Google Professional-Cloud-Network-Engineer Questions and Answers Guarantee you Oass the Test Easily [Q58-Q82



Google Professional-Cloud-Network-Engineer Questions and Answers Guarantee you Oass the Test Easily Share Latest Professional-Cloud-Network-Engineer DUMP with 213 Questions and Answers

Google Cloud Certified - Professional Cloud Network Engineer certification is a coveted credential for networking professionals who want to master the Google Cloud Platform. It is designed for individuals who have experience in implementing and managing network architectures in the cloud. Google Cloud Certified - Professional Cloud Network Engineer certification validates the expertise of network engineers in designing, implementing, and managing secure and scalable cloud network solutions.

Google Professional-Cloud-Network-Engineer certification is an excellent way for IT professionals to demonstrate their expertise in networking technologies and solutions on the Google Cloud Platform. By passing this certification exam, candidates can validate their skills and knowledge in this area, which can help them to advance their careers and open up new opportunities in the rapidly growing cloud computing industry.

Q58. 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.)

- * Connect both projects using Cloud VPN.
- * Connect the VPCs in project code-dev and data-dev using VPC Network Peering.
- * Enable Shared VPC in one project (e. g., code-dev), and make the second project (e. g., data-dev) a service project.

* 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.

* 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.

Q59. You need to configure a static route to an on-premises resource behind a Cloud VPN gateway that is configured for policy-based routing using the gcloud command.

Which next hop should you choose?

- * The default internet gateway
- * The IP address of the Cloud VPN gateway
- * The name and region of the Cloud VPN tunnel
- * The IP address of the instance on the remote side of the VPN tunnel

Reference:

https://cloud.google.com/vpn/docs/how-to/creating-static-vpns

Q60. 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?

- * Tune TCP parameters on the on-premises servers.
- * Compress files using utilities like tar to reduce the size of data being sent.
- * Remove the -m flag from the gsutil command to enable single-threaded transfers.
- * Use the perfdiag parameter in your gsutil command to enable faster performance: gsutil perfdiag gs://[BUCKET_NAME].

Q61. 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?

- * Double the MTU on your on-premises VPN gateway from 1460 bytes to 2920 bytes.
- * Create two VPN tunnels on the same Cloud VPN gateway that point to the same destination VPN gateway IP address.

* 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.

* 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. https://cloud.google.com/network-connectivity/docs/vpn/concepts/classic-topologies#redundancy-options

Q62. Recently, your networking team enabled Cloud CDN for one of the external-facing services that is exposed through an external Application Load Balancer. The application team has already defined which content should be cached within the responses. Upon

testing the load balancer, you did not observe any change in performance after the Cloud CDN enablement. You need to resolve the issue. What should you do?

* Configure the CACHE_MAX_STATIC caching mode on Cloud CDN to ensure Cloud CDN caches content depending on responses from the backends.

* Configure the USE_ORIGIN_HEADERS caching mode on Cloud CDN to ensure Cloud CDN caches content based on response headers from the backends.

* Configure the CACHE_ALL_STATIC caching mode on Cloud CDN to ensure Cloud CDN caches all static content as well as content defined by the backends.

* Configure the FORCE_CACHE_ALL caching mode on Cloud CDN to ensure all appropriate content is cached.

When enabling Cloud CDN, for caching behavior to follow the application-defined caching headers, you need to configure the USE_ORIGIN_HEADERS caching mode. This setting ensures that the Cloud CDN respects the cache control headers specified by the backend, allowing the application-defined caching rules to dictate what content gets cached. This is often required when specific caching directives are already set by the application.

Q63. You have a storage bucket that contains the following objects:

– folder-a/image-a-1.jpg

– folder-a/image-a-2.jpg

– folder-b/image-b-1.jpg

– folder-b/image-b-2.jpg

Cloud CDN is enabled on the storage bucket, and all four objects have been successfully cached.

You want to remove the cached copies of all the objects with the prefix folder-a, using the minimum number of commands.

What should you do?

- * Add an appropriate lifecycle rule on the storage bucket.
- * Issue a cache invalidation command with pattern /folder-a/*.
- * Make sure that all the objects with prefix folder-a are not shared publicly.
- * Disable Cloud CDN on the storage bucket. Wait 90 seconds. Re-enable Cloud CDN on the storage bucket.

Q64. You need to create a new VPC network that allows instances to have IP addresses in both the 10.1.1.0/24 network and the 172.16.45.0/24 network.

What should you do?

- * Configure global load balancing to point 172.16.45.0/24 to the correct instance.
- * Create unique DNS records for each service that sends traffic to the desired IP address.
- * Configure an alias-IP range of 172.16.45.0/24 on the virtual instances within the VPC subnet of 10.1.1.0/24.
- * Use VPC peering to allow traffic to route between the 10.1.0.0/24 network and the 172.16.45.0/24 network.

Q65. Your organization recently re-architected your cloud environment to use Network Connectivity Center. However, an error occurred when you tried to add a new VPC named vpc-dev as a spoke. The error indicated that there was an issue with an existing spoke and the IP space of a VPC named vpc-pre-prod. You must complete the migration quickly and efficiently. What should you do?

* Remove the conflicting VPC spoke for vpc-pre-prod from the set of VPC spokes in Network Connectivity Center. Add the VPC spoke for vpc-dev. Add the previously removed vpc-pre-prod as a VPC spoke.

* Delete the VMs associated with the conflicting subnets, then delete the conflicting subnets in vpc-dev. Recreate the subnets with a

new IP range and redeploy the previously deleted VMs in the new subnets. Add the VPC spoke for vpc-dev.

* Exclude the conflicting IP range by using the –exclude-export-ranges flag when creating the VPC spoke for vpc-dev.

* Exclude the conflicting IP range by using the –exclude-export-ranges flag in the hub when attaching the VPC spoke for vpc-dev.

The most efficient way to resolve the conflict is to temporarily remove the conflicting vpc-pre-prod spoke, add the vpc-dev spoke, and then re-add vpc-pre-prod. This ensures that the migration happens quickly without the need to change IP ranges or delete resources.

Q66. You have deployed a proof-of-concept application by manually placing instances in a single Compute Engine zone. You are now moving the application to production, so you need to increase your application availability and ensure it can autoscale.

How should you provision your instances?

- * Create a single managed instance group, specify the desired region, and select Multiple zones for the location.
- * Create a managed instance group for each region, select Single zone for the location, and manually distribute instances across the zones in that region.
- * Create an unmanaged instance group in a single zone, and then create an HTTP load balancer for the instance group.

* Create an unmanaged instance group for each zone, and manually distribute the instances across the desired zones.

https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-managed-instances

Q67. You recently configured Google Cloud Armor security policies to manage traffic to your application. You discover that Google Cloud Armor is incorrectly blocking some traffic to your application. You need to identity the web application firewall (WAF) rule that is incorrectly blocking traffic. What should you do?

- * Enable firewall logs, and view the logs in Firewall Insights.
- * Enable HTTP(S) Load Balancing logging with sampling rate equal to 1, and view the logs in Cloud Logging.
- * Enable VPC Flow Logs, and view the logs in Cloud Logging.
- * Enable Google Cloud Armor audit logs, and view the logs on the Activity page in the Google Cloud Console.

Q68. In your Google Cloud organization, you have two folders: Dev and Prod. You want a scalable and consistent way to enforce the following firewall rules for all virtual machines (VMs) with minimal cost:

Port 8080 should always be open for VMs in the projects in the Dev folder.

Any traffic to port 8080 should be denied for all VMs in your projects in the Prod folder.

What should you do?

* Create and associate a firewall policy with the Dev folder with a rule to open port 8080. Create and associate a firewall policy with the Prod folder with a rule to deny traffic to port 8080.

* Create a Shared VPC for the Dev projects and a Shared VPC for the Prod projects. Create a VPC firewall rule to open port 8080 in the Shared VPC for Dev. Create a firewall rule to deny traffic to port 8080 in the Shared VPC for Prod. Deploy VMs to those Shared VPCs.

* In all VPCs for the Dev projects, create a VPC firewall rule to open port 8080. In all VPCs for the Prod projects, create a VPC firewall rule to deny traffic to port 8080.

* Use Anthos Config Connector to enforce a security policy to open port 8080 on the Dev VMs and deny traffic to port 8080 on the Prod VMs.

Q69. Your company has a single Virtual Private Cloud (VPC) network deployed in Google Cloud with access from on-premises locations using Cloud Interconnect connections. Your company must be able to send traffic to Cloud Storage only through the Interconnect links while accessing other Google APIs and services over the public internet. What should you do?

* Use the default public domains for all Google APIs and services.

* Use Private Service Connect to access Cloud Storage, and use the default public domains for all other Google APIs and services.

* Use Private Google Access, with restricted.googleapis.com virtual IP addresses for Cloud Storage and private.googleapis.com for all other Google APIs and services.

* Use Private Google Access, with private.googleapis.com virtual IP addresses for Cloud Storage and restricted.googleapis.com virtual IP addresses for all other Google APIs and services.

Q70. 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?

* Create 2 VPCs, each with their own regions and individual subnets. Create 2 VPN gateways to establish connectivity between these regions.

* Create 2 VPCs, each with their own region and individual subnets. Use external IP addresses on the instances to establish connectivity between these regions.

* Create 1 VPC with 2 regional subnets. Create a global load balancer to establish connectivity between the regions.

* Create 1 VPC with 2 regional subnets. Deploy workloads in these subnets and have them communicate using private RFC1918 IP addresses.

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.

Q71. You are troubleshooting connectivity issues between Google Cloud and a public SaaS provider. Connectivity between the two environments is through the public internet. Your users are reporting intermittent connection errors when using TCP to connect; however, ICMP tests show no failures. According to users, errors occur around the same time every day. You want to troubleshoot and gather information by using Google Cloud tools that are most likely to provide insights into what is occurring within Google Cloud. What should you do?

* Create a Connectivity Test by using TCP, the source IP address of your test VM, and the destination IP address of the public SaaS provider. Review the live data plane analysis and take the next steps based on the test results.

* Enable and review Cloud Logging on your Cloud NAT gateway. Look for logs with errors matching the destination IP address of the public SaaS provider.

* Enable the Firewall insights API. Set the deny rule insights observation period to one day. Review the insights to assure there are no firewall rules denying traffic.

* Enable and review Cloud Logging for Cloud Armor. Look for logs with errors matching the destination IP address of the public SaaS provider.

When troubleshooting connectivity issues, especially over public internet connections with intermittent errors, Connectivity Tests in Network Intelligence Center are crucial. This tool allows you to simulate the connectivity and understand the data plane status of Google Cloud resources. Since ICMP tests pass but TCP tests fail intermittently, using Connectivity Tests with TCP parameters will provide detailed insight into possible network issues like route misconfigurations, peering issues, or other transient problems affecting only specific protocols.

Q72. You need to ensure your personal SSH key works on every instance in your project. You want to accomplish this as efficiently as possible.

What should you do?

- * Upload your public ssh key to the project Metadata.
- * Upload your public ssh key to each instance Metadata.
- * Create a custom Google Compute Engine image with your public ssh key embedded.

* Use gcloud compute ssh to automatically copy your public ssh key to the instance.

https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys

Q73. You are configuring load balancing for a standard three-tier (web, application, and database) application. You have configured an external HTTP(S) load balancer for the web servers. You need to configure load balancing for the application tier of servers.

What should you do?

- * Configure a forwarding rule on the existing load balancer for the application tier.
- * Configure equal cost multi-path routing on the application servers.
- * Configure a new internal HTTP(S) load balancer for the application tier.
- * Configure a URL map on the existing load balancer to route traffic to the application tier.

Q74. You have several microservices running in a private subnet in an existing Virtual Private Cloud (VPC). You need to create additional serverless services that use Cloud Run and Cloud Functions to access the microservices. The network traffic volume between your serverless services and private microservices is low. However, each serverless service must be able to communicate with any of your microservices. You want to implement a solution that minimizes cost. What should you do?

* Deploy your serverless services to the serverless VPC. Peer the serverless service VPC to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.

* Create a serverless VPC access connector for each serverless service. Configure the connectors to allow traffic between the serverless services and your existing microservices.

* Deploy your serverless services to the existing VPC. Configure firewall rules to allow traffic between the serverless services and your existing microservices.

* Create a serverless VPC access connector. Configure the serverless service to use the connector for communication to the microservices.

Q75. 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?

- * Create two firewall rules: one to block all traffic with priority 0, and another to allow port 22 with priority 1000.
- * Create two firewall rules: one to block all traffic with priority 65536, and another to allow port 3389 with priority 1000.

* Create a single firewall rule to allow port 22 with priority 1000.

* Create a single firewall rule to allow port 3389 with priority 1000.

Reference:

https://geekflare.com/gcp-firewall-configuration/

Q76. 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?

* Double the MTU on your on-premises VPN gateway from 1460 bytes to 2920 bytes.

* Create two VPN tunnels on the same Cloud VPN gateway that point to the same destination VPN gateway IP address.

* 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.

* 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. Explanation/Reference:

Q77. You are using a third-party next-generation firewall to inspect traffic. You created a custom route of 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.

https://cloud.google.com/vpc/docs/private-access-options#pga Private Google Access VM instances that only have internal IP addresses (no external IP addresses) can use Private Google Access. They can reach the _external IP addresses_ of Google APIs and services.

Q78. 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:

- * Each on-premises router is configured with a unique ASN.
- * Each on-premises router is configured with the same routes and priorities.
- * Both on-premises routers are configured with a VPN connected to a single Cloud Router.
- * BGP sessions are established between both on-premises routers and the Cloud Router.
- * Only 1 of the on-premises router & #8217;s routes are being added to the routing table.

What is the most likely cause of this problem?

- * The on-premises routers are configured with the same routes.
- * A firewall is blocking the traffic across the second VPN connection.
- * You do not have a load balancer to load-balance the network traffic.
- * The ASNs being used on the on-premises routers are different.

Q79. You have a storage bucket that contains two objects. Cloud CDN is enabled on the bucket, and both objects have been successfully cached. Now you want to make sure that one of the two objects will not be cached anymore, and will always be served to the internet directly from the origin.

What should you do?

* Ensure that the object you don't want to be cached anymore is not shared publicly.

* Create a new storage bucket, and move the object you don't want to be checked anymore inside it. Then edit the bucket setting and enable the privateattribute.

* Add an appropriate lifecycle rule on the storage bucket containing the two objects.

* Add a Cache-Controlentry with value private to the metadata of the object you don't want to be cached anymore. Invalidate all the previously cached copies.

 $Explanation/Reference:\ https://developers.google.com/web/ilt/pwa/caching-files-with-service-worker$

Q80. You need to enable Private Google Access for use by some subnets within your Virtual Private Cloud (VPC). Your security team set up the VPC to send all internet-bound traffic back to the on- premises data center for inspection before egressing to the internet, and is also implementing VPC Service Controls in the environment for API-level security control. You have already enabled the subnets for Private Google Access. What configuration changes should you make to enable Private Google Access while adhering to your security team's requirements?

* Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range.

Create a custom route that points Google's restricted API address range to the default internet gateway as the next hop. * Create a private DNS zone with a CNAME record for *.googleapis.com to restricted.googleapis.com, with an A record pointing to Google's restricted API address range.

Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.

* Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record painting to Google's private AP address range.

Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.

* Create a private DNS zone with a CNAME record for *.googleapis.com to private.googleapis.com, with an A record pointing to Google's private API address range.

Create a custom route that points Google's private API address range to the default internet gateway as the next hop.

Q81. You just finished your company's migration to Google Cloud and configured an architecture with 3 Virtual Private Cloud (VPC) networks: one for Sales, one for Finance, and one for Engineering. Every VPC contains over 100 Compute Engine instances, and now developers using instances in the Sales VPC and the Finance VPC require private connectivity between each other. You need to allow communication between Sales and Finance without compromising performance or security. What should you do?

- * Configure an HA VPN gateway between the Finance VPC and the Sales VPC.
- * Configure the instances that require communication between each other with an external IP address.
- * Create a VPC Network Peering connection between the Finance VPC and the Sales VPC.
- * Configure Cloud NAT and a Cloud Router in the Sales and Finance VPCs.

Q82. 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?

- * AS-Path
- * Community
- * Local Preference
- * Multi-exit Discriminator

https://cloud.google.com/router/docs/concepts/overview

Dumps for Free Professional-Cloud-Network-Engineer Practice Exam Questions: <u>https://www.vceprep.com/Professional-Cloud-Network-Engineer-latest-vce-prep.html</u>]