

## [December-2022Free Braindump2go Professional-Cloud-Network-Engineer VCE Dumps Free Professional-Cloud-Network-Engineer 155Q Download[Q95-Q143

December/2022 Latest Braindump2go Professional-Cloud-Network-Engineer Exam Dumps with PDF and VCE Free Updated Today! Following are some new [Braindump2go Professional-Cloud-Network-Engineer Real Exam Questions!](#)

**QUESTION 95**Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with access from your on-premises network using Cloud Interconnect. You must configure access only to Google APIs and services that are supported by VPC Service Controls through hybrid connectivity with a service level agreement (SLA) in place. What should you do?  
A. Configure the existing Cloud Routers to advertise the Google API's public virtual IP addresses.  
B. Use Private Google Access for on-premises hosts with restricted.googleapis.com virtual IP addresses.  
C. Configure the existing Cloud Routers to advertise a default route, and use Cloud NAT to translate traffic from your on-premises network.  
D. Add Direct Peering links, and use them for connectivity to Google APIs that use public virtual IP addresses.  
Answer: B  
Explanation:  
<https://cloud.google.com/vpc/docs/configure-private-google-access-hybrid>

**QUESTION 96**Your company's security team tends to use managed services when possible. You need to build a dashboard to show the number of deny hits that occur against configured firewall rules without increasing operational overhead. What should you do?  
A. Configure Firewall Rules Logging. Use Firewall Insights to display the number of hits.  
B. Configure Firewall Rules Logging. View the logs in Cloud Logging, and create a custom dashboard in Cloud Monitoring to display the number of hits.  
C. Configure a firewall appliance from the Google Cloud Marketplace. Route all traffic through this appliance, and apply the firewall rules at this layer. Use the firewall appliance to display the number of hits.  
D. Configure Packet Mirroring on the VPC. Apply a filter with an IP address list of the Denied Firewall rules. Configure an intrusion detection system (IDS) appliance as the receiver to display the number of hits.  
Answer: A  
Explanation:  
<https://cloud.google.com/network-intelligence-center/docs/firewall-insights/concepts/overview>

**QUESTION 97**You are configuring your Google Cloud environment to connect to your on-premises network. Your configuration must be able to reach Cloud Storage APIs and your Google Kubernetes Engine nodes across your private Cloud Interconnect network. You have already configured a Cloud Router with your Interconnect VLAN attachments. You now need to set up the appropriate router advertisement configuration on the Cloud Router. What should you do?  
A. Configure the route advertisement to the default setting.  
B. On the on-premises router, configure a static route for the storage API virtual IP address which points to the Cloud Router's link-local IP address.  
C. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisements. Leave all other options as their default settings.  
D. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisements. Advertise all visible subnets to the Cloud Router.  
Answer: B  
Explanation:  
<https://cloud.google.com/vpc/docs/private-google-access-hybrid>

**QUESTION 98**You are configuring load balancing for a standard three-tier (web, application, and database) application. You have configured an external HTTP(S) load balancer for the web servers. You need to configure load balancing for the application tier of servers. What should you do?  
A. Configure a forwarding rule on the existing load balancer for the application tier.  
B. Configure equal cost multi-path routing on the application servers.  
C. Configure a new internal HTTP(S) load balancer for the application tier.  
D. Configure a URL map on the existing load balancer to route traffic to the application tier.  
Answer: D  
Explanation:URL maps are used to direct traffic to the back ends and this would be where the application is located.

**QUESTION 99**Your organization has a new security policy that requires you to monitor all egress traffic payloads from your virtual machines in region us-west2. You deployed an intrusion detection system (IDS) virtual appliance in the same region to meet the new policy. You now need to integrate the IDS into the environment to monitor all egress traffic payloads from us-west2. What should you do?  
A. Enable firewall logging, and forward all filtered egress firewall logs to the IDS.  
B. Enable VPC Flow Logs. Create a sink in Cloud Logging to send filtered egress VPC Flow Logs to the IDS.  
C. Create an internal TCP/UDP load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.  
D. Create an internal HTTP(S) load balancer for Packet Mirroring, and add a packet mirroring policy filter for egress traffic.  
Answer: C  
Explanation:IDS requires at least 1 packet mirroring policy attached to it.  
<https://cloud.google.com/intrusion-detection-system/docs/overview>  
<https://cloud.google.com/vpc/docs/packet-mirroring>

**QUESTION 100**You are developing an HTTP API hosted on a Compute Engine virtual machine instance that must be invoked only by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?  
A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Clients should use this IP address to connect to the service.  
B. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[INSTANCE\\_NAME\].\[ZONE.c\].\[PROJECT\\_ID\].internal/](https://[INSTANCE_NAME].[ZONE.c].[PROJECT_ID].internal/).  
C. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rule. Then, define an A record in Cloud DNS. Clients should use the

name of the A record to connect to the service.D. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url [https://\[API\\_NAME\]/\[API\\_VERSION/](https://[API_NAME]/[API_VERSION/).Answer: BExplanation:Virtual Private Cloud networks on Google Cloud have an internal DNS service that lets instances in the same network access each other by using internal DNS namesThis name can be used for access: [INSTANCE\_NAME].[ZONE].c.[PROJECT\_ID].internal

[https://cloud.google.com/compute/docs/internal-dns#access\\_by\\_internal\\_dns](https://cloud.google.com/compute/docs/internal-dns#access_by_internal_dns)QUESTION 101You recently deployed Cloud VPN to connect your on-premises data center to Google Cloud. You need to monitor the usage of this VPN and set up alerts in case traffic exceeds the maximum allowed. You need to be able to quickly decide whether to add extra links or move to a Dedicated Interconnect. What should you do?A. In the Network Intelligence Center, check for the number of packet drops on the VPN.B. In the Google Cloud Console, use Monitoring Query Language to create a custom alert for bandwidth utilization.C. In the Monitoring section of the Google Cloud Console, use the Dashboard section to select a default dashboard for VPN usage.D. In the VPN section of the Google Cloud Console, select the VPN under hybrid connectivity, and then select monitoring to display utilization on the dashboard.Answer: AQUESTION 102You have applications running in the us-west1 and us-east1 regions. You want to build a highly available VPN that provides 99.99% availability to connect your applications from your project to the cloud services provided by your partner's project while minimizing the amount of infrastructure required. Your partner's services are also in the us-west1 and us-east1 regions. You want to implement the simplest solution. What should you do?A. Create one Cloud Router and one HA VPN gateway in each region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways. Enable global dynamic routing in each VPC.B. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC. Create one OpenVPN Access Server in each region of your partner's VPC. Connect your VPN gateway to your partner's servers.C. Create one OpenVPN Access Server in each region of your VPC and your partner's VPC. Connect your servers to the partner's servers.D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VPC. Connect your VPN gateways to the partner's gateways with a pair of tunnels. Enable global dynamic routing in each VPC.Answer: DExplanation:<https://cloud.google.com/static/network-connectivity/docs/vpn/images/ha-vpn-gcp-to-on-prem-2-a.svg>  
<https://cloud.google.com/network-connectivity/docs/vpn/concepts/topologies>QUESTION 103You need to create the network infrastructure to deploy a highly available web application in the us-east1 and us-west1 regions. The application runs on Compute Engine instances, and it does not require the use of a database. You want to follow Google-recommended practices. What should you do?A. Create one VPC with one subnet in each region.Create a regional network load balancer in each region with a static IP address.Enable Cloud CDN on the load balancers.Create an A record in Cloud DNS with both IP addresses for the load balancers.B. Create one VPC with one subnet in each region.Create a global load balancer with a static IP address.Enable Cloud CDN and Google Cloud Armor on the load balancer.Create an A record using the IP address of the load balancer in Cloud DNS.C. Create one VPC in each region, and peer both VPCs.Create a global load balancer.Enable Cloud CDN on the load balancer.Create a CNAME for the load balancer in Cloud DNS.D. Create one VPC with one subnet in each region.Create an HTTP(S) load balancer with a static IP address.Choose the standard tier for the network.Enable Cloud CDN on the load balancer.Create a CNAME record using the load balancer's IP address in Cloud DNS.Answer: CQUESTION 104You are the network administrator responsible for hybrid connectivity at your organization. Your developer team wants to use Cloud SQL in the us-west1 region in your Shared VPC. You configured a Dedicated Interconnect connection and a Cloud Router in us-west1, and the connectivity between your Shared VPC and on-premises data center is working as expected. You just created the private services access connection required for Cloud SQL using the reserved IP address range and default settings. However, your developers cannot access the Cloud SQL instance from on-premises. You want to resolve the issue. What should you do?A. Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.B. Change the VPC routing mode to global.Create a custom route advertisement in your Cloud Router to advertise the Cloud SQL IP address range.C. Create an additional Cloud Router in us-west2.Create a new Border Gateway Protocol (BGP) peering connection to your on-premises data center.Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.D. Change the VPC routing mode to global.Modify the VPC Network Peering connection used for Cloud SQL, and enable the import and export of routes.Answer: AQUESTION 105Your company has separate Virtual Private Cloud (VPC) networks in a single region for two departments: Sales and Finance. The Sales department's VPC network already has connectivity to on-premises locations using HA VPN, and you have confirmed that the subnet ranges do not overlap. You plan to peer both VPC networks to use the same HA tunnels for on-premises connectivity, while providing internet connectivity for the Google Cloud workloads through Cloud NAT. Internet access from the on-premises locations should not flow through Google Cloud. You need to propagate all routes between the Finance department and on-premises locations. What should you do?A. Peer the two VPCs, and use the default configuration for the Cloud Routers.B. Peer the two VPCs, and use Cloud

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

Answer: A

QUESTION 106 You recently noticed a recurring daily spike in network usage in your Google Cloud project. You need to identify the virtual machine (VM) instances and type of traffic causing the spike in traffic utilization while minimizing the cost and management overhead required. What should you do?

A. Enable VPC Flow Logs and send the output to BigQuery for analysis.  
B. Enable Firewall Rules Logging for all allowed traffic and send the output to BigQuery for analysis.  
C. Configure Packet Mirroring to send all traffic to a VM. Use Wireshark on the VM to identify traffic utilization for each VM in the VPC.  
D. Deploy a third-party network appliance and configure it as the default gateway. Use the third-party network appliance to identify users with high network traffic.

Answer: C

QUESTION 107 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.  
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.  
C. 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 AP address range.  
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.

Answer: C

QUESTION 108 You have deployed an HTTP(s) load balancer, but health checks to port 80 on the Compute Engine virtual machine instance are failing, and no traffic is sent to your instances. You want to resolve the problem. Which commands should you run?

A. `gcloud compute instances add-access-config instance-1`  
B. `gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --destination-ranges 130.211.0.0/22,35.191.0.0/16 --direction EGRESS`  
C. `gcloud compute firewall-rules create allow-lb --network load-balancer --allow tcp --source-ranges 130.211.0.0/22,35.191.0.0/16 --direction INGRESS`  
D. `gcloud compute health-checks update http health-check --unhealthy-threshold 10`

Answer: A

QUESTION 109 You deployed a hub-and-spoke architecture in your Google Cloud environment that uses VPC Network Peering to connect the spokes to the hub. For security reasons, you deployed a private Google Kubernetes Engine (GKE) cluster in one of the spoke projects with a private endpoint for the control plane. You configured authorized networks to be the subnet range where the GKE nodes are deployed. When you attempt to reach the GKE control plane from a different spoke project, you cannot access it. You need to allow access to the GKE control plane from the other spoke projects. What should you do?

A. Add a firewall rule that allows port 443 from the other spoke projects.  
B. Enable Private Google Access on the subnet where the GKE nodes are deployed.  
C. Configure the authorized networks to be the subnet ranges of the other spoke projects.  
D. Deploy a proxy in the spoke project where the GKE nodes are deployed and connect to the control plane through the proxy.

Answer: C

QUESTION 110 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

QUESTION 111 In your Google Cloud organization, you have two folders: Dev and Prod. You want a scalable and consistent way to enforce the following firewall rules for all virtual machines (VMs) with minimal cost: Port 8080 should always be open for VMs in the projects in the Dev folder. Any traffic to port 8080 should be denied for all VMs in your projects in the Prod folder. What should you do?

A. Create and associate a firewall policy with the Dev folder with a rule to open port 8080. Create and associate a firewall policy with the Prod folder with a rule to deny traffic to port 8080.  
B. Create a Shared VPC for the Dev projects and a Shared VPC for the Prod projects. Create a VPC firewall rule to open port 8080 in the Shared VPC for Dev. Create a firewall rule to deny traffic to port 8080 in the

Shared VPC for Prod. Deploy VMs to those Shared VPCs.C. In all VPCs for the Dev projects, create a VPC firewall rule to open port 8080. In all VPCs for the Prod projects, create a VPC firewall rule to deny traffic to port 8080.D. Use Anthos Config Connector to enforce a security policy to open port 8080 on the Dev VMs and deny traffic to port 8080 on the Prod VMs.

Answer: AQUESTION 112You 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. B. C. D.

Answer: CQUESTION 113Your company's on-premises network is connected to a VPC using a Cloud VPN tunnel. You have a static route of 0.0.0.0/0 with the VPN tunnel as its next hop defined in the VPC. All internet bound traffic currently passes through the on-premises network. You configured Cloud NAT to translate the primary IP addresses of Compute Engine instances in one region. Traffic from those instances will now reach the internet directly from their VPC and not from the on-premises network. Traffic from the virtual machines (VMs) is not translating addresses as expected. What should you do?A. Lower the TCP Established Connection Idle Timeout for the NAT gateway.B. Add firewall rules that allow ingress and egress of the external NAT IP address, have a target tag that is on the Compute Engine instances, and have a priority value higher than the priority value of the default route to the VPN gateway.C. Add a default static route to the VPC with the default internet gateway as the next hop, the network tag associated with the Compute Engine instances, and a higher priority than the priority of the default route to the VPN tunnel.D. Increase the default min-ports-per-vm setting for the Cloud NAT gateway.

Answer: AQUESTION 114You are designing a Partner Interconnect hybrid cloud connectivity solution with geo-redundancy across two metropolitan areas. You want to follow Google-recommended practices to set up the following region/metro pairs:(region 1/metro 1)(region 2/metro 2)What should you do?A. Create a Cloud Router in region 1 with two VLAN attachments connected to metro1-zone1-x.Create a Cloud Router in region 2 with two VLAN attachments connected to metro1-zone2-x.B. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone1-x.Create a Cloud Router in region 2 with two VLAN attachments connected to metro2-zone2-x.C. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone2-x.Create a Cloud Router in region 2 with one VLAN attachment connected to metro2-zone2-x.D. Create a Cloud Router in region 1 with one VLAN attachment connected to metro1-zone1-x and one VLAN attachment connected to metro1-zone2-x.Create a Cloud Router in region 2 with one VLAN attachment connected to metro2-zone1-x and one VLAN attachment to metro2-zone2-x.

Answer: BQUESTION 115You are designing the network architecture for your organization. Your organization has three developer teams: Web, App, and Database. All of the developer teams require access to Compute Engine instances to perform their critical tasks. You are part of a small network and security team that needs to provide network access to the developers. You need to maintain centralized control over network resources, including subnets, routes, and firewalls. You want to minimize operational overhead. How should you design this topology?A. Configure a host project with a Shared VPC. Create service projects for Web, App, and Database.B. Configure one VPC for Web, one VPC for App, and one VPC for Database. Configure HA VPN between each VPC.C. Configure three Shared VPC host projects, each with a service project: one for Web, one for App, and one for Database.D. Configure one VPC for Web, one VPC for App, and one VPC for Database. Use VPC Network Peering to connect all VPCs in a full mesh.

Answer: CQUESTION 116Your company has 10 separate Virtual Private Cloud (VPC) networks, with one VPC per project in a single region in Google Cloud. Your security team requires each VPC network to have private connectivity to the main on-premises location via a Partner Interconnect connection in the same region. To optimize cost and operations, the same connectivity must be shared with all projects. You must ensure that all traffic between different projects, on-premises locations, and the internet can be inspected using the same third-party appliances. What should you do?A. Configure the third-party appliances with multiple interfaces and specific Partner Interconnect VLAN attachments per project. Create the relevant routes on the third-party appliances and VPC networks.B. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network. Create separate VPC networks for on-premises and internet connectivity.Create the relevant routes on the third-party appliances and VPC networks.C. Consolidate all existing projects' subnetworks into a single VPC. Create separate VPC networks for on-premises and internet connectivity. Configure the third-party appliances with multiple interfaces, with each interface connected to a separate VPC network. Create the relevant routes on the third-party appliances and VPC networks.D. Configure the third-party appliances with multiple interfaces. Create a hub VPC network for all projects, and create separate VPC networks for on-premises and internet connectivity. Create the relevant routes on the third-party appliances and VPC networks. Use VPC Network Peering to connect all projects' VPC networks to the hub VPC. Export custom routes from the hub VPC and import on all projects' VPC networks.

Answer: DQUESTION 117You 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 on-premises 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 on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

Answer: AQUESTION 118Your organization uses a hub-and-spoke architecture with critical Compute Engine instances in your Virtual Private Clouds (VPCs). You are responsible for the design of Cloud DNS in Google Cloud. You need to be able to resolve Cloud DNS private zones from your on-premises data center and enable on-premises name resolution from your hub-and-spoke VPC design. What should you do?A. Configure a private DNS zone in the hub VPC, and configure DNS forwarding to the on-premises server.Configure DNS peering from the spoke VPCs to the hub VPC.B. Configure a DNS policy in the hub VPC to allow inbound query forwarding from the spoke VPCs.Configure the spoke VPCs with a private zone, and set up DNS peering to the hub VPC.C. Configure a DNS policy in the spoke VPCs, and configure your on-premises DNS as an alternate DNS server.Configure the hub VPC with a private zone, and set up DNS peering to each of the spoke VPCs.D. Configure a DNS policy in the hub VPC, and configure the on-premises DNS as an alternate DNS server.Configure the spoke VPCs with a private zone, and set up DNS peering to the hub VPC.

Answer: CQUESTION 119You have a Cloud Storage bucket in Google Cloud project XYZ. The bucket contains sensitive data.You need to design a solution to ensure that only instances belonging to VPCs under project XYZ can access the data stored in this Cloud Storage bucket. What should you do?A. Configure Private Google Access to privately access the Cloud Storage service using private IP addresses.B. Configure a VPC Service Controls perimeter around project XYZ, and include storage.googleapis.com as a restricted service in the service perimeter.C. Configure Cloud Storage with projectPrivate Access Control List (ACL) that gives permission to the project team based on their roles.D. Configure Private Service Connect to privately access Cloud Storage from all VPCs under project XYZ.

Answer: CQUESTION 120You are maintaining a Shared VPC in a host project. Several departments within your company have infrastructure in different service projects attached to the Shared VPC and use Identity and Access Management (IAM) permissions to manage the cloud resources in those projects. VPC Network Peering is also set up between the Shared VPC and a common services VPC that is not in a service project. Several users are experiencing failed connectivity between certain instances in different Shared VPC service projects and between certain instances and the internet. You need to validate the network configuration to identify whether a misconfiguration is the root cause of the problem. What should you do?A. Review the VPC audit logs in Cloud Logging for the affected instances.B. Use Secure Shell (SSH) to connect to the affected Compute Engine instances, and run a series of PING tests to the other affected endpoints and the 8.8.8.8 IPv4 address.C. Run Connectivity Tests from Network Intelligence Center to check connectivity between the affected endpoints in your network and the internet.D. Enable VPC Flow Logs for all VPCs, and review the logs in Cloud Logging for the affected instances.

Answer: CQUESTION 121Your organization has Compute Engine instances in us-east1, us-west2, and us-central1. Your organization also has an existing Cloud Interconnect physical connection in the East Coast of the United States with a single VLAN attachment and Cloud Router in us-east1. You need to provide a design with high availability and ensure that if a region goes down, you still have access to all your other Virtual Private Cloud (VPC) subnets. You need to accomplish this in the most cost-effective manner possible. What should you do?A. Configure your VPC routing in regional mode. Add an additional Cloud Interconnect VLAN attachment in the us-east1 region, and configure a Cloud Router in us-east1.B. Configure your VPC routing in global mode.Add an additional Cloud Interconnect VLAN attachment in the us-east1 region, and configure a Cloud Router in us-east1.C. Configure your VPC routing in global mode.Add an additional Cloud Interconnect VLAN attachment in the us-west2 region, and configure a Cloud Router in us-west2.D. Configure your VPC routing in regional mode. Add additional Cloud Interconnect VLAN attachments in the us-west2 and us-central1 regions, and configure Cloud Routers in us-west2 and us-central1.

Answer: BQUESTION 122You 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 identify the web application firewall (WAF) rule that is incorrectly blocking traffic. What should you do?A. Enable firewall logs, and view the logs in Firewall Insights.B. Enable HTTP(S) Load Balancing logging with sampling rate equal to 1, and view the logs in Cloud Logging.C. Enable VPC Flow Logs, and view the logs in Cloud Logging.D. Enable Google Cloud Armor audit logs, and view the logs on the Activity page in the Google Cloud Console.

Answer: AQUESTION 123You are the Organization Admin for your company. One of your engineers is responsible for setting up multiple host projects across multiple folders and sharing subnets with service projects. You need to enable the engineer's Identity and Access Management (IAM) configuration to

complete their task in the fewest number of steps. What should you do?A. Set up the engineer with Compute Shared VPC Admin IAM role at the folder level.B. Set up the engineer with Compute Shared VPC Admin IAM role at the organization level.C. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the folder level.D. Set up the engineer with Compute Shared VPC Admin IAM role and Project IAM Admin role at the organization level.**Answer: B**QUESTION 124You 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**QUESTION 125You are designing a new global application using Compute Engine instances that will be exposed by a global HTTP(S) load balancer. You need to secure your application from distributed denial-of-service and application layer (layer 7) attacks. What should you do?A. Configure VPC Service Controls and create a secure perimeter. Define fine-grained perimeter controls and enforce that security posture across your Google Cloud services and projects.B. Configure a Google Cloud Armor security policy in your project, and attach it to the backend service to secure the application.C. Configure VPC firewall rules to protect the Compute Engine instances against distributed denial-of-service attacks.D. Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at the organization level.**Answer: C**QUESTION 126Your organization's security policy requires that all internet-bound traffic return to your on-premises data center through HA VPN tunnels before egressing to the internet, while allowing virtual machines (VMs) to leverage private Google APIs using private virtual IP addresses 199.36.153.4/30. You need to configure the routes to enable these traffic flows. What should you do?A. Configure a custom route 0.0.0.0/0 with a priority of 500 whose next hop is the default internet gateway. Configure another custom route 199.36.153.4/30 with priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.B. Configure a custom route 0.0.0.0/0 with a priority of 1000 whose next hop is the internet gateway. Configure another custom route 199.36.153.4/30 with a priority of 500 whose next hop is the VPN tunnel back to the on-premises data center.C. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 1000. Configure a custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the default internet gateway.D. Announce a 0.0.0.0/0 route from your on-premises router with a MED of 500. Configure another custom route 199.36.153.4/30 with a priority of 1000 whose next hop is the VPN tunnel back to the on-premises data center.**Answer: A**QUESTION 127Your company has defined a resource hierarchy that includes a parent folder with subfolders for each department. Each department defines their respective project and VPC in the assigned folder and has the appropriate permissions to create Google Cloud firewall rules. The VPCs should not allow traffic to flow between them. You need to block all traffic from any source, including other VPCs, and delegate only the intra-VPC firewall rules to the respective departments. What should you do?A. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 0.B. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 1000.C. Create two hierarchical firewall policies per department's folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to allow, and another lower-priority rule that blocks traffic from any other source.D. Create two hierarchical firewall policies per department's folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to goto\_next, and another lower-priority rule that blocks traffic from any other source.**Answer: B**QUESTION 128You 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**QUESTION 129You are configuring an HA VPN connection between your Virtual Private Cloud (VPC) and on-premises network. The VPN gateway is named VPN\_GATEWAY\_1. You need to restrict VPN tunnels created in the project to only connect to your on-premises VPN public IP address: 203.0.113.1/32.What should you do?A. Configure a firewall rule accepting 203.0.113.1/32, and set a target tag equal to VPN\_GATEWAY\_1.B. Configure the Resource Manager constraint constraints/compute.restrictVpnPeerIPs to use an allowList consisting of only the 203.0.113.1/32 address.C. Configure a Google Cloud Armor security policy, and create a policy rule to allow 203.0.113.1/32.D. Configure an access control list on the peer VPN gateway to deny all traffic except 203.0.113.1/32, and attach it to the primary external interface.**Answer: B**QUESTION 130Your company has recently installed a Cloud VPN tunnel between your on-premises data center and your Google Cloud Virtual Private Cloud (VPC). You need to configure access to the Cloud Functions API for your on-premises servers. The configuration

must meet the following requirements:- Certain data must stay in the project where it is stored and not be exfiltrated to other projects.- Traffic from servers in your data center with RFC 1918 addresses do not use the internet to access Google Cloud APIs.- All DNS resolution must be done on-premises.The solution should only provide access to APIs that are compatible with VPC Service Controls.What should you do?A. Create an A record for private.googleapis.com using the 199.36.153.8/30 address range. Create a CNAME record for \*.googleapis.com that points to the A record.Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.Remove the default internet gateway from the VPC where your Cloud VPN tunnel terminates.B. Create an A record for restricted.googleapis.com using the 199.36.153.4/30 address range.Create a CNAME record for \*.googleapis.com that points to the A record.Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.Configure your on-premises firewalls to allow traffic to the restricted.googleapis.com addresses.C. Create an A record for restricted.googleapis.com using the 199.36.153.4/30 address range. Create a CNAME record for \*.googleapis.com that points to the A record.Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.Remove the default internet gateway from the VPC where your Cloud VPN tunnel terminates.D. Create an A record for private.googleapis.com using the 199.36.153.8/30 address range.Create a CNAME record for \*.googleapis.com that points to the A record.Configure your on-premises routers to use the Cloud VPN tunnel as the next hop for the addresses you used in the A record.Configure your on-premises firewalls to allow traffic to the private.googleapis.com addresses.

Answer: CQUESTION 131You need to configure a Google Kubernetes Engine (GKE) cluster. The initial deployment should have 5 nodes with the potential to scale to 10 nodes. The maximum number of Pods per node is 8. The number of services could grow from 100 to up to 1024. How should you design the IP schema to optimally meet this requirement?A. Configure a /28 primary IP address range for the node IP addresses. Configure a (25 secondary IP range for the Pods. Configure a /22 secondary IP range for the Services.B. Configure a /28 primary IP address range for the node IP addresses. Configure a /25 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.C. Configure a /28 primary IP address range for the node IP addresses. Configure a /28 secondary IP range for the Pods. Configure a /21 secondary IP range for the Services.D. Configure a /28 primary IP address range for the node IP addresses. Configure a /24 secondary IP range for the Pads. Configure a /22 secondary IP range for the Services.

Answer: AQUESTION 132You are migrating a three-tier application architecture from on-premises to Google Cloud. As a first step in the migration, you want to create a new Virtual Private Cloud (VPC) with an external HTTP(S) load balancer. This load balancer will forward traffic back to the on-premises compute resources that run the presentation tier. You need to stop malicious traffic from entering your VPC and consuming resources at the edge, so you must configure this policy to filter IP addresses and stop cross-site scripting (XSS) attacks. What should you do?A. Create a Google Cloud Armor policy, and apply it to a backend service that uses an unmanaged instance group backend.B. Create a hierarchical firewall ruleset, and apply it to the VPC's parent organization resource node.C. Create a Google Cloud Armor policy, and apply it to a backend service that uses an internet network endpoint group (NEG) backend.D. Create a VPC firewall ruleset, and apply it to all instances in unmanaged instance groups.

Answer: CQUESTION 133You just finished your company's migration to Google Cloud and configured an architecture with 3 Virtual Private Cloud (VPC) networks: one for Sales, one for Finance, and one for Engineering. Every VPC contains over 100 Compute Engine instances, and now developers using instances in the Sales VPC and the Finance VPC require private connectivity between each other. You need to allow communication between Sales and Finance without compromising performance or security. What should you do?A. Configure an HA VPN gateway between the Finance VPC and the Sales VPC.B. Configure the instances that require communication between each other with an external IP address.C. Create a VPC Network Peering connection between the Finance VPC and the Sales VPC.D. Configure Cloud NAT and a Cloud Router in the Sales and Finance VPCs.

Answer: CQUESTION 134You have provisioned a Partner Interconnect connection to extend connectivity from your on-premises data center to Google Cloud. You need to configure a Cloud Router and create a VLAN attachment to connect to resources inside your VPC. You need to configure an Autonomous System number (ASN) to use with the associated Cloud Router and create the VLAN attachment.What should you do?A. Use a 4-byte private ASN 4200000000-4294967294.B. Use a 2-byte private ASN 64512-65535.C. Use a public Google ASN 15169.D. Use a public Google ASN 16550.

Answer: BQUESTION 135You are configuring a new application that will be exposed behind an external load balancer with both IPv4 and IPv6 addresses and support TCP pass-through on port 443. You will have backends in two regions: us-west1 and us-east1. You want to serve the content with the lowest possible latency while ensuring high availability and autoscaling. Which configuration should you use?A. Use global SSL Proxy Load Balancing with backends in both regions.B. Use global TCP Proxy Load Balancing with backends in both regions.C. Use global external HTTP(S) Load Balancing with backends in both regions.D. Use Network Load Balancing in both regions, and use DNS-based load balancing to direct traffic to the closest region.

Answer: DQUESTION 136In your project my-project, you have two subnets in a Virtual Private Cloud (VPC):



subnet-a with IP range 10.128.0.0/20 and subnet-b with IP range 172.16.0.0/24. You need to deploy database servers in subnet-A. You will also deploy the application servers and web servers in subnet-b. You want to configure firewall rules that only allow database traffic from the application servers to the database servers. What should you do? A. Create network tag app-server and service account sa-db@my-project.iam.gserviceaccount.com. Add the tag to the application servers, and associate the service account with the database servers. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules top:3306 --source-tags app-server --target-service-accounts sa-db@my-project.iam.gserviceaccount.com` B. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com. Associate service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-service-accounts sa-app@my-project.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com` C. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.com. Associate the service account sa-app with the application servers, and associate the service account sa-db with the database servers. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-ranges 10.128.0.0/20 --source-service-accounts sa-app@my-project.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com` D. Create network tags app-server and db-server. Add the app-server tag to the application servers, and add the db-server tag to the database servers. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules tcp:3306 --source-ranges 10.128.0.0/20 --source-tags app-server --target-tags db-server` Answer: D

QUESTION 137 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

QUESTION 138 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

QUESTION 139 You successfully provisioned a single Dedicated Interconnect. The physical connection is at a colocation facility closest to us-west2. Seventy-five percent of your workloads are in us-east4, and the remaining twenty-five percent of your workloads are in us-central1. All workloads have the same network traffic profile. You need to minimize data transfer costs when deploying VLAN attachments. What should you do? A. Keep the existing Dedicated interconnect. Deploy a VLAN attachment to a Cloud Router in us-west2, and use VPC global routing to access workloads in us-east4 and us-central1. B. Keep the existing Dedicated Interconnect. Deploy a VLAN attachment to a Cloud Router in us-east4, and deploy another VLAN attachment to a Cloud Router in us-central1. C. Order a new Dedicated Interconnect for a colocation facility closest to us-east4, and use VPC global routing to access workloads in us-central1. D. Order a new Dedicated Interconnect for a colocation facility closest to us-central1, and use VPC global routing to access workloads in us-east4. Answer: C

QUESTION 140 You are designing a hybrid cloud environment. Your Google Cloud environment is interconnected with your on-premises network using HA VPN and Cloud Router in a central transit hub VPC. The Cloud Router is configured with the default settings. Your on-premises DNS server is located at 192.168.20.88. You need to ensure that your Compute Engine resources in multiple spoke VPCs can resolve on-premises private hostnames using the domain corp.altostrat.com while also resolving Google Cloud hostnames. You want to follow Google-recommended practices. What should you do? A. Create a private forwarding zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. Configure VPC peering in the spoke VPCs to peer with the hub VPC. B. Create a private forwarding zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. C. Create a private forwarding zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com that points to



192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. Create a hub-and-spoke VPN deployment in each spoke VPC to connect back to the on-premises network directly. D. Create a private forwarding zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com that points to 192.168.20.88. Associate the zone with the hub VPC. Create a private peering zone in Cloud DNS for `corp.altostrat.com` called corp-altostrat-com associated with the spoke VPCs, with the hub VPC as the target. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19. Create a hub and spoke VPN deployment in each spoke VPC to connect back to the hub VPC. Answer: AQUESTION 141 You have the following firewall ruleset applied to all instances in your Virtual Private Cloud (VPC):

Direction	Action
egress	deny
egress	deny
ingress	allow

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

Direction	Action	Address range	Port	Logging
egress	deny	192.0.2.42/32	80	True

You are using a new user account. You must assign the appropriate identity and Access Management (IAM) user roles to this new user account before updating the firewall rule. The new user account must be able to apply the update and view firewall logs. What should you do? A. Assign the compute.securityAdmin and logging.viewer role to the new user account. Apply the new firewall rule with a priority of 50. B. Assign the compute.securityAdmin and logging.bucketWriter role to the new user account. Apply the new firewall rule with a priority of 150. C. Assign the compute.orgSecurityPolicyAdmin and logging.viewer role to the new user account. Apply the new firewall rule with a priority of 50. D. Assign the compute.orgSecurityPolicyAdmin and logging.bucketWriter role to the new user account. Apply the new firewall rule with a priority of 150. Answer: AQUESTION 142

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: BQUESTION 143

Your company has provisioned 2000 virtual machines (VMs) in the private subnet of your Virtual Private Cloud (VPC) in the us-east1 region. You need to configure each VM to have a minimum of 128 TCP connections to a public repository so that users can download software updates and packages over the internet. You need to implement a Cloud NAT gateway so that the VMs are able to perform outbound NAT to the internet. You must ensure that all VMs can simultaneously connect to the public repository and download software updates and packages. Which two methods can you use to accomplish this? (Choose two.) A. Configure the NAT gateway in manual allocation mode, allocate 2 NAT IP addresses, and update the minimum number of ports per VM to 256. B. Create a second Cloud NAT gateway with the default minimum number of ports configured per VM to 64. C. Use the default Cloud NAT gateway's NAT proxy to dynamically scale using a single NAT IP address. D. Use the default Cloud NAT gateway to automatically scale to the required number of NAT IP addresses, and update the minimum number of ports per VM to 128. E. Configure the NAT gateway in manual allocation mode, allocate 4 NAT IP addresses, and update the minimum number of ports per VM to 128. Answer: ABResources From: 1.2022 Latest

Braindump2go Professional-Cloud-Network-Engineer Exam Dumps (PDF & VCE) Free Share:

<https://www.braindump2go.com/professional-cloud-network-engineer.html> 2.2022 Latest Braindump2go

Professional-Cloud-Network-Engineer PDF and VCE Dumps Free Share:

<https://drive.google.com/drive/folders/17dsDXhDZ-V4yNHGCxaK8CbOU8XqEB9AL?usp=sharing> 3.2022 Free Braindump2go

Professional-Cloud-Network-Engineer Exam Questions Download:

[https://www.braindump2go.com/free-online-pdf/Professional-Cloud-Network-Engineer-PDF-Dumps\(95-143\).pdf](https://www.braindump2go.com/free-online-pdf/Professional-Cloud-Network-Engineer-PDF-Dumps(95-143).pdf) Free Resources

from Braindump2go, We Devoted to Helping You 100% Pass All Exams!