[Q65-Q82 Accurate & Verified 2025 New Professional-Cloud-Network-Engineer Answers As Experienced in the Actual Test!

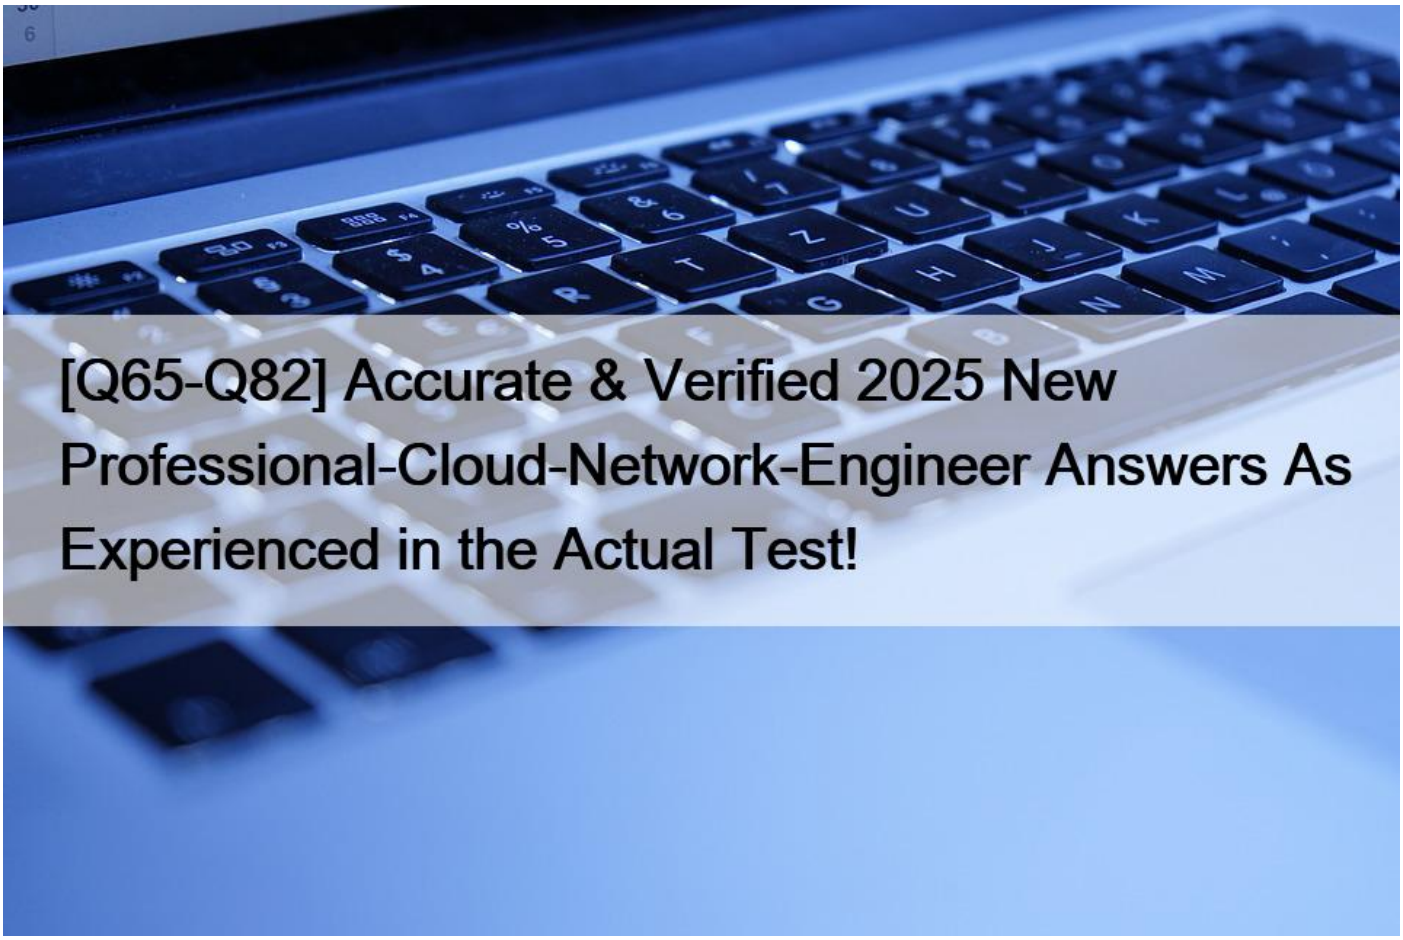

Accurate & Verified 2025 New Professional-Cloud-Network-Engineer Answers As Experienced in the Actual Test!
Professional-Cloud-Network-Engineer Certification Sample Questions certification Exam

Google Professional-Cloud-Network-Engineer certification exam is a challenging and comprehensive exam that requires significant preparation. It is a two-hour, multiple-choice exam that consists of 50 questions. Professional-Cloud-Network-Engineer exam is available in English and Japanese and can be taken at any of the authorized testing centers worldwide.

NO.65 Your 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?
* 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.
* 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.
* 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.
* 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.

NO.66 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?
* Create an allow on match ingress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.
* Create an allow on match egress firewall rule with the target tag “web-server” to allow all IP addresses for TCP port 80.
* 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.
* 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.

NO.67 You have setup a shared VPC and you have created three projects; Host Project, Service Project-1 and Service Project-2. You have created two subnets, subnet-1 in us-west1 and subnet-

2 in us-central1 in the Host Project. Only subnet-1 has been shared with Service Project -1 but when you go to VPC networks in

Service Project-1 you also see subnet-2 which hasn’t been shared with Service Project-1. Please select the correct option from below why is subnet-2 available to Service Project-1. Note Host Project is the Host Project in the shared VPC, Service Project-1 and Service project-2 are the Service Projects in the shared VPC.
* The current user has Shared VPC Admin role and with Shared VPC Admin role all the networks are available.
* It is a bug in Google Cloud, please report it.
* By default all subnets are available.
* Remove Shared Network admin role to the current user.

Option A is the Correct choice because , if the current user has Shared VPC Admin role then all the networks in the shared VPC is the available to the user irrespective of subnet level sharing permission with the Service Projects.

Option B is Incorrect because , it is not a bug .

Option C is Incorrect because ,all the subnets would be available if the current user has Shared Admin role.

Option D is Incorrect because ,Shared Network Admin role doesn’t exist.

NO.68 You create a Google Kubernetes Engine private cluster and want to use kubectl to get the status of the pods. In one of your instances you notice the master is not responding, even though the cluster is up and running.

What should you do to solve the problem?
* Assign a public IP address to the instance.
* Create a route to reach the Master, pointing to the default internet gateway.
* Create the appropriate firewall policy in the VPC to allow traffic from Master node IP address to the instance.
* Create the appropriate master authorized network entries to allow the instance to communicate to the master.

NO.69 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?
* Assign each user the editor role.
* Assign each user the compute.networkAdmin role.
* Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get.
* Give each user the following permissions only: compute.interconnectAttachments.create, compute.interconnectAttachments.get, compute.routers.create, compute.routers.get, compute.routers.update.

Explanation/Reference:

NO.70 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?
* Update the TTL for the zone.
* Set the zone to the TRANSFER state.
* Disable DNSSEC at your domain registar.
* Transfer ownership of the domain to a new registar.

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.

NO.71 You need to create a GKE cluster in an existing VPC that is accessible from on-premises. You must meet the following requirements:

IP ranges for pods and services must be as small as possible.

The nodes and the master must not be reachable from the internet.

You must be able to use kubectl commands from on-premises subnets to manage the cluster.

How should you create the GKE cluster?
* * Create a private cluster that uses VPC advanced routes.

* Set the pod and service ranges as /24.

* Set up a network proxy to access the master.
* * Create a VPC-native GKE cluster using GKE-managed IP ranges.

* Set the pod IP range as /21 and service IP range as /24.

* Set up a network proxy to access the master.
* * Create a VPC-native GKE cluster using user-managed IP ranges.

* Enable a GKE cluster network policy, set the pod and service ranges as /24.

* Set up a network proxy to access the master.

* Enable master authorized networks.
* * Create a VPC-native GKE cluster using user-managed IP ranges.

* Enable privateEndpoint on the cluster master.

* Set the pod and service ranges as /24.

* Set up a network proxy to access the master.
* Enable master authorized networks.

Explanation:

Creating GKE private clusters with network proxies for controller access When you create a GKE private cluster with a private cluster controller endpoint, the cluster’s controller node is inaccessible from the public internet, but it needs to be accessible for administration. By default, clusters can access the controller through its private endpoint, and authorized networks can be defined within the VPC network. To access the controller from on-premises or another VPC network, however, requires additional steps. This is because the VPC network that hosts the controller is owned by Google and cannot be accessed from resources connected through another VPC network peering connection, Cloud VPN or Cloud Interconnect.
https://cloud.google.com/solutions/creating-kubernetes-engine-private-clusters-with-net-proxies

NO.72 Your company has a Virtual Private Cloud (VPC) with two Dedicated Interconnect connections in two different regions: us-west1 and us-east1. Each Dedicated Interconnect connection is attached to a Cloud Router in its respective region by a VLAN attachment. You need to configure a high availability failover path. By default, all ingress traffic from the on-premises environment should flow to the VPC using the us-west1 connection. If us-west1 is unavailable, you want traffic to be rerouted to us-east1. How should you configure the multi-exit discriminator (MED) values to enable this failover path?
* Use regional routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority

of 1
* Use global routing. Set the us-east1 Cloud Router to a base priority of 100, and set the us-west1 Cloud Router to a base priority of 1
* Use regional routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1
* Use global routing. Set the us-east1 Cloud Router to a base priority of 1000, and set the us-west1 Cloud Router to a base priority of 1

NO.73 You are using a third-party next-generation firewall to inspect traffic. You created a custom route of 0.0.0.0/0 to route egress traffic to the firewall. You want to allow your VPC instances without public IP addresses to access the BigQuery and Cloud Pub/Sub APIs, without sending the traffic through the firewall.

Which two actions should you take? (Choose two.)
* Turn on Private Google Access at the subnet level.
* Turn on Private Google Access at the VPC level.
* Turn on Private Services Access at the VPC level.
* Create a set of custom static routes to send traffic to the external IP addresses of Google APIs and services via the default internet gateway.
* Create a set of custom static routes to send traffic to the internal IP addresses of Google APIs and services via the default internet gateway.
Explanation/Reference: https://cloud.google.com/vpc/docs/private-access-options

NO.74 You are configuring the final elements of a migration effort where resources have been moved from on-premises to Google Cloud. While reviewing the deployed architecture, you noticed that DNS resolution is failing when queries are being sent to the on-premises environment. You log in to a Compute Engine instance, try to resolve an on-premises hostname, and the query fails. DNS queries are not arriving at the on-premises DNS server. You need to use managed services to reconfigure Cloud DNS to resolve the DNS error. What should you do?
* Validate that the Compute Engine instances are using the Metadata Service IP address as their resolver. Configure an outbound forwarding zone for the on-premises domain pointing to the on-premises DNS server. Configure Cloud Router to advertise the Cloud DNS proxy range to the on-premises network.
* Validate that there is network connectivity to the on-premises environment and that the Compute Engine instances can reach other on-premises resources. If errors persist, remove the VPC Network Peerings and recreate the peerings after validating the routes.
* Review the existing Cloud DNS zones, and validate that there is a route in the VPC directing traffic destined to the IP address of the DNS servers. Recreate the existing DNS forwarding zones to forward all queries to the on-premises DNS servers.
* Ensure that the operating systems of the Compute Engine instances are configured to send DNS queries to the on-premises DNS servers directly.
To resolve DNS resolution issues for on-premises domains from Google Cloud, you should use Cloud DNS outbound forwarding zones. This setup forwards DNS requests for specific domains to on-premises DNS servers. Cloud Router is needed to advertise the range for the DNS proxy service back to the on-premises environment, ensuring that DNS queries from Compute Engine instances reach the on-premises DNS servers.

NO.75 You have created a firewall with rules that only allow traffic over HTTP, HTTPS, and SSH ports. While testing, you specifically try to reach the server over multiple ports and protocols; however, you do not see any denied connections in the firewall logs. You want to resolve the issue.

What should you do?
* Enable logging on the default Deny Any Firewall Rule.
* Enable logging on the VM Instances that receive traffic.
* Create a logging sink forwarding all firewall logs with no filters.
* Create an explicit Deny Any rule and enable logging on the new rule.

NO.76 Your team deployed two applications in GKE that are exposed through an external Application Load Balancer. When queries are sent to www.mountkirkgames.com/sales and www.mountkirkgames.com/get-an-analysis, the correct pages are displayed. However, you have received complaints that www.mountkirkgames.com yields a 404 error. You need to resolve this error. What should you do?
* Review the Ingress YAML file. Define the default backend. Reapply the YAML.
* Review the Ingress YAML file. Add a new path rule for the * character that directs to the base service. Reapply the YAML.
* Review the Service YAML file. Define a default backend. Reapply the YAML.
* Review the Service YAML file. Add a new path rule for the * character that directs to the base service. Reapply the YAML.
The 404 error is occurring because there is no default backend defined for requests to the root URL. Defining the default backend in the Ingress YAML file ensures that requests to www.mountkirkgames.com are routed to the correct service.

NO.77 You need to restrict access to your Google Cloud load-balanced application so that only specific IP addresses can connect.

What should you do?
* Create a secure perimeter using the Access Context Manager feature of VPC Service Controls and restrict access to the source IP range of the allowed clients and Google health check IP ranges.
* Create a secure perimeter using VPC Service Controls, and mark the load balancer as a service restricted to the source IP range of the allowed clients and Google health check IP ranges.
* Tag the backend instances “application,” and create a firewall rule with target tag “application” and the source IP range of the allowed clients and Google health check IP ranges.
* Label the backend instances “application,” and create a firewall rule with the target label “application” and the source IP range of the allowed clients and Google health check IP ranges.
https://link.springer.com/chapter/10.1007/978-1-4842-1004-8_4

NO.78 You work for one of the biggest digital media company in USA .The company management has decided to move 90 TB of backups and archival data to Google Cloud. They are looking for long term cost effective archival storage for disaster recovery in Google Cloud . Please select the right solution.
* Storage Transfer and Nearline storage
* Transfer Appliance and Coldline storage
* gsutil and Cloud storage
* Transfer Appliance and Nearline storage
Option B is the correct choice because ,Transfer Appliance is the best choice moving large volume of data and since they are looking for long term cost effective disaster recovery solution , coldline is the best option.

Option A is Incorrect because Storage Transfer is used to import online data into Cloud Storage .

Your online data source can be an Amazon Simple Storage Service (Amazon S3) bucket, an HTTP/HTTPS location, or a Cloud Storage bucket Option C is Incorrect because , gsutil isn’t recommended for large volume of data transfer ,It will take a very long time for data transfer depending on the bandwidth.

Option D is Incorrect because , Coldline is a more cost effective archival storage for disaster recovery.

NO.79 The security team has disabled external SSH access into production virtual machines in GCP.

The operations team needs to remotely manage the VMs and other resources. What can they do?
* Develop a new access request process that grants temporary SSH access to cloud VMs when an operations engineer needs to perform a task.
* Grant the operations team access to use Google Cloud Shell.
* Have the development team build an API service that allows the operations team to execute specific remote procedure calls to

accomplish their tasks.
* Configure a VPN connection to GCP to allow SSH access to the cloud VMs.
Grant the operations team access to use Google Cloud Shell.

B (Correct Answer) – Grant the operations engineers access to use Google Cloud Shell.

All the engineer asked is remote access the VMs just like using SSH, so if the machines still have an external IP address, the engineers can access them via SSH using Google Cloud Shell.

This is easies effective way to meet the requirements. All other answers are possible options that might require more setup than worthwhile for your needs.

NO.80 You work for a university that is migrating to GCP.

These are the cloud requirements:

* On-premises connectivity with 10 Gbps

* Lowest latency access to the cloud

* Centralized Networking Administration Team

New departments are asking for on-premises connectivity to their projects. You want to deploy the most cost-efficient interconnect solution for connecting the campus to Google Cloud.

What should you do?
* Use Shared VPC, and deploy the VLAN attachments and Interconnect in the host project.
* Use Shared VPC, and deploy the VLAN attachments in the service projects. Connect the VLAN attachment to the Shared VPC’s host project.
* Use standalone projects, and deploy the VLAN attachments in the individual projects. Connect the VLAN attachment to the standalone projects’ Interconnects.
* Use standalone projects and deploy the VLAN attachments and Interconnects in each of the individual projects.
https://cloud.google.com/interconnect/docs/how-to/dedicated/using-interconnects-other-projects Using Cloud Interconnect with Shared VPC You can use Shared VPC to share your VLAN attachment in a project with other VPC networks. Choosing Shared VPC is preferable if you need to create many projects and would like to prevent individual project owners from managing their connectivity back to your on-premises network. In this scenario, the host project contains a common Shared VPC network usable by VMs in service projects. Because VMs in the service projects use this network, Service Project Admins don’t need to create other VLAN attachments or Cloud Routers in the service projects. In this scenario, you must create VLAN attachments and Cloud Routers for a Cloud Interconnect connection only in the Shared VPC host project. The combination of a VLAN attachment and its associated Cloud Router are unique to a given Shared VPC network.
https://cloud.google.com/network-connectivity/docs/interconnect/how-to/enabling-multiple-networks-access-same-attachment#using_with

https://cloud.google.com/vpc/docs/shared-vpc

NO.81 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.)
* Activate the Service Networking API in your project.
* Activate the Cloud Datastore API in your project.

* Create a private connection to a service producer.
* Create a custom static route to allow the traffic to reach the Cloud SQL API.
* Enable Private Google Access.

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

NO.82 You 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?

* 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.
* Configure a Google Cloud Armor security policy in your project, and attach it to the backend service to secure the application.
* Configure VPC firewall rules to protect the Compute Engine instances against distributed denial-of-service attacks.
* Configure hierarchical firewall rules for the global HTTP(S) load balancer public IP address at the organization level.

Certification Topics of Professional-Cloud-Network-Engineer Exam PDF Recently Updated Questions:
https://www.exams4sures.com/Google/Professional-Cloud-Network-Engineer-practice-exam-dumps.html]