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

QUESTION 1

You have a web application that is currently hosted in the us-central1 region. Users experience high latency when traveling in Asia. You\\'ve configured a network load balancer, but users have not experienced a performance improvement. You want to decrease the latency.

What should you do?

A. Configure a policy-based route rule to prioritize the traffic.

B. Configure an HTTP load balancer, and direct the traffic to it.

C. Configure Dynamic Routing for the subnet hosting the application.

D. Configure the TTL for the DNS zone to decrease the time between updates.

Correct Answer: B

QUESTION 2

You created a new VPC network named Dev with a single subnet. You added a firewall rule for the network Dev to allow HTTP traffic only and enabled logging. When you try to log in to an instance in the subnet via Remote Desktop Protocol, the login fails. You look for the Firewall rules logs in Stackdriver Logging, but you do not see any entries for blocked traffic. You want to see the logs for blocked traffic.

What should you do?

A. Check the VPC flow logs for the instance.

B. Try connecting to the instance via SSH, and check the logs.

C. Create a new firewall rule to allow traffic from port 22, and enable logs.

D. Create a new firewall rule with priority 65500 to deny all traffic, and enable logs.

Correct Answer: D

Ingress packets in VPC Flow Logs are sampled after ingress firewall rules. If an ingress firewall rule denies inbound packets, those packets are not sampled by VPC Flow Logs. We want to see the logs for blocked traffic so we have to look for them in firewall logs. https://cloud.google.com/vpc/docs/flow-logs#key_properties

QUESTION 3

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.

Correct Answer: AE

QUESTION 4

Your organization has a single project that contains multiple Virtual Private Clouds (VPCs). You need to secure API access to your Cloud Storage buckets and BigQuery datasets by allowing API access only from resources in your corporate public networks. What should you do?

A. Create an access context policy that allows your VPC and corporate public network IP ranges, and then attach the policy to Cloud Storage and BigQuery.

B. Create a VPC Service Controls perimeter for your project with an access context policy that allows your corporate public network IP ranges.

C. Create a firewall rule to block API access to Cloud Storage and BigQuery from unauthorized networks.

D. Create a VPC Service Controls perimeter for each VPC with an access context policy that allows your corporate public network IP ranges.

Correct Answer: B

QUESTION 5

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?

A. 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

B. 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

C. 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

D. 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

Correct Answer: A

QUESTION 6

Your company has defined a resource hierarchy that includes a parent folder with subfolders for each department. Each department defines their respective project and VPC in the assigned folder and has the appropriate permissions to create Google Cloud firewall rules. The VPCs should not allow traffic to flow between them. You need to block all traffic from any source, including other VPCs, and delegate only the intra-VPC firewall rules to the respective departments. What should you do?

A. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 0.

B. Create a VPC firewall rule in each VPC to block traffic from any source, with priority 1000.

C. Create two hierarchical firewall policies per department\\'s folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to allow, and another lower-priority rule that blocks traffic from any other source.

D. Create two hierarchical firewall policies per department\\'s folder with two rules in each: a high-priority rule that matches traffic from the private CIDRs assigned to the respective VPC and sets the action to goto_next, and another lower-priority rule that blocks traffic from any other source.

Correct Answer: B

QUESTION 7

Your company is planning a migration to Google Kubernetes Engine. Your application team informed you that they require a minimum of 60 Pods per node and a maximum of 100 Pods per node Which Pod per node CIDR range should you use?

A. /24

B. /25

C. /26

D. /28

Correct Answer: B

To determine the Pod per node CIDR range, you need to calculate how many IP addresses are required for each node, and then choose the smallest CIDR range that can accommodate that number. A CIDR range of /n means that there are 2^(32-n) IP addresses available in that range. For example, a /24 range has 2^(32-24) = 256 IP addresses. According to the question, the application team requires a minimum of 60 Pods per node and a maximum of 100 Pods per node. Therefore, you need to choose a CIDR range that can provide at least 100 IP addresses per node, but not more than necessary. A /25 range has 2^(32-25) = 128 IP addresses, which is enough for 100 Pods per node. A /26 range has 2^(32-26) = 64 IP addresses, which is not enough for 60 Pods per node. A /24 range has 256 IP addresses, which is more than needed and wastes IP address space. A /28 range has 2^(32-28) = 16 IP addresses, which is far too small for any node. Therefore, the best option is B. /25. This is also consistent with the Google Kubernetes Engine documentation, which states that each node is allocated a /24 range of IP addresses for Pods by default, but the maximum number of Pods per node is 1101. This means that there are approximately twice as many available IP addresses as possible Pods, which is similar to the ratio of 128 to 100 in the /25 range.

1: Configure maximum Pods per node | Google Kubernetes Engine (GKE) | Google Cloud

QUESTION 8

You are deploying an application that runs on Compute Engine instances. You need to determine how to expose your application to a new customer You must ensure that your application meets the following requirements

1.

Maps multiple existing reserved external IP addresses to the Instance

2.

Processes IP Encapsulating Security Payload (ESP) traffic

What should you do?

A. Configure a target pool, and create protocol forwarding rules for each external IP address.

B. Configure a backend service, and create an external network load balancer for each external IP address

C. Configure a target instance, and create a protocol forwarding rule for each external IP address to be mapped to the instance.

D. Configure the Compute Engine Instances\\' network Interface external IP address from None to Ephemeral Add as many external IP addresses as required

Correct Answer: C

The correct answer is C. Configure a target instance, and create a protocol forwarding rule for each external IP address to be mapped to the instance.

This answer is based on the following facts:

A target instance is a Compute Engine instance that handles traffic from one or more forwarding rules1. You can use target instances to forward traffic to a single VM instance from one or more external IP addresses2. A protocol forwarding

rule specifies the IP protocol and port range for the traffic that you want to forward3. You can use protocol forwarding rules to forward traffic of any IP protocol, including ESP4.

The other options are not correct because:

QUESTION 9

You are migrating to Cloud DNS and want to import your BIND zone file. Which command should you use?

A. gcloud dns record-sets import ZONE_FILE --zone MANAGED_ZONE

B. gcloud dns record-sets import ZONE_FILE --replace-origin-ns --zone MANAGED_ZONE

C. gcloud dns record-sets import ZONE_FILE --zone-file-format --zone MANAGED_ZONE

D. gcloud dns record-sets import ZONE_FILE --delete-all-existing --zone MANAGED ZONE

Correct Answer: C

https://cloud.google.com/sdk/gcloud/reference/dns/record-sets/import

QUESTION 10

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

https://cloud.google.com/vpc/docs/using-vpc#create-auto-network We create one VPC network in auto mode that creates one subnet in each Google Cloud region automatically. So, region us-east1 and europe-west1 are in the same network and they can communicate using their internal IP address even though they are in different Regions. They take advantage of Google\\'s global fiber network.

QUESTION 11

Your company offers a popular gaming service. Your instances are deployed with private IP addresses, and external access is granted through a global load balancer. You have recently engaged a traffic-scrubbing service and want to restrict

your origin to allow connections only from the traffic-scrubbing service.

What should you do?

A. Create a Cloud Armor Security Policy that blocks all traffic except for the traffic-scrubbing service.

B. Create a VPC Firewall rule that blocks all traffic except for the traffic-scrubbing service.

C. Create a VPC Service Control Perimeter that blocks all traffic except for the traffic-scrubbing service.

D. Create IPTables firewall rules that block all traffic except for the traffic-scrubbing service.

Correct Answer: A

Global load balancer will proxy the connection . thus no trace of session origin IP. you should use Cloud Armor to geofence your service. https://cloud.google.com/load-balancing/docs/https

QUESTION 12

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 perfdiag gs://[BUCKET NAME].

Correct Answer: A

https://cloud.google.com/solutions/tcp-optimization-for-network-performance-in-gcp-and-hybrid
https://cloud.google.com/blog/products/gcp/5-steps-to-better-gcp-network-performance?hl=ml

QUESTION 13

Your company runs an enterprise platform on-premises using virtual machines (VMS). Your internet customers have created tens of thousands of DNS domains panting to your public IP addresses allocated to the Vtvls Typically, your customers hard-code your IP addresses In their DNS records You are now planning to migrate the platform to Compute Engine and you want to use Bring your Own IP you want to minimize disruption to the Platform What Should you d0?

A. Create a VPC and request static external IP addresses from Google Cloud Assagn the IP addresses to the Compute Engine instances. Notify your customers of the new IP addresses so they can update their DNS

B. Verify ownership of your IP addresses. After the verification, Google Cloud advertises and provisions the IP prefix for you_ Assign the IP addresses to the Compute Engine Instances

C. Create a VPC With the same IP address range as your on-premises network Asson the IP addresses to the Compute Engine Instances.

D. Verify ownership of your IP addresses. Use live migration to import the prefix Assign the IP addresses to Compute Engine instances.

Correct Answer: D

The correct answer is D because it allows you to use your own public IP addresses in Google Cloud without disrupting the platform or requiring your customers to update their DNS records. Option A is incorrect because it involves changing the IP addresses and notifying the customers, which can cause disruption and errors. Option B is incorrect because it does not use live migration, which is a feature that lets you control when Google starts advertising routes for your prefix. Option C is incorrect because it does not involve bringing your own IP addresses, but rather using Google-provided IP addresses. References: Bring your own IP addresses Professional Cloud Network Engineer uide Bring your own IP addresses (BYOIP) to Azure with Custom IP Prefix

QUESTION 14

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?

A. 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.

B. 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.

C. 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. Eet up a network proxy to access the master. Enable master authorized networks.

D. 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.

Correct Answer: D

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

QUESTION 15

You have the networking configuration shown. In the diagram Two VLAN attachments associated With two Dedicated Interconnect connections terminate on the same Cloud Router (mycloudrouter). The Interconnect connections terminate on two separate on-premises routers. You advertise the same prefixes from the Border Gateway Protocol (BOP) sessions associated with each Of the VLAN attachments.

You notice an asymmetric traffic flow between the two Interconnect connections. Which of the following actions should you take to troubleshoot the asymmetric traffic flow?

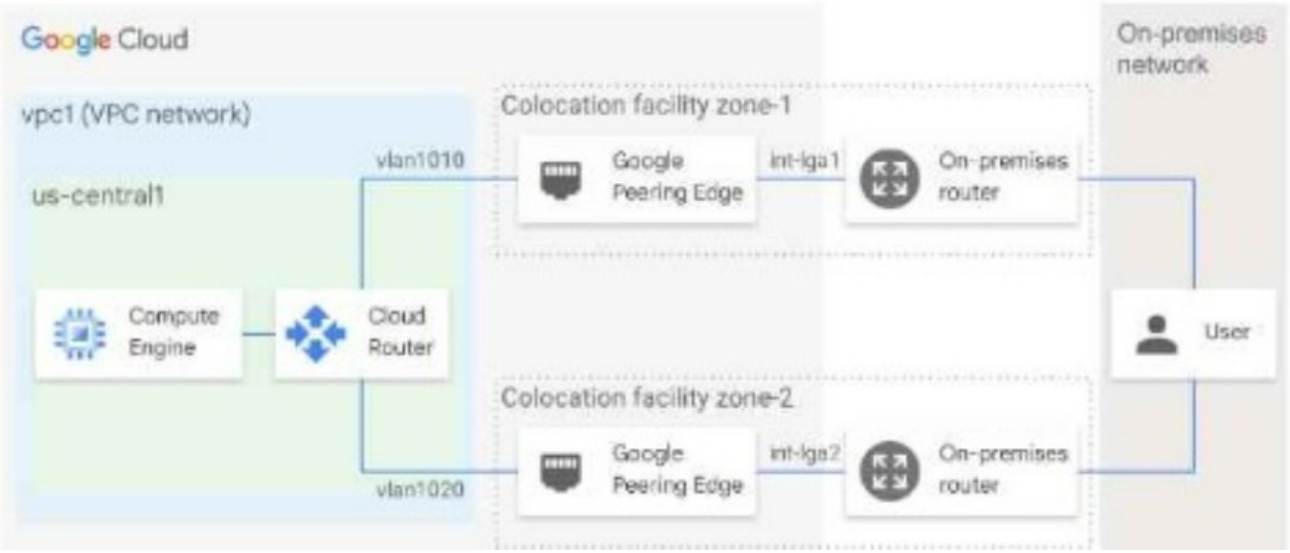

A. From the Google Cloud console, navigate to the Hybrid Connectivity select the Cloud Router, and view BGP sessions.

B. From the Cloud CLI, run gcloud compute -Protect_ID router get--status mycloudrouter ---region REGION and review the results.

C. From the Google Cloud console, navigate to Cloud Logging to view VPC Flow Logs and review the results

D. From the Cloud CLI. run gcloud compute routers describe mycloudrouter --region REGION and review the results

Correct Answer: A

The correct answer is B. From the Cloud CLI, run gcloud compute --project_ID router get-status mycloudrouter --region REGION and review the results. This command will show you the BGP session status, the advertised and learned routes,

and the last error for each VLAN attachment. You can use this information to troubleshoot the asymmetric traffic flow and identify any issues with the BGP configuration or the Interconnect connections.

The other options are not correct because:

Option A will only show you the BGP session status, but not the advertised and learned routes or the last error for each VLAN attachment. Option C will only show you the VPC Flow Logs, which are useful for monitoring and troubleshooting

network performance and security issues within your VPC network, but not for your Interconnect connections. Option D will only show you the basic information about the Cloud Router, such as its name, region, network, and BGP settings, but

not the detailed status of each VLAN attachment.