are configured with the custom metadata enable-oslogin value set to FALSE and to block project-wide SSH keys. None of the instances are set with any SSH key, and no project-wide SSH keys have been configured. Firewall rules are set up to allow SSH sessions from any IP address range. You want to SSH into one instance.
What should you do?

A. Open the Cloud Shell SSH into the instance using gcloud compute ssh.
B. Set the custom metadata enable-oslogin to TRUE, and SSH into the instance using a third-party tool like putty or ssh.
C. Generate a new SSH key pai
D. Verify the format of the private key and add it to the instanc
E. SSH into the instance using a third-party tool like putty or ssh.
F. Generate a new SSH key pai
G. Verify the format of the public key and add it to the projec
H. SSH into the instance using a third-party tool like putty or ssh.

Answer: A

NEW QUESTION 60
You are responsible for enabling Private Google Access for the virtual machine (VM) instances in your Virtual Private Cloud (VPC) to access Google APIs. All VM instances have only a private IP address and need to access Cloud Storage. You need to ensure that all VM traffic is routed back to your on-premises data center for traffic scrubbing via your existing Cloud Interconnect connection. However, VM traffic to Google APIs should remain in the VPC. What should you do?

A. Delete the default route in your VPC.Create a private Cloud DNS zone for googleapis.com, create a CNAME for *.googleapis.com to restricted googleapis.com, and create an A record for restricted googleapis com that resolves to the addresses in 199.36.153.4/30.Create a static route in your VPC for the range 199.36.153.4/30 with the default internet gateway as the next hop.
B. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).Create a public Cloud DNS zone with a CNAME for *.google.com to private googleapis com, create a CNAME for * googleapis.com to private googleapis com, and create an A record for Private googleapis.com that resolves to the addresses in 199.36.153 8/30.Create a static route in your VPC for the range 199 .36.153.8/30 with the default internet gateway as the next hop.
C. Configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP) with a lower priority (MED) than the default VPC route.Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googieapis.com to private googleapis com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30.Create a static route in your VPC for the range 199.36. 153.8/30 with the default internet gateway as the next hop.
D. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP).Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googieapis.com to Private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30.Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.

Answer: C

NEW QUESTION 65
After a network change window one of your company’s applications stops working. The application uses an on-premises database server that no longer receives any traffic from the application. The database server IP address is 10.2.1.25. You examine the change request, and the only change is that 3 additional VPC subnets were created. The new VPC subnets created are 10.1.0.0/16, 10.2.0.0/16, and 10.3.1.0/24/ The on-premises router is advertising 10.0.0.0/8.
What is the most likely cause of this problem?

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

Answer: B

NEW QUESTION 69
You are designing a shared VPC architecture. Your network and security team has strict controls over which routes are exposed between departments. Your Production and Staging departments can communicate with each other, but only via specific networks. You want to follow Google-recommended practices.
How should you design this topology?

A. Create 2 shared VPCs within the shared VPC Host Project, and enable VPC peering between the
B. Use firewall rules to filter access between the specific networks.

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

Answer: D

NEW QUESTION 70
You want to configure a NAT to perform address translation between your on-premises network blocks and GCP.
Which NAT solution should you use?

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

Answer: A

NEW QUESTION 74
You want to create a service in GCP using IPv6. What should you do?

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

Answer: C

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

NEW QUESTION 79
You have an HA VPN connection with two tunnels running in active/passive mode between your Virtual Private Cloud (VPC) and on-premises network. Traffic over the connection has recently increased from 1 gigabit per second (Gbps) to 4 Gbps, and you notice that packets are being dropped. You need to configure your VPN connection to Google Cloud to support 4 Gbps. What should you do?

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

Answer: D

NEW QUESTION 80
One instance in your VPC is configured to run with a private IP address only. You want to ensure that even if this instance is deleted, its current private IP address will not be automatically assigned to a different instance.
In the GCP Console, what should you do?

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

Answer: C

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

NEW QUESTION 85
You are planning a large application deployment in Google Cloud that includes on-premises connectivity. The application requires direct connectivity between workloads in all regions and on-premises locations without address translation, but all RFC 1918 ranges are already in use in the on-premises locations. What should you do?

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

Answer: D

NEW QUESTION 90
You successfully provisioned a single Dedicated Interconnect. The physical connection is at a colocation facility closest to us-west2. Seventy-five percent of your workloads are in us-east4, and the remaining twenty-five percent of your workloads are in us-central1. All workloads have the same network traffic profile. You need to minimize data transfer costs when deploying VLAN attachments. What should you do?

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

Answer: C

NEW QUESTION 94
You are adding steps to a working automation that uses a service account to authenticate. You need to drive the automation the ability to retrieve files from a Cloud Storage bucket. Your organization requires using the least privilege possible.
What should you do?

A. Grant the compute.instanceAdmin to your user account.
B. Grant the iam.serviceAccountUser to your user account.
C. Grant the read-only privilege to the service account for the Cloud Storage bucket.
D. Grant the cloud-platform privilege to the service account for the Cloud Storage bucket.

Answer: C

NEW QUESTION 97
You are creating a new application and require access to Cloud SQL from VPC instances without public IP addresses.
Which two actions should you take? (Choose two.)

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

Answer: CE

Explanation:
https://cloud.google.com/sql/docs/mysql/configure-private-services-access#console_1
C: If you are using private IP for any of your Cloud SQL instances, you only need to configure private services access one time for every Google Cloud project that has or needs to connect to a Cloud SQL instance. If your Google Cloud project has a Cloud SQL instance, you can either configure it yourself or let Cloud SQL do it for you to use private IP. Cloud SQL configures private services access for you when all the conditions below are true:
https://cloud.google.com/sql/docs/postgres/configure-private-services-access#before_you_begin
E: You can enable Private Google access on a subnet level and any VMs on that subnet can access Google APIs by using their internal IP address.
https://cloud.google.com/vpc/docs/configure-private-google-access

NEW QUESTION 101
You recently deployed your application in Google Cloud. You need to verify your Google Cloud network configuration before deploying your on-premises workloads. You want to confirm that your Google Cloud network configuration allows traffic to flow from your cloud resources to your on- premises network. This validation should also analyze and diagnose potential failure points in your Google Cloud network configurations without sending any data plane test traffic. What should you do?

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

Answer: C

NEW QUESTION 102
You are the network administrator responsible for hybrid connectivity at your organization. Your developer team wants to use Cloud SQL in the us-west1 region in your Shared VPC. You configured a Dedicated Interconnect connection and a Cloud Router in us-west1, and the connectivity between your Shared VPC and on-premises data center is working as expected. You just created the private services access connection required for Cloud SQL using the reserved IP address range and default settings. However, your developers cannot access the Cloud SQL instance from on-premises. You want to resolve the issue. What should you do?

A. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.
B. Change the VPC routing mode to global.Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.
C. Create an additional Cloud Router in us-west2.Create a new Border Gateway Protocol (BGP) peering connection to your on-premises data cente
D. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export ofroutes.
E. Change the VPC routing mode to global.Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.

Answer: A

NEW QUESTION 107
You are designing a Google Kubernetes Engine (GKE) cluster for your organization. The current cluster size is expected to host 10 nodes, with 20 Pods per node and 150 services. Because of the migration of new services over the next 2 years, there is a planned growth for 100 nodes, 200 Pods per node, and 1500 services. You want to use VPC-native clusters with alias IP ranges, while minimizing address consumption.
How should you design this topology?

A. Create a subnet of size/25 with 2 secondary ranges of: /17 for Pods and /21 for Service
B. Create a VPC-native cluster and specify those ranges.

C. Create a subnet of size/28 with 2 secondary ranges of: /24 for Pods and /24 for Service
D. Create a VPC-native cluster and specify those range
E. When the services are ready to be deployed, resize the subnets.
F. Use gcloud container clusters create [CLUSTER NAME]--enable-ip-alias to create a VPC-native cluster.
G. Use gcloud container clusters create [CLUSTER NAME] to create a VPC-native cluster.

Answer: A

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

NEW QUESTION 111
Your company has separate Virtual Private Cloud (VPC) networks in a single region for two departments: Sales and Finance. The Sales department's VPC network already has connectivity to on-premises locations using HA VPN, and you have confirmed that the subnet ranges do not overlap. You plan to peer both VPC networks to use the same HA tunnels for on-premises connectivity, while providing internet connectivity for the Google Cloud workloads through Cloud NAT. Internet access from the on-premises locations should not flow through Google Cloud. You need to propagate all routes between the Finance department and on-premises locations. What should you do?

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

Answer: A

NEW QUESTION 113
You are configuring a new instance of Cloud Router in your Organization’s Google Cloud environment to allow connection across a new Dedicated Interconnect to your data center Sales, Marketing, and IT each have a service project attached to the Organization’s host project.
Where should you create the Cloud Router instance?

A. VPC network in all projects
B. VPC network in the IT Project
C. VPC network in the Host Project
D. VPC network in the Sales, Marketing, and IT Projects

Answer: C

NEW QUESTION 117
You have two Google Cloud projects in a perimeter to prevent data exfiltration. You need to move a third project inside the perimeter; however, the move could negatively impact the existing environment. You need to validate the impact of the change. What should you do?

A. Enable Firewall Rules Logging inside the third project.
B. Modify the existing VPC Service Controls policy to include the new project in dry run mode.
C. Monitor the Resource Manager audit logs inside the perimeter.
D. Enable VPC Flow Logs inside the third project, and monitor the logs for negative impact.

Answer: B

NEW QUESTION 119
Your company's security team wants to limit the type of inbound traffic that can reach your web servers to protect against security threats. You need to configure the firewall rules on the web servers within your Virtual Private Cloud (VPC) to handle HTTP and HTTPS web traffic for TCP only. What should you do?

A. Create an allow on match ingress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.
B. Create an allow on match egress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.
C. Create an allow on match ingress firewall rule with the target tag “web-server” to allow all IP addresses for TCP ports 80 and 443.
D. Create an allow on match egress firewall rule with the target tag “web-server" to allow web server IP addresses for TCP ports 60 and 443.

Answer: C

NEW QUESTION 121
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 addres
D. Create a second tunnel on the existing Cloud VPN gateway that forwards the same IP range, but points at the new on-premises gateway IP.
E. Add a second Cloud VPN gateway in a different region than the existing VPN gatewa
F. Create a new tunnel on the second Cloud VPN gateway that forwards the same IP range, but points to the existing on-premises VPN gateway IP address.

Answer: C

Explanation:
https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#redundancy-options

NEW QUESTION 122
You need to create a new VPC network that allows instances to have IP addresses in both the 10.1.1.0/24 network and the 172.16.45.0/24 network.
What should you do?

A. Configure global load balancing to point 172.16.45.0/24 to the correct instance.
B. Create unique DNS records for each service that sends traffic to the desired IP address.
C. Configure an alias-IP range of 172.16.45.0/24 on the virtual instances within the VPC subnet of 10.1.1.0/24.
D. Use VPC peering to allow traffic to route between the 10.1.0.0/24 network and the 172.16.45.0/24network.

Answer: C

NEW QUESTION 126
Your organization has a single project that contains multiple Virtual Private Clouds (VPCs). You need to secure API access to your Cloud Storage buckets and BigQuery datasets by allowing API access only from resources in your corporate public networks. What should you do?

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

Answer: B

NEW QUESTION 131
Your on-premises data center has 2 routers connected to your Google Cloud environment through a VPN on each router. All applications are working correctly; however, all of the traffic is passing across a single VPN instead of being load-balanced across the 2 connections as desired.
During troubleshooting you find:
• Each on-premises router is configured with a unique ASN.
• Each on-premises router is configured with the same routes and priorities.
• Both on-premises routers are configured with a VPN connected to a single Cloud Router.
• BGP sessions are established between both on-premises routers and the Cloud Router.
• Only 1 of the on-premises router's routes are being added to the routing table. What is the most likely cause of this problem?

A. The on-premises routers are configured with the same routes.
B. A firewall is blocking the traffic across the second VPN connection.
C. You do not have a load balancer to load-balance the network traffic.
D. The ASNs being used on the on-premises routers are different.

Answer: D

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

NEW QUESTION 134
You need to ensure your personal SSH key works on every instance in your project. You want to accomplish this as efficiently as possible.
What should you do?

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

Answer: A

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

NEW QUESTION 135
You have provisioned a Dedicated Interconnect connection of 20 Gbps with a VLAN attachment of 10 Gbps. You recently noticed a steady increase in ingress traffic on the Interconnect connection from the on-premises data center. You need to ensure that your end users can achieve the full 20 Gbps throughput as quickly as possible. Which two methods can you use to accomplish this? (Choose two.)

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

Answer: CE

NEW QUESTION 136

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

Professional-Cloud-Network-Engineer Practice Exam Features:

* Professional-Cloud-Network-Engineer Questions and Answers Updated Frequently

* Professional-Cloud-Network-Engineer Practice Questions Verified by Expert Senior Certified Staff

* Professional-Cloud-Network-Engineer Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* Professional-Cloud-Network-Engineer Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year