

PROFESSIONAL-CLOUD-NETWORK-ENGINEER^{Q&As}

Professional Cloud Network Engineer

Pass Google PROFESSIONAL-CLOUD-NETWORK-ENGINEER Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/professional-cloud-network-engineer.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Google
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

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.

Correct Answer: D

Reference: <https://cloud.google.com/vpc/docs/shared-vpc>

QUESTION 2

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.

Correct Answer: B

QUESTION 3

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 perfdiaggs://[BUCKET_NAME].

Correct Answer: A

QUESTION 4

You are creating a new application and require access to Cloud SQL from VPC instances without public IP addresses.

Which two actions should you take? (Choose two.)

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

Correct Answer: AC

Reference: <https://cloud.google.com/sql/docs/mysql/private-ip>

QUESTION 5

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

How should you design this topology?

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

Correct Answer: D

Explanation:

VPC Network Peering enables you to peer VPC networks so that workloads in different VPC networks can communicate in private RFC 1918 space. Traffic stays within Google's network and doesn't traverse the public internet.

Reference: <https://cloud.google.com/vpc/docs/vpc-peering>

QUESTION 6

You need to give each member of your network operations team least-privilege access to create, modify, and delete Cloud Interconnect VLAN attachments.

What should you do?

- A. Assign each user the editor role.
- B. Assign each user the compute.networkAdmin role.
- C. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get.
- D. Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get, compute.routers.create, compute.routers.get, compute.routers.update.

Correct Answer: C

QUESTION 7

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

What should you do?

- A. Create a Cloud VPN instance. Create a policy-based VPN tunnel per subnet. Configure the appropriate local and remote traffic selectors to match your local and remote networks. Create the appropriate static routes.
- B. Create a Cloud VPN instance. Create a policy-based VPN tunnel. Configure the appropriate local and remote traffic selectors to match your local and remote networks. Configure the appropriate static routes.
- C. Create a Cloud VPN instance. Create a route-based VPN tunnel. Configure the appropriate local and remote traffic selectors to match your local and remote networks. Configure the appropriate static routes.
- D. Create a Cloud VPN instance. Create a route-based VPN tunnel. Configure the appropriate local and remote traffic selectors to 0.0.0.0/0. Configure the appropriate static routes.

Correct Answer: D

Reference: <https://cloud.google.com/vpn/docs/concepts/choosing-networks-routing>

QUESTION 8

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

How should you deploy this service in GCP?

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

Correct Answer: B

Reference: <https://cloud.google.com/compute/docs/instance-groups/adding-an-instance-group-to-a-loadbalancer>

QUESTION 9

You need to define an address plan for a future new GKE cluster in your VPC. This will be a VPC-native cluster, and the default Pod IP range allocation will be used. You must pre-provision all the needed VPC subnets and their respective IP address ranges before cluster creation. The cluster will initially have a single node, but it will be scaled to a maximum of three nodes if necessary. You want to allocate the minimum number of Pod IP addresses.

Which subnet mask should you use for the Pod IP address range?

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

Correct Answer: D

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

QUESTION 10

You have an application that is running in a managed instance group. Your development team has released an updated instance template which contains a new feature which was not heavily tested. You want to minimize impact to users if there is a bug in the new template.

How should you update your instances?

- A. Manually patch some of the instances, and then perform a rolling restart on the instance group.
- B. Using the new instance template, perform a rolling update across all instances in the instance group. Verify the new feature once the rollout completes.
- C. Deploy a new instance group and canary the updated template in that group. Verify the new feature in the new canary instance group, and then update the original instance group.
- D. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Correct Answer: C

Reference: <https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managedinstances>

QUESTION 11

You are deploying a global external TCP load balancing solution and want to preserve the source IP address of the original layer 3 payload.

Which type of load balancer should you use?

- A. HTTP(S) load balancer
- B. Network load balancer
- C. Internal load balancer
- D. TCP/SSL proxy load balancer

Correct Answer: B

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

QUESTION 12

Your on-premises data center has 2 routers connected to your GCP through a VPN on each router. All applications are working correctly; however, all of the traffic is passing across a single VPN instead of being load-balanced across the 2 connections as desired.

During troubleshooting you find:

1.

Each on-premises router is configured with the same ASN.

2.

Each on-premises router is configured with the same routes and priorities.

3.

Both on-premises routers are configured with a VPN connected to a single Cloud Router.

4.

The VPN logs have no-proposal-chosen lines when the VPNs are connecting.

5.

BGP session is not established between one on-premises router and the Cloud Router.

What is the most likely cause of this problem?

- A. One of the VPN sessions is configured incorrectly.
- B. A firewall is blocking the traffic across the second VPN connection.
- C. You do not have a load balancer to load-balance the network traffic.
- D. BGP sessions are not established between both on-premises routers and the Cloud Router.

Correct Answer: C

QUESTION 13

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.

Correct 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.

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

QUESTION 14

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`.
- D. Check the email for the account of the NOC contact that you specified during the ordering process.
- E. Contact your cross-connect provider and inform them that Google automatically sent the LOA/CFA to them via email, and to complete the connection.

Correct Answer: DE

QUESTION 15

Your on-premises data center has 2 routers connected to your Google Cloud environment through a VPN on each router. All applications are working correctly; however, all of the traffic is passing across a single VPN instead of being load-balanced across the 2 connections as desired.

During troubleshooting you find:

1.

Each on-premises router is configured with a unique ASN.

2.

Each on-premises router is configured with the same routes and priorities.

3.

Both on-premises routers are configured with a VPN connected to a single Cloud Router.

4.

BGP sessions are established between both on-premises routers and the Cloud Router.

5.

Only 1 of the on-premises router's routes are being added to the routing table.

What is the most likely cause of this problem?

- A. The on-premises routers are configured with the same routes.
- B. A firewall is blocking the traffic across the second VPN connection.

C. You do not have a load balancer to load-balance the network traffic.

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

Correct Answer: D

[PROFESSIONAL-CLOUD-NETWORK-ENGINEER VCE Dumps](#)

[PROFESSIONAL-CLOUD-NETWORK-ENGINEER Practice Test](#)

[PROFESSIONAL-CLOUD-NETWORK-ENGINEER Braindumps](#)