



Vendor: Google

Exam Code: Professional-Cloud-Network-Engineer

Exam Name: Professional Cloud Network Engineer

Version: DEMO

QUESTION 1

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

How should you deploy this service in GCP?

- A. Create a managed instance group from one of the images of the on-premises servers, and link this instance group to a target pool behind your load balancer.
- B. Create a target pool, add all backend instances to this target pool, and deploy the target pool behind your load balancer.
- C. Deploy a third-party virtual appliance as frontend to these servers that will accommodate the significant differences between these backend servers.
- D. Use GCP's ECMP capability to load-balance traffic to the backend servers by installing multiple equal- priority static routes to the backend servers.

Answer: B

Explanation:

External TCP/UDP Network Load Balancing can use either a backend service or a target pool to define the group of backend instances that receive incoming traffic.

Target pools work with forwarding rules that handle TCP and UDP traffic. You must create a target pool before you can use it with a forwarding rule.

<https://cloud.google.com/load-balancing/docs/target-pools>

QUESTION 2

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

What is the most likely cause of this problem?

- A. The instance has been configured with multiple interfaces.
- B. An external IP address has been configured on the instance.
- C. You have created static routes that use RFC1918 ranges.
- D. The instance is accessible by a load balancer external IP address.

Answer: B

Explanation:

The existence of an external IP address on an interface always takes precedence and always performs one-to-one NAT, without using Cloud NAT.

<https://cloud.google.com/nat/docs/overview#specifications>

QUESTION 3

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

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

- A. AS-Path
- B. Community
- C. Local Preference

D. Multi-exit Discriminator

Answer: D

Explanation:

You can configure 2 different MED values for each BGP neighbor in your single on-prem router, to influence ISP(GCP)'s 2 separate routers to select which path they send traffic towards you. The lower MED value is preferred.

<https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13759-37.html>

QUESTION 4

You are increasing your usage of Cloud VPN between on-premises and GCP, and you want to support more traffic than a single tunnel can handle. You want to increase the available bandwidth using Cloud VPN.

What should you do?

- A. Double the MTU on your on-premises VPN gateway from 1460 bytes to 2920 bytes.
- B. Create two VPN tunnels on the same Cloud VPN gateway that point to the same destination VPN gateway IP address.
- C. Add a second on-premises VPN gateway with a different public IP address.
Create a second tunnel on the existing Cloud VPN gateway that forwards the same IP range, but points at the new on-premises gateway IP.
- D. Add a second Cloud VPN gateway in a different region than the existing VPN gateway.
Create a new tunnel on the second Cloud VPN gateway that forwards the same IP range, but points to the existing on-premises VPN gateway IP address.

Answer: C

Explanation:

Option 1: Scale the on-premises VPN gateway

<https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#option-1>

QUESTION 5

You are disabling DNSSEC for one of your Cloud DNS-managed zones. You removed the DS records from your zone file, waited for them to expire from the cache, and disabled DNSSEC for the zone. You receive reports that DNSSEC validating resolves are unable to resolve names in your zone.

What should you do?

- A. Update the TTL for the zone.
- B. Set the zone to the TRANSFER state.
- C. Disable DNSSEC at your domain registrar.
- D. Transfer ownership of the domain to a new registrar.

Answer: C

Explanation:

Before disabling DNSSEC for a managed zone you want to use, you must deactivate DNSSEC at your domain registrar to ensure that DNSSEC-validating resolvers can still resolve names in the zone.

<https://cloud.google.com/dns/docs/dnssec-config>

QUESTION 6

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

During troubleshooting you find:

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

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

- A. The traffic is matching the expected ingress rule.
- B. The traffic is matching the expected egress rule.
- C. The traffic is not matching the expected ingress rule.
- D. The traffic is not matching the expected egress rule.

Answer: C

Explanation:

Ingress packets are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs.

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

QUESTION 7

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

What is the most likely cause of the problem?

- A. You have not configured compression in Cloud CDN.
- B. You have configured the web servers and Cloud CDN with different compression types.
- C. The web servers behind the load balancer are configured with different compression types.
- D. You have to configure the web servers to compress responses even if the request has a Via header.

Answer: D

Explanation:

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

<https://cloud.google.com/cdn/docs/troubleshooting-steps>

QUESTION 8

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

What should you do?

- A. Configure a policy-based route rule to prioritize the traffic.
- B. Configure an HTTP load balancer, and direct the traffic to it.
- C. Configure Dynamic Routing for the subnet hosting the application.
- D. Configure the TTL for the DNS zone to decrease the time between updates.

Answer: B

Explanation:

HTTP LB - backends can be in any region and any VPC network (Premium tier). On the other hand, Network LB - the backend service must also be in the same region and VPC network as the forwarding rule.

QUESTION 9

You have an application running on Compute Engine that uses BigQuery to generate some results that are stored in Cloud Storage. You want to ensure that none of the application instances have external IP addresses.

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

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

Answer: AE

Explanation:

Private Google Access interaction

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

Specifications

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

QUESTION 10

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

How should you design this topology?

- A. Create 2 shared VPCs within the shared VPC Host Project, and enable VPC peering between them.
Use firewall rules to filter access between the specific networks.
- B. Create 2 shared VPCs within the shared VPC Host Project, and create a Cloud VPN/Cloud Router between them.

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

Answer: D

Explanation:

Building on the initial reference architecture, Shared VPC host projects and multiple service projects let administrators delegate administrative responsibilities—such as creating and managing instances—to Service Project Admins while maintaining centralized control over network resources like subnets, routes, and firewalls.

<https://cloud.google.com/solutions/best-practices-vpc-design#single-host-project-multiple-service-projects-single-shared-vpc>

QUESTION 11

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:

GCP recommends creating VPC peering for establishing communication between two organizations in GCP.

QUESTION 12

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

What should you do?

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

Answer: D

Explanation:

After you convert an auto mode network to custom mode, you must review all API calls and gcloud commands that implicitly reference any subnet that was automatically created while the network was in auto mode. API calls and commands will need to be modified so that they reference the subnet explicitly.

<https://cloud.google.com/vpc/docs/using-vpc#switch-network-mode>

QUESTION 13

You have recently been put in charge of managing identity and access management for your organization. You have several projects and want to use scripting and automation wherever possible. You want to grant the editor role to a project member.

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

- A. GetIamPolicy() via REST API
- B. setIamPolicy() via REST API
- C. `gcloud pubsub add-iam-policy-binding Sprojectname --member user:Susername -- role roles/editor`
- D. `gcloud projects add-iam-policy-binding Sprojectname --member user:Susername --role roles/editor`
- E. Enter an email address in the Add members field, and select the desired role from the drop-down menu in the GCP Console.

Answer: DE

Explanation:

<https://cloud.google.com/iam/docs/granting-changing-revoking-access#granting-gcloud-manual>

<https://cloud.google.com/iam/docs/granting-changing-revoking-access#access-control-via-console>

QUESTION 14

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:

As the question states that the RTT is 100ms thus low transfer rate is due to the TCP window size that is too small. And the solution is to increase the window size.

QUESTION 15

You work for a multinational enterprise that is moving to GCP.

These are the cloud requirements:

- An on-premises data center located in the United States in Oregon and New York with Dedicated Interconnects connected to Cloud regions us-west1 (primary HQ) and us-east4 (backup)
- Multiple regional offices in Europe and APAC
- Regional data processing is required in europe-west1 and australia-southeast1
- Centralized Network Administration Team

Your security and compliance team requires a virtual inline security appliance to perform L7 inspection for URL filtering. You want to deploy the appliance in us-west1.

What should you do?

- A. Create 2 VPCs in a Shared VPC Host Project.
Configure a 2-NIC instance in zone us-west1-a in the Host Project.
Attach NIC0 in VPC #1 us-west1 subnet of the Host Project.
Attach NIC1 in VPC #2 us-west1 subnet of the Host Project.
Deploy the instance.
Configure the necessary routes and firewall rules to pass traffic through the instance.
- B. Create 2 VPCs in a Shared VPC Host Project.
Configure a 2-NIC instance in zone us-west1-a in the Service Project.
Attach NIC0 in VPC #1 us-west1 subnet of the Host Project.
Attach NIC1 in VPC #2 us-west1 subnet of the Host Project.
Deploy the instance.
Configure the necessary routes and firewall rules to pass traffic through the instance.
- C. Create 1 VPC in a Shared VPC Host Project.
Configure a 2-NIC instance in zone us-west1-a in the Host Project.
Attach NIC0 in us-west1 subnet of the Host Project.
Attach NIC1 in us-west1 subnet of the Host Project
Deploy the instance.
Configure the necessary routes and firewall rules to pass traffic through the instance.
- D. Create 1 VPC in a Shared VPC Service Project.
Configure a 2-NIC instance in zone us-west1-a in the Service Project.
Attach NIC0 in us-west1 subnet of the Service Project.
Attach NIC1 in us-west1 subnet of the Service Project ?Deploy the instance.
Configure the necessary routes and firewall rules to pass traffic through the instance.

Answer: A

Explanation:

You cannot attach 2 NICs of same appliance to same VPC. The two NICs must be attached to different VPCs.

Each interface is attached to a different VPC network, giving that instance access to different VPC networks in Google Cloud Platform (GCP). You cannot attach multiple network interfaces to the same VPC network.

<https://cloud.google.com/vpc/docs/create-use-multiple-interfaces>

QUESTION 16

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

How should you design this topology?

- A. Create a subnet of size/25 with 2 secondary ranges of: /17 for Pods and /21 for Services. Create a VPC-native cluster and specify those ranges.
- B. Create a subnet of size/28 with 2 secondary ranges of: /24 for Pods and /24 for Services. Create a VPC-native cluster and specify those ranges. When the services are ready to be deployed, resize the subnets.
- C. Use gcloud container clusters create [CLUSTER NAME]--enable-ip-alias to create a VPC-native cluster.
- D. Use gcloud container clusters create [CLUSTER NAME] to create a VPC-native cluster.

Answer: A

Explanation:

When you create a VPC-native cluster, you specify a subnet in a VPC network. The cluster uses three unique subnet IP address ranges:

It uses the subnet's primary IP address range for all node IP addresses.

It uses one secondary IP address range for all Pod IP addresses.

It uses another secondary IP address range for all Service (cluster IP) addresses.

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

QUESTION 17

Your company has recently expanded their EMEA-based operations into APAC. Globally distributed users report that their SMTP and IMAP services are slow. Your company requires end-to-end encryption, but you do not have access to the SSL certificates.

Which Google Cloud load balancer should you use?

- A. SSL proxy load balancer
- B. Network load balancer
- C. HTTPS load balancer
- D. TCP proxy load balancer

Answer: D

Explanation:

<https://cloud.google.com/load-balancing/docs/choosing-load-balancer>

SSL offload yes >> SSL proxy

SSL offload no >> TCP proxy

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



10% Discount Coupon Code: ASTR14