

DUMPSBOSS.COM

Professional Cloud Network Engineer

Google Professional-Cloud-Network-Engineer

Version Demo

Total Demo Questions: 10

Total Premium Questions: 81

Buy Premium PDF

<https://dumpsboss.com>

support@dumpsboss.com

dumpsboss.com

QUESTION NO: 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.

ANSWER: D**Explanation:**

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

QUESTION NO: 2

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

QUESTION NO: 3

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: B E

QUESTION NO: 4

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

Explanation:

Reference: <https://geekflare.com/gcp-firewall-configuration/>

QUESTION NO: 5

Your company just completed the acquisition of Altostrat (a current GCP customer). Each company has a separate organization in GCP and has implemented a custom DNS solution. Each organization will retain its current domain and host names until after a full transition and architectural review is done in one year. These are the assumptions for both GCP environments.

- Each organization has enabled full connectivity between all of its projects by using Shared VPC.
- Both organizations strictly use the 10.0.0.0/8 address space for their instances, except for bastion hosts (for accessing the instances) and load balancers for serving web traffic.
- There are no prefix overlaps between the two organizations.
- Both organizations already have firewall rules that allow all inbound and outbound traffic from the 10.0.0.0/8 address space.
- Neither organization has Interconnects to their on-premises environment.

You want to integrate networking and DNS infrastructure of both organizations as quickly as possible and with minimal downtime.

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

- A. Provision Cloud Interconnect to connect both organizations together.
- B. Set up some variant of DNS forwarding and zone transfers in each organization.
- C. Connect VPCs in both organizations using Cloud VPN together with Cloud Router.
- D. Use Cloud DNS to create A records of all VMs and resources across all projects in both organizations.
- E. Create a third organization with a new host project, and attach all projects from your company and Altostrat to it using shared VPC.

ANSWER: C D

QUESTION NO: 6

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: C D

Explanation:

Reference: <https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations>

QUESTION NO: 7

In your company, two departments with separate GCP projects (code-dev and data-dev) in the same organization need to allow full cross-communication between all of their virtual machines in GCP. Each department has one VPC in its project and wants full control over their network. Neither department intends to recreate its existing computing resources. You want to implement a solution that minimizes cost.

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

- A. Connect both projects using Cloud VPN.
- B. Connect the VPCs in project code-dev and data-dev using VPC Network Peering.
- C. Enable Shared VPC in one project (e. g., code-dev), and make the second project (e. g., data-dev) a service project.
- D. Enable firewall rules to allow all ingress traffic from all subnets of project code-dev to all instances in project data-dev, and vice versa.
- E. Create a route in the code-dev project to the destination prefixes in project data-dev and use nexthop as the default gateway, and vice versa.

ANSWER: C E

QUESTION NO: 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.

ANSWER: B

Explanation:

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

QUESTION NO: 9

You have enabled HTTP(S) load balancing for your application, and your application developers have reported that HTTP(S) requests are not being distributed correctly to your Compute Engine Virtual Machine instances. You want to find data about how the request are being distributed.

Which two methods can accomplish this? (Choose two.)

- A.** On the Load Balancer details page of the GCP Console, click on the Monitoring tab, select your backend service, and look at the graphs.
- B.** In Stackdriver Error Reporting, look for any unacknowledged errors for the Cloud Load Balancers service.
- C.** In Stackdriver Monitoring, select Resources > Metrics Explorer and search for https/request_bytes_count metric.
- D.** In Stackdriver Monitoring, select Resources > Google Cloud Load Balancers and review the Key Metrics graphs in the dashboard.
- E.** In Stackdriver Monitoring, create a new dashboard and track the https/backend_request_count metric for the load balancer.

ANSWER: A D

QUESTION NO: 10

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