

Google

Exam Questions Professional-Cloud-Network-Engineer

Google Cloud Certified - Professional Cloud Network Engineer



NEW QUESTION 1

You have a Cloud Storage bucket in Google Cloud project XYZ. The bucket contains sensitive data. You need to design a solution to ensure that only instances belonging to VPCs under project XYZ can access the data stored in this Cloud Storage bucket. What should you do?

- A. Configure Private Google Access to privately access the Cloud Storage service using private IP addresses.
- B. Configure a VPC Service Controls perimeter around project XYZ, and include storage.googleapis.com as a restricted service in the service perimeter.
- C. Configure Cloud Storage with projectPrivate Access Control List (ACL) that gives permission to the project team based on their roles.
- D. Configure Private Service Connect to privately access Cloud Storage from all VPCs under project XYZ.

Answer: C

NEW QUESTION 2

You have an application that is running in a managed instance group. Your development team has released an updated instance template which contains a new feature which was not heavily tested. You want to minimize impact to users if there is a bug in the new template. How should you update your instances?

- A. Manually patch some of the instances, and then perform a rolling restart on the instance group.
- B. Using the new instance template, perform a rolling update across all instances in the instance group. Verify the new feature once the rollout completes.
- C. Deploy a new instance group and canary the updated template in that group.
- D. Verify the new feature in the new canary instance group, and then update the original instance group.
- E. Perform a canary update by starting a rolling update and specifying a target size for your instances to receive the new template.
- F. Verify the new feature on the canary instances, and then roll forward to the rest of the instances.

Answer: D

Explanation:

<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups#starting> <https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups>

NEW QUESTION 3

You have just deployed your infrastructure on Google Cloud. You now need to configure the DNS to meet the following requirements: Your on-premises resources should resolve your Google Cloud zones. Your Google Cloud resources should resolve your on-premises zones. You need the ability to resolve “.internal” zones provisioned by Google Cloud. What should you do?

- A. Configure an outbound server policy, and set your alternative name server to be your on-premises DNS resolver.
- B. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.
- C. Configure both an inbound server policy and outbound DNS forwarding zones with the target as the on-premises DNS resolver.
- D. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
- E. Configure an outbound DNS server policy, and set your alternative name server to be your on-premises DNS resolver.
- F. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google Cloud's DNS resolver.
- G. Configure Cloud DNS to DNS peer with your on-premises DNS resolver.
- H. Configure your on-premises DNS resolver to forward Google Cloud zone queries to Google's public DNS 8.8.8.8.

Answer: A

NEW QUESTION 4

You converted an auto mode VPC network to custom mode. Since the conversion, some of your Cloud Deployment Manager templates are no longer working. You want to resolve the problem. What should you do?

- A. Apply an additional IAM role to the Google API's service account to allow custom mode networks.
- B. Update the VPC firewall to allow the Cloud Deployment Manager to access the custom mode networks.
- C. Explicitly reference the custom mode networks in the Cloud Armor whitelist.
- D. Explicitly reference the custom mode networks in the Deployment Manager templates.

Answer: D

NEW QUESTION 5

You need to enable Cloud CDN for all the objects inside a storage bucket. You want to ensure that all the objects in the storage bucket can be served by the CDN. What should you do in the GCP Console?

- A. Create a new cloud storage bucket, and then enable Cloud CDN on it.
- B. Create a new TCP load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- C. Create a new SSL proxy load balancer, select the storage bucket as a backend, and then enable Cloud CDN on the backend.
- D. Create a new HTTP load balancer, select the storage bucket as a backend, enable Cloud CDN on the backend, and make sure each object inside the storage bucket is shared publicly.

Answer: D

Explanation:

https://cloud.google.com/load-balancing/docs/https/adding-backend-buckets-to-load-balancers#using_cloud_cdn Cloud CDN needs HTTP(S) Load Balancers and Cloud Storage bucket has to be shared publicly.
<https://cloud.google.com/cdn/docs/setting-up-cdn-with-bucket>

NEW QUESTION 6

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: AE

NEW QUESTION 7

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?

- A. 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.
- B. 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.
- C. 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. Change the custom route that points the default route (0/0) to the default internet gateway as the next hop.
- D. 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.

Answer: C

NEW QUESTION 8

You need to configure a Google Kubernetes Engine (GKE) cluster. The initial deployment should have 5 nodes with the potential to scale to 10 nodes. The maximum number of Pods per node is 8. The number of services could grow from 100 to up to 1024. How should you design the IP schema to optimally meet this requirement?

- A. Configure a /28 primary IP address range for the node IP addresses
- B. Configure a /25 secondary IP range for the Pod
- C. Configure a /22 secondary IP range for the Services.
- D. Configure a /28 primary IP address range for the node IP addresses
- E. Configure a /25 secondary IP range for the Pod
- F. Configure a /21 secondary IP range for the Services.
- G. Configure a /28 primary IP address range for the node IP addresses
- H. Configure a /28 secondary IP range for the Pod
- I. Configure a /21 secondary IP range for the Services.
- J. Configure a /28 primary IP address range for the node IP addresses
- K. Configure a /24 secondary IP range for the Pod
- L. Configure a /22 secondary IP range for the Services.

Answer: A

NEW QUESTION 9

You want to establish a dedicated connection to Google that can access Cloud SQL via a public IP address and that does not require a third-party service provider. Which connection type should you choose?

- A. Carrier Peering
- B. Direct Peering
- C. Dedicated Interconnect
- D. Partner Interconnect

Answer: B

Explanation:

When established, Direct Peering provides a direct path from your on-premises network to Google services, including Google Cloud products that can be exposed through one or more public IP addresses. Traffic from Google's network to your on-premises network also takes that direct path, including traffic from VPC networks in your projects. Google Cloud customers must request that direct egress pricing be enabled for each of their projects after they have established Direct Peering with Google. For more information, see Pricing.

NEW QUESTION 10

You have configured a service on Google Cloud that connects to an on-premises service via a Dedicated Interconnect. Users are reporting recent connectivity issues. You need to determine whether the traffic is being dropped because of firewall rules or a routing decision. What should you do?

- A. Use the Network Intelligence Center Connectivity Tests to test the connectivity between the VPC and the on-premises network.
- B. Use Network Intelligence Center Network Topology to check the traffic flow, and replay the traffic from the time period when the connectivity issue occurred.
- C. Configure VPC Flow Log
- D. Review the logs by filtering on the source and destination.
- E. Configure a Compute Engine instance on the same VPC as the service running on Google Cloud to run a traceroute targeted at the on-premises service.

Answer: B

NEW QUESTION 10

Your company has provisioned 2000 virtual machines (VMs) in the private subnet of your Virtual Private Cloud (VPC) in the us-east1 region. You need to configure each VM to have a minimum of 128 TCP connections to a public repository so that users can download software updates and packages over the internet. You need to implement a Cloud NAT gateway so that the VMs are able to perform outbound NAT to the internet. You must ensure that all VMs can simultaneously connect to the public repository and download software updates and packages. Which two methods can you use to accomplish this? (Choose two.)

- A. Configure the NAT gateway in manual allocation mode, allocate 2 NAT IP addresses, and update the minimum number of ports per VM to 256.
- B. Create a second Cloud NAT gateway with the default minimum number of ports configured per VM to 64.
- C. Use the default Cloud NAT gateway's NAT proxy to dynamically scale using a single NAT IP address.
- D. Use the default Cloud NAT gateway to automatically scale to the required number of NAT IP addresses, and update the minimum number of ports per VM to 128.
- E. Configure the NAT gateway in manual allocation mode, allocate 4 NAT IP addresses, and update the minimum number of ports per VM to 128.

Answer: AB

NEW QUESTION 11

You have created an HTTP(S) load balanced service. You need to verify that your backend instances are responding properly. How should you configure the health check?

- A. Set request-path to a specific URL used for health checking, and set proxy-header to PROXY_V1.
- B. Set request-path to a specific URL used for health checking, and set host to include a custom host header that identifies the health check.
- C. Set request-path to a specific URL used for health checking, and set response to a string that the backend service will always return in the response body.
- D. Set proxy-header to the default value, and set host to include a custom host header that identifies the health check.

Answer: C

Explanation:

https://cloud.google.com/load-balancing/docs/health-check-concepts#content-based_health_checks

NEW QUESTION 12

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?

- A. Enable logging on the default Deny Any Firewall Rule.
- B. Enable logging on the VM Instances that receive traffic.
- C. Create a logging sink forwarding all firewall logs with no filters.
- D. Create an explicit Deny Any rule and enable logging on the new rule.

Answer: D

Explanation:

https://cloud.google.com/vpc/docs/firewall-rules-logging#egress_deny_example

You can only enable Firewall Rules Logging for rules in a Virtual Private Cloud (VPC) network. Legacy networks are not supported. Firewall Rules Logging only records TCP and UDP connections. Although you can create a firewall rule applicable to other protocols, you cannot log their connections. You cannot enable Firewall Rules Logging for the implied deny ingress and implied allow egress rules. Log entries are written from the perspective of virtual machine (VM) instances. Log entries are only created if a firewall rule has logging enabled and if the rule applies to traffic sent to or from the VM. Entries are created according to the connection logging limits on a best effort basis. The number of connections that can be logged in a given interval is based on the machine type. Changes to firewall rules can be viewed in VPC audit logs. <https://cloud.google.com/vpc/docs/firewall-rules-logging#specifications>

NEW QUESTION 17

You have applications running in the us-west1 and us-east1 regions. You want to build a highly available VPN that provides 99.99% availability to connect your applications from your project to the cloud services provided by your partner's project while minimizing the amount of infrastructure required. Your partner's services are also in the us-west1 and us-east1 regions. You want to implement the simplest solution. What should you do?

- A. Create one Cloud Router and one HA VPN gateway in each region of your VPC and your partner's VP
- B. Connect your VPN gateways to the partner's gateway
- C. Enable global dynamic routing in each VPC.
- D. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VP
- E. Create one OpenVPN Access Server in each region of your partner's VP
- F. Connect your VPN gateway to your partner's servers.
- G. Create one OpenVPN Access Server in each region of your VPC and your partner's VP
- H. Connect your servers to the partner's servers.
- I. Create one Cloud Router and one HA VPN gateway in the us-west1 region of your VPC and your partner's VP
- J. Connect your VPN gateways to the partner's gateways with a pair of tunnel
- K. Enable global dynamic routing in each VPC.

Answer: A

NEW QUESTION 19

You are designing a hybrid cloud environment for your organization. Your Google Cloud environment is interconnected with your on-premises network using Cloud HA VPN and Cloud Router. The Cloud Router is configured with the default settings. Your on-premises DNS server is located at 192.168.20.88 and is protected by a firewall, and your Compute Engine resources are located at 10.204.0.0/24. Your Compute Engine resources need to resolve on-premises private hostnames using the domain corp.altostrat.com while still resolving Google Cloud hostnames. You want to follow Google-recommended practices. What should you do?

- A. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24. Set a custom route advertisement on the Cloud Router for 10.204.0.0/24
- B. Create a private forwarding zone in Cloud DNS for 'corp.altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 35.199.192.0/19 Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

C. Create a private forwarding zone in Cloud DNS for 'corp .altostrat.com' called corp-altostrat-com that points to 192.168.20.88. Configure your on-premises firewall to accept traffic from 10.204.0.0/24. Modify the /etc/resolv.conf file on your Compute Engine instances to point to 192.168.20.88

D. Create a private zone in Cloud DNS for 'corp altostrat.com' called corp-altostrat-com. Configure DNS Server Policies and create a policy with Alternate DNS servers to 192.168.20.88. Configure your on-premises firewall to accept traffic from 35.199.192.0/19. Set a custom route advertisement on the Cloud Router for 35.199.192.0/19.

Answer: D

NEW QUESTION 22

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

NEW QUESTION 24

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

NEW QUESTION 29

You have provisioned a Partner Interconnect connection to extend connectivity from your on-premises data center to Google Cloud. You need to configure a Cloud Router and create a VLAN attachment to connect to resources inside your VPC. You need to configure an Autonomous System number (ASN) to use with the associated Cloud Router and create the VLAN attachment. What should you do?

- A. Use a 4-byte private ASN 4200000000-4294967294.
- B. Use a 2-byte private ASN 64512-65535.
- C. Use a public Google ASN 15169.
- D. Use a public Google ASN 16550.

Answer: B

NEW QUESTION 34

You need to configure the Border Gateway Protocol (BGP) session for a VPN tunnel you just created between two Google Cloud VPCs, 10.1.0.0/16 and 172.16.0.0/16. You have a Cloud Router (router-1) in the 10.1.0.0/16 network and a second Cloud Router (router-2) in the 172.16.0.0/16 network. Which configuration should you use for the BGP session?

A. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	if-tunnel-a-to-b-if-0	169.254.0.254	169.254.0.254	65502
router-2	if-tunnel-b-to-a-if-0	169.254.0.254	169.254.0.254	65501

B. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	if-tunnel-a-to-b-if-0	10.1.0.1	172.16.0.1	15052
router-2	if-tunnel-b-to-a-if-0	172.16.0.1	10.1.0.1	15501

C. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	if-tunnel-a-to-b-if-0	169.254.20.1	169.254.20.2	65002
router-2	if-tunnel-b-to-a-if-0	169.254.20.2	169.254.20.1	65001

D. C:\Users\Admin\Desktop\Data\Odt data\Untitled.jpg

Router	BGP Interface Name	BGP IP	BGP Peer IP	Peer ASN
router-1	if-tunnel-a-to-b-if-0	172.16.0.254	10.1.0.254	16552
router-2	if-tunnel-b-to-a-if-0	10.1.0.254	172.16.0.254	16551

Answer: C

NEW QUESTION 36

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.

Answer: B

NEW QUESTION 37

You are responsible for designing a new connectivity solution for your organization's enterprise network to access and use Google Workspace. You have an existing Shared VPC with Compute Engine instances in us-west1. Currently, you access Google Workspace via your service provider's internet access. You want to set up a direct connection between your network and Google. What should you do?

- A. Order a Dedicated Interconnect connection in the same metropolitan area
- B. Create a VLAN attachment, a Cloud Router in us-west1, and a Border Gateway Protocol (BGP) session between your Cloud Router and your router.
- C. Order a Direct Peering connection in the same metropolitan area
- D. Configure a Border Gateway Protocol (BGP) session between Google and your router.
- E. Configure HA VPN in us-west1. Configure a Border Gateway Protocol (BGP) session between your Cloud Router and your on-premises data center.
- F. Order a Carrier Peering connection in the same metropolitan area
- G. Configure a Border Gateway Protocol (BGP) session between Google and your router.

Answer: B

NEW QUESTION 42

You are configuring your Google Cloud environment to connect to your on-premises network. Your configuration must be able to reach Cloud Storage APIs and your Google Kubernetes Engine nodes across your private Cloud Interconnect network. You have already configured a Cloud Router with your Interconnect VLAN attachments. You now need to set up the appropriate router advertisement configuration on the Cloud Router. What should you do?

- A. Configure the route advertisement to the default setting.
- B. On the on-premises router, configure a static route for the storage API virtual IP address which points to the Cloud Router's link-local IP address.
- C. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisement
- D. Leave all other options as their default settings.
- E. Configure the route advertisement to the custom setting, and manually add prefix 199.36.153.8/30 to the list of advertisement
- F. Advertise all visible subnets to the Cloud Router.

Answer: C

NEW QUESTION 44

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 subnet
- B. Create 2 VPN gateways to establish connectivity between these regions.
- C. Create 2 VPCs, each with their own region and individual subnet
- D. Use external IP addresses on the instances to establish connectivity between these regions.
- E. Create 1 VPC with 2 regional subnet
- F. Create a global load balancer to establish connectivity between the regions.
- G. Create 1 VPC with 2 regional subnet
- H. Deploy workloads in these subnets and have them communicate using private RFC1918 IP addresses.

Answer: D

Explanation:

<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.

NEW QUESTION 49

You are responsible for enabling Private Google Access for the virtual machine (VM) instances in your Virtual Private Cloud (VPC) to access Google APIs. All VM instances have only a private IP address and need to access Cloud Storage. You need to ensure that all VM traffic is routed back to your on-premises data center for traffic scrubbing via your existing Cloud Interconnect connection. However, VM traffic to Google APIs should remain in the VPC. What should you do?

- A. Delete the default route in your VPC. Create a private Cloud DNS zone for googleapis.com, create a CNAME for *.googleapis.com to restricted googleapis.com, and create an A record for restricted googleapis.com that resolves to the addresses in 199.36.153.4/30. Create a static route in your VPC for the range 199.36.153.4/30 with the default internet gateway as the next hop.
- B. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP). Create a public Cloud DNS zone with a CNAME for *.google.com to private googleapis.com, create a CNAME for * googleapis.com to private googleapis.com, and create an A record for Private googleapis.com that resolves to the addresses in 199.36.153.8/30. Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.
- C. Configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP) with a lower priority (MED) than the default VPC route. Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googleapis.com to private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30. Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.
- D. Delete the default route in your VPC and configure your on-premises router to advertise 0.0.0.0/0 via Border Gateway Protocol (BGP). Create a private Cloud DNS zone for googleapis.com, create a CNAME for * googleapis.com to Private googleapis.com, and create an A record for private.googleapis.com that resolves to the addresses in 199.36.153.8/30. Create a static route in your VPC for the range 199.36.153.8/30 with the default internet gateway as the next hop.

Answer: C

NEW QUESTION 51

You have the following private Google Kubernetes Engine (GKE) cluster deployment:

```
gcloud container clusters describe customer-1-cluster --zone us-central1-c

...

clusterIpv4Cidr: 192.168.36.0/24
endpoint: 192.168.38.2
ipAllocationPolicy:
  clusterIpv4Cidr: 192.168.36.0/24
  clusterIpv4CidrBlock: 192.168.36.0/24
  clusterSecondaryRangeName: customer-1-pods
  servicesIpv4Cidr: 192.168.37.0/24
  servicesIp4CidrBlock: 192.168.37.0/24
  servicesSecondaryRangeName: customer-1-svc
  useIpAliases: true

...

masterAuthorizedNetworksConfig:

...

privateClusterConfig:
  enablePrivateEndpoint: true
  enablePrivateNodes: true
  masterIpv4CidrBlock: 192.168.38.0/28
  privateEndpoint: 192.168.38.2
  publicEndpoint: 35.224.37.17

...

servicesIpv4Cidr: 192.162.37.0/24

...

subnetwork: customer-1-nodes
zone: us-central1-c
```

You have a virtual machine (VM) deployed in the same VPC in the subnetwork kubernetes-management with internal IP address 192.168.40 2/24 and no external IP address assigned. You need to communicate with the cluster master using kubectl. What should you do?

- A. Add the network 192.168.40.0/24 to the masterAuthorizedNetworksConfi
- B. Configure kubectl to communicate with the endpoint 192.168.38.2.
- C. Add the network 192.168.38.0/28 to the masterAuthorizedNetworksConfi
- D. Configure kubectl to communicate with the endpoint 192.168.38.2
- E. Add the network 192.168.36.0/24 to the masterAuthorizedNetworksConfi
- F. Configure kubectl to communicate with the endpoint 192.168.38.2
- G. Add an external IP address to the VM, and add this IP address in the masterAuthorizedNetworksConfig. Configure kubectl to communicate with the endpoint 35.224.37.17.

Answer: A

NEW QUESTION 56

After a network change window one of your company's applications stops working. The application uses an on-premises database server that no longer receives any traffic from the application. The database server IP address is 10.2.1.25. You examine the change request, and the only change is that 3 additional VPC subnets were created. The new VPC subnets created are 10.1.0.0/16, 10.2.0.0/16, and 10.3.1.0/24/ The on-premises router is advertising 10.0.0.0/8. What is the most likely cause of this problem?

- A. The less specific VPC subnet route is taking priority.
- B. The more specific VPC subnet route is taking priority.
- C. The on-premises router is not advertising a route for the database server.
- D. A cloud firewall rule that blocks traffic to the on-premises database server was created during the change.

Answer: B

NEW QUESTION 61

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: B

Explanation:

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

NEW QUESTION 66

You deployed a hub-and-spoke architecture in your Google Cloud environment that uses VPC Network Peering to connect the spokes to the hub. For security reasons, you deployed a private Google Kubernetes Engine (GKE) cluster in one of the spoke projects with a private endpoint for the control plane. You configured authorized networks to be the subnet range where the GKE nodes are deployed. When you attempt to reach the GKE control plane from a different spoke project, you cannot access it. You need to allow access to the GKE control plane from the other spoke projects. What should you do?

- A. Add a firewall rule that allows port 443 from the other spoke projects.
- B. Enable Private Google Access on the subnet where the GKE nodes are deployed.
- C. Configure the authorized networks to be the subnet ranges of the other spoke projects.
- D. Deploy a proxy in the spoke project where the GKE nodes are deployed and connect to the control plane through the proxy.

Answer: C

NEW QUESTION 69

You recently deployed Cloud VPN to connect your on-premises data center to Google Cloud. You need to monitor the usage of this VPN and set up alerts in case traffic exceeds the maximum allowed. You need to be able to quickly decide whether to add extra links or move to a Dedicated Interconnect. What should you do?

- A. In the Network Intelligence Center, check for the number of packet drops on the VPN.
- B. In the Google Cloud Console, use Monitoring Query Language to create a custom alert for bandwidth utilization.
- C. In the Monitoring section of the Google Cloud Console, use the Dashboard section to select a default dashboard for VPN usage.
- D. In the VPN section of the Google Cloud Console, select the VPN under hybrid connectivity, and then select monitoring to display utilization on the dashboard.

Answer: A

NEW QUESTION 74

You want to apply a new Cloud Armor policy to an application that is deployed in Google Kubernetes Engine (GKE). You want to find out which target to use for your Cloud Armor policy.

Which GKE resource should you use?

- A. GKE Node
- B. GKE Pod
- C. GKE Cluster
- D. GKE Ingress

Answer: D

Explanation:

Cloud Armour is applied at load balancers Configuring Google Cloud Armor through Ingress. <https://cloud.google.com/kubernetes-engine/docs/how-to/ingress-features> Security policy features Google Cloud Armor security policies have the following core features: You can optionally use the QUIC protocol with load balancers that use Google Cloud Armor. You can use Google Cloud Armor with external HTTP(S) load balancers that are in either Premium Tier or Standard Tier. You can use security policies with GKE and the default Ingress controller.

NEW QUESTION 78

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.

Answer: C

Explanation:

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

NEW QUESTION 81

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?

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

Answer: A

NEW QUESTION 85

You configured Cloud VPN with dynamic routing via Border Gateway Protocol (BGP). You added a custom route to advertise a network that is reachable over the VPN tunnel. However, the on-premises clients still cannot reach the network over the VPN tunnel. You need to examine the logs in Cloud Logging to confirm that the appropriate routers are being advertised over the VPN tunnel. Which filter should you use in Cloud Logging to examine the logs?

- A. resource.type= "gce_router"
- B. resource.type= "gce_network_region"
- C. resource.type= "vpn_tunnel"
- D. resource.type= "vpn_gateway"

Answer: C

NEW QUESTION 89

You are planning a large application deployment in Google Cloud that includes on-premises connectivity. The application requires direct connectivity between workloads in all regions and on-premises locations without address translation, but all RFC 1918 ranges are already in use in the on-premises locations. What should you do?

- A. Use multiple VPC networks with a transit network using VPC Network Peering.
- B. Use overlapping RFC 1918 ranges with multiple isolated VPC networks.
- C. Use overlapping RFC 1918 ranges with multiple isolated VPC networks and Cloud NAT.
- D. Use non-RFC 1918 ranges with a single global VPC.

Answer: D

NEW QUESTION 90

You are using the gcloud command line tool to create a new custom role in a project by copying a predefined role. You receive this error message: INVALID_ARGUMENT: Permission resourcemanager.projects.list is not valid What should you do?

- A. Add the resourcemanager.projects.get permission, and try again.
- B. Try again with a different role with a new name but the same permissions.
- C. Remove the resourcemanager.projects.list permission, and try again.
- D. Add the resourcemanager.projects.setIamPolicy permission, and try again.

Answer: C

NEW QUESTION 92

You want to implement an IPsec tunnel between your on-premises network and a VPC via Cloud VPN. You need to restrict reachability over the tunnel to specific local subnets, and you do not have a device capable of speaking Border Gateway Protocol (BGP). Which routing option should you choose?

- A. Dynamic routing using Cloud Router
- B. Route-based routing using default traffic selectors
- C. Policy-based routing using a custom local traffic selector
- D. Policy-based routing using the default local traffic selector

Answer: C

NEW QUESTION 97

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 identify the web application firewall (WAF) rule that is incorrectly blocking traffic. What should you do?

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

Answer: A

NEW QUESTION 98

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?

- A. The default internet gateway
- B. The IP address of the Cloud VPN gateway
- C. The name and region of the Cloud VPN tunnel
- D. The IP address of the instance on the remote side of the VPN tunnel

Answer: C

Explanation:

When you create a route based tunnel using the Cloud Console, Classic VPN performs both of the following tasks: Sets the tunnel's local and remote traffic selectors to any IP address (0.0.0.0/0) For each range in Remote network IP ranges, Google Cloud creates a custom static route whose destination (prefix) is the range's CIDR, and whose next hop is the tunnel.

<https://cloud.google.com/network-connectivity/docs/vpn/how-to/creating-static-vpns>

NEW QUESTION 102

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?

- A. Use the default public domains for all Google APIs and services.
- B. Use Private Service Connect to access Cloud Storage, and use the default public domains for all other Google APIs and services.
- C. 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.
- D. 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.

Answer: B

NEW QUESTION 106

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?

- A. Use Shared VPC, and deploy the VLAN attachments and Interconnect in the host project.
- B. Use Shared VPC, and deploy the VLAN attachments in the service project
- C. Connect the VLAN attachment to the Shared VPC's host project.
- D. Use standalone projects, and deploy the VLAN attachments in the individual project
- E. Connect the VLAN attachment to the standalone projects' Interconnects.
- F. Use standalone projects and deploy the VLAN attachments and Interconnects in each of the individual projects.

Answer: A

Explanation:

<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-sa>

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

NEW QUESTION 109

Your company has separate Virtual Private Cloud (VPC) networks in a single region for two departments: Sales and Finance. The Sales department's VPC network already has connectivity to on-premises locations using HA VPN, and you have confirmed that the subnet ranges do not overlap. You plan to peer both VPC networks to use the same HA tunnels for on-premises connectivity, while providing internet connectivity for the Google Cloud workloads through Cloud NAT. Internet access from the on-premises locations should not flow through Google Cloud. You need to propagate all routes between the Finance department and on-premises locations. What should you do?

- A. Peer the two VPCs, and use the default configuration for the Cloud Routers.
- B. Peer the two VPCs, and use Cloud Router's custom route advertisements to announce the peered VPC network ranges to the on-premises locations.
- C. Peer the two VPC
- D. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network
- E. Use Cloud Router's custom route advertisements to announce a default route to the on-premises locations.
- F. Peer the two VPC
- G. Configure VPC Network Peering to export custom routes from Sales and import custom routes on Finance's VPC network
- H. Use Cloud Router's custom route advertisements to announce the peered VPC network ranges to the on-premises locations.

Answer: A

NEW QUESTION 113

Your software team is developing an on-premises web application that requires direct connectivity to Compute Engine Instances in GCP using the RFC 1918 address space. You want to choose a connectivity solution from your on-premises environment to GCP, given these specifications:

- Your ISP is a Google Partner Interconnect provider.
- Your on-premises VPN device's internet uplink and downlink speeds are 10 Gbps.
- A test VPN connection between your on-premises gateway and GCP is performing at a maximum speed of 500 Mbps due to packet losses.
- Most of the data transfer will be from GCP to the on-premises environment.
- The application can burst up to 1.5 Gbps during peak transfers over the Interconnect.
- Cost and the complexity of the solution should be minimal.

How should you provision the connectivity solution?

- A. Provision a Partner Interconnect through your ISP.
- B. Provision a Dedicated Interconnect instead of a VPN.
- C. Create multiple VPN tunnels to account for the packet losses, and increase bandwidth using ECMP.
- D. Use network compression over your VPN to increase the amount of data you can send over your VPN.

Answer: A

Explanation:

Direct Interconnect will be too expensive and also an overkill for this requirement. Managing multiple tunnels that too with packet loss consideration is complex also. Whereas partner interconnect fits the bill with providing required bandwidth but not super expensive also once setup not too complex too manage.

NEW QUESTION 117

In your project my-project, you have two subnets in a Virtual Private Cloud (VPC): subnet-a with IP range 10.128.0.0/20 and subnet-b with IP range 172.16.0.0/24. You need to deploy database servers in subnet-a. You will also deploy the application servers and web servers in subnet-b. You want to configure firewall rules that only allow database traffic from the application servers to the database servers. What should you do?

- A. Create network tag app-server and service account sa-db@my-project.iam.gserviceaccount.co
- B. Add the tag to the application servers, and associate the service account with the database server
- C. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules top:3306 --source-tags app-server --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- D. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.co
- E. Associate service account sa-app with the application servers, and associate the service account sa-db with the database server
- F. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-service-accounts sa-app@democloud-idp-demo.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- G. Create service accounts sa-app@my-project.iam.gserviceaccount.com and sa-db@my-project.iam.gserviceaccount.co
- H. Associate the service account sa-app with the application servers, and associate the service account sa-db with the database server
- I. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --allow TCP:3306 --source-ranges 10.128.0.0/20 --source-service-accounts sa-app@my-project.iam.gserviceaccount.com --target-service-accounts sa-db@my-project.iam.gserviceaccount.com`
- J. Create network tags app-server and db-server
- K. Add the app-server tag to the application servers, and add the db-server tag to the database server
- L. Run the following command: `gcloud compute firewall-rules create app-db-firewall-rule --action allow --direction ingress --rules tcp:3306 --source-ranges 10.128.0.0/20 --source-tags app-server --target-tags db-server`

Answer: D

NEW QUESTION 119

You have two Google Cloud projects in a perimeter to prevent data exfiltration. You need to move a third project inside the perimeter; however, the move could negatively impact the existing environment. You need to validate the impact of the change. What should you do?

- A. Enable Firewall Rules Logging inside the third project.
- B. Modify the existing VPC Service Controls policy to include the new project in dry run mode.
- C. Monitor the Resource Manager audit logs inside the perimeter.
- D. Enable VPC Flow Logs inside the third project, and monitor the logs for negative impact.

Answer: B

NEW QUESTION 121

You have an application hosted on a Compute Engine virtual machine instance that cannot communicate with a resource outside of its subnet. When you review the flow and firewall logs, you do not see any denied traffic listed.

During troubleshooting you find:

- Flow logs are enabled for the VPC subnet, and all firewall rules are set to log.
- The subnetwork logs are not excluded from Stackdriver.
- The instance that is hosting the application can communicate outside the subnet.
- Other instances within the subnet can communicate outside the subnet.
- The external resource initiates communication. What is the most likely cause of the missing log lines?

- A. The traffic is matching the expected ingress rule.
- B. The traffic is matching the expected egress rule.
- C. The traffic is not matching the expected ingress rule.
- D. The traffic is not matching the expected egress rule.

Answer: C

NEW QUESTION 124

You are designing the network architecture for your organization. Your organization has three developer teams: Web, App, and Database. All of the developer teams require access to Compute Engine instances to perform their critical tasks. You are part of a small network and security team that needs to provide network access to the developers. You need to maintain centralized control over network resources, including subnets, routes, and firewalls. You want to minimize

operational overhead. How should you design this topology?

- A. Configure a host project with a Shared VP
- B. Create service projects for Web, App, and Database.
- C. Configure one VPC for Web, one VPC for App, and one VPC for Databas
- D. Configure HA VPN between each VPC.
- E. Configure three Shared VPC host projects, each with a service project: one for Web, one for App, and one for Database.
- F. Configure one VPC for Web, one VPC for App, and one VPC for Databas
- G. Use VPC Network Peering to connect all VPCs in a full mesh.

Answer: C

NEW QUESTION 129

You are developing an HTTP API hosted on a Compute Engine virtual machine instance that must be invoked only by multiple clients within the same Virtual Private Cloud (VPC). You want clients to be able to get the IP address of the service. What should you do?

- A. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwarding rul
- B. Clients should use this IP address to connect to the service.
- C. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url `https://[INSTANCE_NAME].[ZONE].c.[PROJECT_ID].internal/`.
- D. Reserve a static external IP address and assign it to an HTTP(S) load balancing service's forwardingrul
- E. Then, define an A record in Cloud DN
- F. Clients should use the name of the A record to connect to the service.
- G. Ensure that clients use Compute Engine internal DNS by connecting to the instance name with the url `https://[API_NAME]/[API_VERSION]/`.

Answer: C

NEW QUESTION 132

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?

- A. Configure global load balancing to point 172.16.45.0/24 to the correct instance.
- B. Create unique DNS records for each service that sends traffic to the desired IP address.
- C. 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.
- D. Use VPC peering to allow traffic to route between the 10.1.0.0/24 network and the 172.16.45.0/24network.

Answer: C

NEW QUESTION 137

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 (
- D. g., code-dev), and make the second project (
- E. g., data-dev) a service project.
- F. 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.
- G. 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: BD

NEW QUESTION 138

You have configured a Compute Engine virtual machine instance as a NAT gateway. You execute the following command:

```
gcloud compute routes create no-ip-internet-route \
--network custom-network1 \
--destination-range 0.0.0.0/0 \
--next-hop instance nat-gateway \
--next-hop instance-zone us-central1-a \
--tags no-ip --priority 800
```

You want existing instances to use the new NAT gateway. Which command should you execute?

- A. `sudo sysctl -w net.ipv4.ip_forward=1`
- B. `gcloud compute instances add-tags [existing-instance] --tags no-ip`
- C. `gcloud builds submit --config=cloudbuild.waml --substitutions=TAG_NAME=no-ip`
- D. `gcloud compute instances create example-instance --network custom-network1 --subnet subnet-us-central --no-address --zone us-central1-a --image-family debian-9 --image-project debian-cloud --tags no-ip`

Answer: B

Explanation:

<https://cloud.google.com/sdk/gcloud/reference/compute/routes/create>

In order to apply a route to an existing instance we should use a tag to bind the route to it.

NEW QUESTION 139

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.)

- A. Turn on Private Google Access at the subnet level.
- B. Turn on Private Google Access at the VPC level.
- C. Turn on Private Services Access at the VPC level.
- D. 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.
- E. 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.

Answer: AD

Explanation:

<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.

NEW QUESTION 142

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

Professional-Cloud-Network-Engineer Practice Exam Features:

- * Professional-Cloud-Network-Engineer Questions and Answers Updated Frequently
- * Professional-Cloud-Network-Engineer Practice Questions Verified by Expert Senior Certified Staff
- * Professional-Cloud-Network-Engineer Most Realistic Questions that Guarantee you a Pass on Your First Try
- * Professional-Cloud-Network-Engineer Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The Professional-Cloud-Network-Engineer Practice Test Here](#)