



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.pass4itsure.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



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/professional-cloud-network-engineer.html>

2024 Latest pass4itsure PROFESSIONAL-CLOUD-NETWORK-ENGINEER

PDF and VCE dumps Download

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



**QUESTION 1**

You want Cloud CDN to serve the <https://www.example.com/images/spacetime.png> static image file that is hosted in a private Cloud Storage bucket. You are using the VSE ORIG.-X_NZADERS cache mode. You receive an HTTP 403 error when opening the file in your browser and you see that the HTTP response has a Cache-control: private, max-age=0 header. How should you correct this issue?

- A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role
- B. Change the cache mode to cache all content.
- C. Increase the default time-to-live (TTL) for the backend service.
- D. Enable negative caching for the backend bucket

Correct Answer: A

The correct answer is A. Configure a Cloud Storage bucket permission that gives the Storage Legacy Object Reader role.

This answer is based on the following facts:

Cloud CDN can serve private content from Cloud Storage buckets, but you need to grant the appropriate permissions to the Google-managed service account that represents your load balancer¹.

The Storage Legacy Object Reader role grants read access to objects in a bucket².

The Cache-control: private header indicates that the object is not publicly readable and requires authentication³.

The USE_ORIGIN_HEADERS cache mode instructs Cloud CDN to cache responses based on the Cache-Control and Expires headers from the origin server⁴. Changing the cache mode, increasing the TTL, or enabling negative caching will

not affect the 403 error.

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: C

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



QUESTION 3

You are designing an IP address scheme for new private Google Kubernetes Engine (GKE) clusters. Due to IP address exhaustion of the RFC 1918 address space in your enterprise, you plan to use privately used public IP space for the new clusters. You want to follow Google-recommended practices. What should you do after designing your IP scheme?

- A. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters. Re-use the secondary address range for the pods across multiple private GKE clusters
- B. Create the minimum usable RFC 1918 primary and secondary subnet IP ranges for the clusters. Re-use the secondary address range for the services across multiple private GKE clusters
- C. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected and
- D. Create privately used public IP primary and secondary subnet ranges for the clusters. Create a private GKE cluster with the following options selected --disable-default-snat, --enable-ip-alias, and --enable-private-nodes

Correct Answer: D

This answer follows the Google-recommended practices for using privately used public IP (PUIP) addresses for GKE Pod address blocks¹. The benefits of this approach are:

It allows you to use any public IP addresses that are not owned by Google or your organization for your Pods, which can help mitigate address exhaustion in your enterprise.

It prevents any external traffic from reaching your Pods, as Google Cloud does not route PUIP addresses to the internet or to other VPC networks by default.

It enables you to use VPC Network Peering to connect your GKE cluster to other VPC networks that use different PUIP addresses, as long as you enable the export and import of custom routes for the peering connection.

It preserves the fully integrated network model of GKE, where Pods can communicate with nodes and other resources in the same VPC network without NAT.

The options that you need to select when creating a private GKE cluster with PUIP

addresses are:

--disable-default-snat: This option disables source NAT for outbound traffic from Pods to destinations outside the cluster's VPC network. This is necessary to prevent Pods from using RFC 1918 addresses as their source IP addresses, which

could cause conflicts with other networks that use the same address space².

--enable-ip-alias: This option enables alias IP ranges for Pods and Services, which allows you to use separate subnet ranges for them. This is required to use PUIP addresses for Pods¹.

--enable-private-nodes: This option creates a private cluster, where nodes do not have external IP addresses and can only communicate with the control plane through a private endpoint. This enhances the security and privacy of your cluster³.

Option A is incorrect because it does not use PUIP addresses for Pods, but rather RFC 1918 addresses. This does not solve the problem of address exhaustion in your enterprise.



Option B is incorrect because it reuses the secondary address range for Services across multiple private GKE clusters, which could cause IP conflicts and routing issues.

Option C is incorrect because it does not specify the options that are needed to create a private GKE cluster with PUIP addresses.

1: Configuring privately used public IPs for GKE | Kubernetes Engine | Google Cloud 2: Using Cloud NAT with GKE | Kubernetes Engine | Google Cloud 3: Private clusters | Kubernetes Engine | Google Cloud

QUESTION 4

Your organization uses a Shared VPC architecture with a host project and three service projects. You have Compute Engine instances that reside in the service projects. You have critical workloads in your on-premises data center. You need to ensure that the Google Cloud instances can resolve on-premises hostnames via the Dedicated Interconnect you deployed to establish hybrid connectivity. What should you do?

- A. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 35.199.192.0/19 to the on-premises environment.
- B. Create a Cloud DNS private forwarding zone in the host project of the Shared VPC that forwards the Private zone to the on-premises DNS servers. In your Cloud Router, add a custom route advertisement for the IP 169.254 169.254 to the on-premises environment.
- C. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. In your Cloud Router, add a custom route advertisement for the IP 169.254 169 254 to the on-premises environment.
- D. Configure a Cloud DNS private zone in the host project of the Shared VPC. Set up DNS forwarding to your Google Cloud private zone on your on-premises DNS servers to point to the inbound forwarder IP address in your host project. Configure a DNS policy in the Shared VPC to allow inbound query forwarding with your on-premises DNS server as the alternative DNS server.

Correct Answer: D

QUESTION 5

You 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.



Correct Answer: B

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

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

[PROFESSIONAL-CLOUD-NETWORK-ENGINEER Exam Questions](#)